

# DISSERTATION

zur Erlangung des akademischen Grades  
Doktor der Naturwissenschaften (Dr. rer. nat.)

## **Investigating neural replay of task representations in the human brain using fMRI**

eingereicht am Fachbereich Erziehungswissenschaft und  
Psychologie der Freien Universität Berlin  
von **John Lennart Wittkuhn, M.Sc.**

Berlin, 2022



Erstgutachter/-in: Dr. Nicolas W. Schuck

Zweitgutachter/-in: Prof. Dr. Radoslaw M. Cichy

Drittgutachter/-in: Prof. Dr. Monika Schönauer

Datum der Disputation: 11.04.2022

## ACKNOWLEDGMENTS

This dissertation was conducted within the Max Planck Research Group “Neural and Computational Basis of Learning, Decision Making and Memory (NeuroCode)”, led by Dr. Nicolas W. Schuck at the Max Planck Institute for Human Development in Berlin, Germany. I was also a doctoral fellow of the International Max Planck Research School on Computational Methods in Psychiatry and Ageing Research (IMPRS COMP2PSYCH). All research was funded by an Independent Max Planck Research Group grant by the Max Planck Society and a Starting Grant by the European Research Council, both awarded to Dr. Nicolas W. Schuck. The work presented in this dissertation was also supported by travel grants by the German Academic Exchange Service (DAAD) and the IMPRS COMP2PSYCH. I am obliged to these sources of funding which made this dissertation possible.

I am very grateful to all the people and institutions that have supported me and my work in the last 4+ years. First and foremost, I want to express my gratitude to my advisor Nico Schuck. Nico, I am deeply grateful for your invaluable commitment in mentoring my dissertation. Without your reliable guidance, far-sighted supervision, and intellectual inspiration, I would not be submitting this thesis today. I am grateful to Prof. Dr. Radoslaw Cichy for evaluating my dissertation, Prof. Dr. Peter Mohr and Prof. Dr. Bernhard Staresina for kindly agreeing to join my doctoral committee, as well as Prof. Dr. Hauke Heekeren for supporting my dissertation proposal at Freie Universität Berlin.

I feel extremely fortunate that I have been part of an amazing group of people in the NeuroCode lab that have supported me in countless ways. Anika, for her exceptional memory, eye for detail and oat bread. Christoph, for sharing the PhD rollercoaster ride and German dialects. Eva, for saving me from an oversupply of admin work and the office plants from an undersupply of water. Gregor, for filling every single MRI slot with motivated participants and chocolate on his office couch. Lena, for carrying a two-session fMRI project through a global pandemic with me, and admitting to the usefulness of Git issues. Nir, for being my favorite devil’s advocate and obsessing about details of study design with me. Ondrej, for sharing the excitement about tools that make science more efficient and reproducible and keeping up the motivation to push for change. Sam C., for staying calm, and telling me to “go for it”. Sam H.M., for advice on writing and impressing me with constructive feedback. Thanks to Fabian, Marit, Moritz, Luianta, Noa, and Shany for immediately joining into this friendly and supportive community. I thoroughly enjoyed working with you all over the past years.

Beyond the NeuroCode lab, I thank the Max Planck Institute for Human Development and the Max Planck Society for providing an outstanding research infrastructure. Special thanks to Michael Krause, for near instant solutions to virtually every technical issue. The COMP2PSYCH faculty and fellows, for enriching discussions and “informal consultations” that let us climb to new academic heights. Simon, for reminding me to not overthink too much and bike rides. Rasmus, for advice on navigating the academic path and completing my doctoral committee. Adina, for teaching me DataLad, which revolutionized my workflow and kept me sane when working on two fMRI projects at the same time. The ReproNim faculty and fellows, for a community where openness and reproducibility go without saying.

Finally, I am most grateful to my family and friends. My parents and brother, for always supporting me. Phine, for her endless patience and encouragement. And Jale, for bringing so much joy.

## SUMMARY

A remarkable discovery made in rodents three decades ago demonstrated that the brain reactivates previous experience during sleep and wakeful rest. This phenomenon, since known as *replay*, has been implicated in a variety of cognitive functions, ranging from spatial navigation and episodic memory consolidation to planning and decision-making. At the same time, research in machine learning (ML) has found that experience replay can substantially improve the performance of artificial agents. Together, these observations have spawned the idea that replay supports behavior by retrieving information from abstract internal representations of the environment. How the interplay between replay and internal task representations supports learning and decision-making in humans is still not well understood. In addition, investigating replay in humans is challenging because replay is fast, sequential, and occurs throughout the brain, but current non-invasive neuroimaging methods offer either sufficient temporal or spatial resolution, but not both.

This dissertation consists of three publications that offer a theoretical and empirical perspective on the role of replay in learning and decision-making as well as methodological advances for the study of replay in humans using functional magnetic resonance imaging (fMRI).

In *Paper I*, we provide a review of the recent computational and neuroscientific literature on replay and elucidate how replay can improve learning and decision-making in both biological and artificial agents. We identified five key computational functions of replay which include faster and data efficient learning, less forgetting, the reorganization and augmentation of experiences, planning, and generalization. We also discussed the benefits of reactivating abstract internal representations instead of recapitulating veridical experiences, and explored the idea that replay could be involved in learning which representation is useful for a given task.

In *Paper II*, we developed and experimentally validated multivariate analysis methods for fMRI that allow studying the sequentiality and speed of fast-paced neural event sequences, like replay, with anatomical precision in humans. The results showed that probabilistic fMRI pattern classifiers make it possible to detect the sequential order of neural image representations at speeds of up to 32 milliseconds between sequence items. Finally, applying these methods to fMRI data from awake resting-state scans, we could differentiate fast from slow neural sequences that occurred at random time points and found evidence for sequential replay of task-related stimulus sequences during post-task rest.

In *Paper III*, we investigated how humans form predictive internal representations of the statistical relationships between task stimuli and if the brain reactivates sequences from these internal representations during short pauses from ongoing behavior. The results showed that humans learned the higher-order relationships between consecutively presented images in the form of a predictive cognitive map that represents each item in terms of the subsequent items that follow from it. Applying the fMRI methods developed in *Paper II* to neural data during short pauses interleaved with task performance, we found that participants reactivated upcoming stimulus sequences that were most likely given this predictive internal representation.

In summary, this dissertation makes three major contributions. First, linking findings about replay in the neuroscience and ML literature spurs new ideas how replay can support learning and decision-making in both biological and artificial agents. Second, novel methods to characterize the speed and sequentiality of replay with spatial specificity using fMRI have the potential to foster future insights into the role of replay in the human brain. Third, studying how replay interacts with internal representations opens avenues to further understand how the reactivation of experience supports adaptive behavior in machines and humans.

## ZUSAMMENFASSUNG

Eine bemerkenswerte Entdeckung, die vor drei Jahrzehnten in Ratten gemacht wurde, zeigte, dass das Gehirn vorherige Erfahrungen während des Schlafs oder wachen Ruhepausen reaktiviert. Dieses Phänomen, seither bekannt als *Replay* (engl. für *Wiedergabe* oder *Reaktivierung*), wurde mit einer Vielfalt von kognitiven Funktionen in Verbindung gebracht, die von räumlicher Navigation und episodischer Gedächtniskonsolidierung bis zu Planungsfähigkeiten und Entscheidungsfindung reichen. Zeitgleich wurde in Forschung zu maschinellem Lernen erkannt, dass Replay von Erfahrungen zu substantiellen Verbesserungen in der Performanz von künstlichen Akteuren (z.B. Robotern) führt. Zusammengefasst haben diese Beobachtungen die Idee hervorgebracht, dass Replay Verhalten unterstützt, indem Informationen einer abstrakten internen Repräsentation der Umgebung reaktiviert werden. Wie genau das Zusammenspiel von Replay und internen Repräsentationen der Aufgabe Lern- und Entscheidungsprozesse unterstützt, wurde noch nicht ausreichend verstanden. AuSSerdem ist die Erforschung von Replay in Menschen herausfordernd, weil Replay schnell und sequentiell ist und überall im Gehirn vorkommen kann. Jedoch bieten nicht-invasive neurowissenschaftliche Bildgebungsverfahren entweder ausreichende zeitliche oder räumliche Auflösung, aber nicht beides.

Diese Dissertation besteht aus drei Publikationen, die eine theoretische und empirische Perspektive auf die Rolle von Replay in Lern- und Entscheidungsprozessen, sowie methodische Fortschritte für die Erforschung von Replay in Menschen mittels funktioneller Magnetresonanztomografie (fMRT) bieten.

In *Artikel I* bieten wir eine Übersicht über die aktuelle komputationale und neurowissenschaftliche Literatur zu Replay und erläutern, wie Replay Lern- und Entscheidungsprozesse sowohl in Menschen als auch künstlichen Akteuren verbessern kann. Wir identifizieren fünf zentrale komputationale Funktionen von Replay, die es erlauben schneller und effizienter zu lernen, weniger zu vergessen, Erfahrungen neu zu organisieren oder zu erweitern, zu planen und Erfahrungen zu generalisieren. Wir diskutieren auch, welche Vorteile es hat, abstrakte interne Repräsentationen zu reaktivieren anstatt Erfahrungen wahrheitsgetreu zu rekapitulieren und untersuchen die Idee, dass Replay dazu beitragen könnte, zu lernen, welche Repräsentation für eine bestimmte Aufgabe hilfreich ist.

In *Artikel II* haben wir multivariate Analysemethoden für fMRT entwickelt und experimentell validiert, die es ermöglichen, die Sequentialität und Geschwindigkeit von schnellen neuronalen Sequenzen, wie Replay, mit anatomischer Präzision in Menschen zu messen. Unsere Ergebnisse zeigen, dass probabilistische Algorithmen zur Klassifizierung von fMRT Aktivierungsmustern es ermöglichen, die sequentielle Reihenfolge von neuronalen Repräsentationen von Bildern auch bei Geschwindigkeiten von nur 32 Millisekunden zwischen einzelnen Elementen der Sequenz zu erkennen. AuSSerdem konnten wir zwischen schnellen und langsamen neuronalen Sequenzen unterscheiden, die an zufälligen Zeitpunkten während einer verlängerten Ruhephase aufgetreten sind. Indem wir diese Methoden auf fMRT Daten während einer wachen Ruhephase nach der Aufgabe anwendeten, konnten wir Evidenz für sequentielles Replay von den Bildern der vorherigen Aufgabe zeigen.

In *Artikel III* untersuchten wir, wie Menschen prädiktive interne Repräsentationen von statischen Zusammenhänge zwischen Bildern lernen und ob das Gehirn Sequenzen von dieser internen Repräsentation in kurzen Pausen während der Aufgabe reaktiviert. Unsere Ergebnisse zeigten, dass Menschen übergeordnete Beziehungen zwischen aufeinanderfolgenden Bildern in Form einer prädiktiven kognitiven Karte lernen, die jedes Bild durch die darauffolgenden Bilder repräsentiert. Indem wir die fMRT Methoden aus *Artikel II* auf neuronale Daten während kurzen Pausen von der laufenden Aufgabe anwendeten, entdeckten wir, dass die Probanden

bevorstehende Sequenzen von Bildern reaktivierten, die entsprechend der prädiktiven internen Repräsentation am wahrscheinlichsten waren.

Zusammenfassend leistet diese Dissertation drei zentrale Beiträge. Erstens wird gezeigt, dass die Verknüpfung von Erkenntnissen zu Replay aus Neurowissenschaft und maschinellem Lernen neue Ideen anregt, wie Replay Lern- und Entscheidungsprozesse in biologischen und künstlichen Akteuren unterstützt. Zweitens haben die neuen Methoden zur Erforschung der Sequentialität und Geschwindigkeit von Replay mittels räumlich spezifischer fMRT das Potenzial, zukünftige Erkenntnisse zur Rolle von Replay in Menschen zu fördern. Drittens eröffnet die Erforschung der Interaktion zwischen Replay und internen Repräsentationen neue Möglichkeiten zu verstehen, wie die Reaktivierung von Erfahrungen adaptives Verhalten von Maschinen und Menschen ermöglicht.

## LIST OF PAPERS

This doctoral thesis is based on the following three original papers:

### *Paper I*

**Wittkuhn, L.**, Chien, S., Hall-McMaster, S., & Schuck, N. W. (2021). Replay in minds and machines. *Neuroscience & Biobehavioral Reviews*, *129*, 367–388. doi: [10.1016/j.neubiorev.2021.08.002](https://doi.org/10.1016/j.neubiorev.2021.08.002)

### *Paper II*

**Wittkuhn, L.**, & Schuck, N. W. (2021). Dynamics of fMRI patterns reflect sub-second activation sequences and reveal replay in human visual cortex. *Nature Communications*, *12*(1795). doi: [10.1038/s41467-021-21970-2](https://doi.org/10.1038/s41467-021-21970-2)

### *Paper III*

**Wittkuhn, L.**, Krippner, L. M., & Schuck, N. W. (2022). Statistical learning of successor representations is related to on-task replay. *bioRxiv*. doi: [10.1101/2022.02.02.478787](https://doi.org/10.1101/2022.02.02.478787)



## CONTENTS

<b>Acknowledgments</b>	<b>iv</b>
<b>Summary</b>	<b>v</b>
<b>Zusammenfassung</b>	<b>vi</b>
<b>List of papers</b>	<b>viii</b>
<b>Contents</b>	<b>1</b>
<b>1 Introduction</b>	<b>2</b>
1.1 The discovery of replay . . . . .	3
1.2 The diversity of replay . . . . .	4
1.3 Key characteristics of replay . . . . .	5
1.4 Replay in humans . . . . .	8
1.5 Replay in artificial agents . . . . .	10
1.6 The content of replay . . . . .	11
<b>2 Research questions</b>	<b>13</b>
<b>3 Overview of papers</b>	<b>15</b>
3.1 Paper I . . . . .	15
3.2 Paper II . . . . .	19
3.3 Paper III . . . . .	21
<b>4 Discussion</b>	<b>24</b>
4.1 Summary and evaluation . . . . .	24
4.2 Limitations and open questions . . . . .	28
4.3 Conclusion . . . . .	35
<b>References</b>	<b>36</b>
<b>Appendices</b>	<b>50</b>
A Declaration of own share . . . . .	51
B Declaration of independent work . . . . .	52
C Paper I . . . . .	53
D Paper II . . . . .	76
E Paper III . . . . .	115

# 1

## INTRODUCTION

Future actions are shaped by past experience. How the brain reuses previously learned information to enable flexible behavior is a central question at the intersection of memory, learning, and decision-making research (Gershman & Daw, 2017). To guide behavior, the brain requires a mechanism to form internal representations of previous experience and later retrieve this stored information to inform action. It is well established that the hippocampus plays a foundational role in memory functioning and mediates the encoding, storage, and retrieval of previous experience (Squire, 1992; Winocur & Moscovitch, 2011). In humans, the hippocampus and surrounding brain areas in the medial temporal lobe (MTL) are thought to support episodic and semantic memory (Scoville & Milner, 1957; Eichenbaum et al., 1999; Tulving, 2002) as well as planning and imagination (Buckner, 2010). In rodents, research on the same brain structures has largely focused on their role in spatial navigation (O’Keefe & Nadel, 1978; Hafting et al., 2005).

At the beginning of the 70s, O’Keefe and Dostrovsky (1971) performed electrophysiological recordings in the hippocampus of rats exploring a spatial environment and discovered neurons that reliably increased their activity when the animal occupied a particular location in space. These spatially-selective cells were called *place cells*, and the spatial location in the environment that consistently activated a particular neuron was referred to as its *place field*. In the following decades, many more cell types in the MTL were identified to represent distinct properties of physical space and contribute to spatial memory and navigation (for reviews, see e.g., McNaughton et al., 2006; Moser et al., 2008). These neurons also include grid cells in the entorhinal cortex whose firing fields tessellate the environment into a regular hexagonal grid pattern and enable the animal to keep track of relative spatial position and orientation (Hafting et al., 2005), and head direction-sensitive cells which could supply the animal with an internal compass (Taube et al., 1990). Together, these cells in the hippocampal-entorhinal system were identified as the neural substrate of a *cognitive map* (Tolman, 1948; O’Keefe & Nadel, 1974, 1978). Current scientific debate revolves around the question if and how place cells also represent locations in non-spatial cognitive maps (for reviews, see e.g., Behrens et al., 2018; Bellmund et al., 2018; O’Keefe & Krupic, 2021) and thereby also play a role in the neural

organization of episodic memories and relational knowledge. In one example landmark study in rodents, Aronov et al. (2017) showed that place cells responded to a linear change in sound frequency in the same way as they typically respond during spatial navigation along a linear track. Together, hippocampal cells may not only construct maps of physical space but also track non-spatial relational information which supports the idea that spatial navigation and episodic memory are both supported by the same neural machinery.

## 1.1 The discovery of replay

The early observations of place cells established the view that they represent the *current* position of the animal in a spatial environment. It was therefore a striking discovery that the same place cells, that were active when the animal navigated throughout the environment, *reactivated* during subsequent sleep when the animal remained stationary (Pavlides & Wilson, 1989). Following studies established that pairs of place cells with overlapping place fields during behavior showed increased co-activity during later sleep (Wilson & McNaughton, 1994; Skaggs & McNaughton, 1996; Kudrimoti et al., 1999; Gerrard et al., 2001), and that sequences of place cells reactivated during sleep matched behavioral sequences during wakefulness (Nádasy et al., 1999; Louie & Wilson, 2001; Lee & Wilson, 2002). Strikingly, the duration of replay events on the order of 50 to 200 milliseconds (ms) reflected a considerable temporal compression relative to the timescale of behavior. These seminal findings (also reviewed in Redish, 1999; Sutherland & McNaughton, 2000) established a classical and prevailing view of replay: The same sequential neural patterns in the hippocampus that represent previous experience are reactivated during sleep at compressed speeds of a few hundred milliseconds. The interpretation of these early findings was strongly influenced by theoretical models of memory consolidation at the time, that was viewed as a two-stage process whereby labile memory traces of recent experiences are repeatedly reactivated during “offline” periods like sleep in order to consolidate the information in the hippocampus and transfer it to the neocortex for long-term storage (Marr, 1971; Buzsáki, 1989).

## 1.2 The diversity of replay

In the following decades, a wealth of studies has established that replay is a much more diverse phenomenon than initially thought and likely plays a much broader role in cognition. Initially discovered during sleep, replay was also found during wakefulness (Kudrimoti et al., 1999; O’Neill et al., 2006; Jackson et al., 2006) (for review, see Carr et al., 2011). Awake replay was also found in the reverse order of experience (Foster & Wilson, 2006; Diba & Buzsáki, 2007; Csicsvari et al., 2007; Karlsson & Frank, 2009; Gupta et al., 2010; Ambrose et al., 2016; Shin et al., 2019). The discovery of backward replay brought the far-reaching implication that replay not necessarily represents a faithful replication of previous behavioral sequences but has the power to abstract from the original experience and reactivate it in novel ways. This insight was further deepened by findings showing that replay also represented remote experiences from environments that the animal did not currently occupy (Karlsson & Frank, 2009) and preferentially reactivated sequences that the animal experienced infrequently (Gupta et al., 2010). Observations of forward replay during wakefulness that started at the current location of the animal and ended at potential goal locations spurred the idea that replay is also involved in decision-making, planning, and deliberation (Johnson & Redish, 2007; Pfeiffer & Foster, 2013; Singer et al., 2013). Adding further to this complexity, Dragoi and Tonegawa (2011, 2013b) reported replay of environments that the animal had not even experienced yet. This so-called *preplay* was documented in several subsequent studies (Ólafsdóttir et al., 2015; Grosmark & Buzsáki, 2016; Farooq et al., 2019) but also contested (Silva et al., 2015), fueling an ongoing debate on how pre-existing memory representations are integrated with novel learning experiences and how they give rise to replay (Eichenbaum, 2015; Dragoi, 2020). Across this diversity of findings in rodents, three key characteristics of replay have been consistently identified that have important methodological implications for the investigation of replay in humans. These characteristics are speed, sequentiality, and anatomical specificity.

## 1.3 Key characteristics of replay

### Replay, fast and slow

The large majority of studies in rodents report replay events that are temporally compressed relative to the timescale of the behavioral experience that they reflect. The average replay speed is often indicated to be between 8 and 10 meters (m) per second (s) which is roughly 20 times faster than the animal's typical movement speed (Nádasdy et al., 1999; Lee & Wilson, 2002; Davidson et al., 2009; Karlsson & Frank, 2009). Replay of larger environments is accomplished by chaining replay events of successive subsequences that still maintain a coherent speed (Davidson et al., 2009). The dominant interpretation of the speed of replay is that the fast co-firing of consecutive cell pairs is conducive for synaptic plasticity (Hebb, 1949). Repeated successive spiking of pre- and postsynaptic neurons within a time window of about 20 ms has been shown to induce synaptic plasticity (Bi & Poo, 1998; Buzsáki, 1989) which strengthens synaptic connections between co-activated neurons (King et al., 1999; Mizunuma et al., 2014). While replay is fast in most observations, some studies suggest that replay speed might actually be much more variable, changing with experience (Deng et al., 2020), and occurring at much slower behavioral timescales (Denovellis et al., 2021). These findings hint at a substantial variability of replay speeds, but their functional relevance, dependence on the pace of experience, and relationship to subjective experience still remain largely unexplored. Regarding the study of replay in humans, this highlights the need for neuroimaging methods which are able to differentiate between neural sequences at varying speeds, and are sensitive to neural dynamics at the order of tens of milliseconds in particular. In the section on human replay below, I will describe how the issue of replay speed challenges fMRI methodology and how we addressed this issue in *Paper II*.

### Replay is sequential

We experience the world in a continuous sequence. Processing the sequential order of events is at the core of both spatial navigation and episodic memory abilities (Foster & Wilson, 2007; Buzsáki & Llinás, 2017). The discovery of place cells has provided fundamental insights into how the brain could represent events as locations in cognitive maps across space and time

(O'Keefe, 1979; Eichenbaum, 2014). Several properties of replay highlight the importance of capturing its sequentiality.

First, as mentioned above, replay is described as forward or backward depending on whether replayed sequences occur in the same or reverse temporal order as behavioral sequences. Importantly, forward and backward replay are thought to differ in their function (Diba & Buzsáki, 2007; Davidson et al., 2009). Forward replay is most commonly found at the beginning of a run and encodes paths towards potential goal locations and could therefore reflect an evaluation of future behavioral trajectories (Diba & Buzsáki, 2007; Johnson & Redish, 2007). Backward replay, in contrast, is often found at reward locations where it depicts trajectories back to the starting location and has therefore been ascribed a role in retrospective evaluation and credit assignment (Foster & Wilson, 2006; Diba & Buzsáki, 2007; Ambrose et al., 2016; Liu et al., 2019). The direction of replay is influenced by various factors, including properties of the task environment like rewards and goal locations, the behavioral state of the animal (sleep, awake rest, or brief pauses from ongoing behavior), the task demands and the putative cognitive process (e.g., memory consolidation, planning, or decision-making) (for review, see e.g., Foster, 2017). In order to further elucidate the roles of forward and backward replay, methods are required that can distinguish the directionality of neural sequences.

Second, a precise characterization of the contents and sequential order of replay events is required to establish links between replayed and behavioral sequences, which might not always be easy to predict. For example, in Wu et al. (2017) rodents replayed sequences that represented trajectories into a shock zone that the animal subsequently avoided. Similarly, in Carey et al. (2019) rodents that were either thirsty or hungry quickly learned to flexibly shift their behavior towards water and food while replayed sequences depicted paths away from the desired outcome. More generally, the consequences of alterations in the sequential order of neural representations has been demonstrated by Robinson et al. (2020). These authors showed that stimulating place cells that had place fields in a rewarded zone, when rats were not actually present in that part of the track, caused the rats to display increased licking behavior, as if the rewarded area had been reached. In addition, rats would run past the rewarded area when the place cell of an earlier location was stimulated, as if the animals were earlier in their running sequence. These findings illustrate the intricate relationships between (re-)activated neural representations and behavior.

Third, certain psychiatric disorders might be characterized by impairments that are specific to the sequential ordering of replay events. For example, in human patients with schizophrenia, the ability to infer sequential relationships in task structure is compromised and related to reduced replay of relational memories (Nour et al., 2021). In order to better understand the role of replay in such disorders, neuroimaging techniques are needed that are not only sensitive to the contents but also the sequential order of replay events.

Together, these findings establish that the sequential order of sequence items is an important characteristic of replay. These insights also call for a precise definition of *reactivation* and *replay*. While *reactivation* describes evidence for non-sequential reinstatement of neural activity patterns, replay is defined as “a specific form of reactivation that includes *sequential* (temporal and/or spatial) information” (for a recent consensus statement, see Genzel et al., 2020). Measurement techniques must not only be able to detect the identity of reactivated experiences, but also have sensitivity for their sequential order. As described further below, capturing the sequentiality of replay is a major open challenge for human fMRI studies.

### Replay is anatomically localized and distributed

Most research on replay has focused on the hippocampus. However, replay not only exists in the hippocampus, but has been found in a wide range of brain areas, including visual cortex (Ji & Wilson, 2006), auditory cortex (Rothschild et al., 2016), prefrontal cortex (PFC) (Euston et al., 2007; Peyrache et al., 2009; Jadhav et al., 2016; Shin et al., 2019; Kaefer et al., 2020), motor cortex (Xu et al., 2019; Eckert et al., 2020), entorhinal cortex (Ólafsdóttir et al., 2016; O’Neill et al., 2017; Trettel et al., 2019), and ventral striatum (Lansink et al., 2008; Lansink et al., 2009). Critically, the degree of coordination (or independence) between anatomically distributed replay events might be linked to their function. For example, Ólafsdóttir et al. (2017) reported coordinated replay between the hippocampus and entorhinal cortex only during immobile “offline” rest periods (often linked to memory consolidation), but not during “online” periods before and after movements (often linked to planning and decision-making). In Kaefer et al. (2020), replay in the medial PFC was found to be temporally uncorrelated with hippocampal replay, suggesting independent replay-mediated computations in the medial PFC. The degree of coordination between replay in the hippocampus and PFC has also been shown to change over the course of learning and was predictive of behavioral trajectories (Shin

et al., 2019). Together, investigating simultaneous, anatomically distributed replay with spatial specificity promises important insights into how the reactivation of memory representations is organized across brain regions and linked to perception and behavior. Accordingly, Foster (2017, p. 598) concluded his review noting that “hippocampal replay offers a window on a degree of complexity that may exist throughout the cortex [...]. This in the end may be its most important contribution”. These findings about replay from work in non-human animals highlight that spatially resolved whole-brain neuroimaging will be necessary in humans to decipher the roles of anatomically distributed replay events in the brain.

## 1.4 Replay in humans

Our understanding of replay in the brain has been largely gained from electrophysiological studies in rodents. While work on replay in rodents has brought fundamental insights, it has strongly focused on replay of hippocampal representations of physical space and its relationship to spatial navigation behavior. This has left several intriguing questions about replay – for instance, its role in non-spatial episodic memories and relationship to subjective experience – unaddressed and pointed out that answers may only be found in humans (for reviews of human replay, see e.g., Zhang et al., 2017; Tambini & Davachi, 2019). Two notable exemplary observations include recent findings that human replay assembles novel sensory observations according to a previously learned sequential structure (Liu et al., 2019), and reflects abstract non-spatial task representations that are based on integrating current sensory input with information stored in memory (Schuck & Niv, 2019). The depth of insights about replay that will be gained from future studies in humans is interlinked with the advancement of non-invasive measurement techniques that allow to capture the key characteristics of replay described above.

### Methods to study replay in humans

The closest link between replay in rodents and humans can be established by studying invasive recordings in patient populations, using intracranial electroencephalography (iEEG) (Engel et al., 2005). The practical scope of iEEG recordings is limited, as they can only be performed in specialized hospitals on a small number of patients who are affected by neurological disorders like epilepsy or Parkinson’s disease, and often receive medication – factors which can limit



the generalizability of findings to non-patient populations. For these reasons, non-invasive neuroimaging methods, such as functional magnetic resonance imaging (fMRI), magnetoencephalography (MEG), and electroencephalography (EEG) are indispensable to advance the understanding of replay in humans. However, the key characteristics of replay described above (i.e., its speed, sequentiality, and anatomical specificity) pose a considerable challenge for non-invasive neuroimaging tools because existing methods typically offer either good temporal resolution or good spatial resolution, but not both. The speed and sequentiality of replay make MEG or EEG sensible methods to choose for the study of replay. In particular, since methods for the detection of replay using MEG have been first introduced (cf. Kurth-Nelson et al., 2015) and further refined (Liu et al., 2021a), the field has witnessed a vast number of replay studies in humans using MEG (Michelmann et al., 2018; Liu et al., 2019; Eldar et al., 2020; Wimmer et al., 2020; Liu et al., 2021b; Nour et al., 2021; Wise et al., 2021; Wimmer et al., 2021; Higgins et al., 2020; Schwartenbeck et al., 2021). While MEG studies have yielded intriguing insights into replay in the human brain, MEG is known to have limited spatial resolution which makes it difficult to confidently localize the anatomical source of the neural signal. This makes it challenging to disambiguate co-occurring replay events in the brain. fMRI, in turn, offers the required spatial specificity. Over recent years, two main approaches to study memory reactivation in humans using fMRI have been established (for reviews, see e.g., Zhang et al., 2017; Tambini & Davachi, 2019). The first approach quantifies reactivation as the similarity of multi-voxel patterns of individual stimuli activated during task performance to post-encoding rest and pre-encoding baseline (Deuker et al., 2013; Staresina et al., 2013; Tambini & Davachi, 2013; Schapiro et al., 2018; Schlichting & Preston, 2014; Sadeh et al., 2019). When multi-voxel patterns activated during behavior are identified during post-encoding rest, and the occurrence of these events is statistically more frequent than the pre-encoding baseline, this is considered as evidence that reactivation has occurred. The second approach examines changes in functional connectivity between hippocampal, cortical, and dopaminergic brain structures and their relationship to post-encoding memory consolidation (Tambini et al., 2010; Tompariy et al., 2015; Murty et al., 2017; Gruber et al., 2016; Hermans et al., 2016; de Voogd et al., 2016). When functional connectivity between task-related brain activations changes from pre- to post-experience measurements and is related to behavioral indicators of memory consolidation, this is thought to reflect replay-mediated interactions between brain areas. Following the

considerations of replay sequentiality described above, the large majority of fMRI studies that employed these analysis approaches reported reactivation of single task-related aspects but did not provide evidence for *sequential* replay directly and were also insensitive to the speed of replay events. First fMRI evidence for sequential replay in humans was reported only more recently by Schuck and Niv (2019). These authors trained fMRI pattern classifiers on previously learned task states in a non-spatial decision-making paradigm, and, analyzing the ordering of classifier predictions in each volume of fMRI data, found evidence for forward-ordered sequential reactivation of task states during post-task rest which was related to improved decisions (Schuck & Niv, 2019). However, questions about how exactly fMRI gives rise to forward and backward sequences and how it could be used to assess the speed of replay still remained. These questions were addressed in *Paper II*.

## 1.5 Replay in artificial agents

Replay is not only relevant to the brain. At around the same time that replay was discovered in neuroscience in the early 90s, “experience replay” was introduced to machine learning (ML) research as a novel technique to improve the efficiency of learning in artificial agents navigating to a rewarding goal location (Lin, 1991)<sup>1</sup>. During experience replay, the learning episodes of an agent are temporally stored in a memory buffer and presented again during later offline periods, which enables the agent to learn from previous interactions with the environment multiple times, requiring less data. In addition, Lin (1991) showed that replaying learning episodes in backward order (i.e., from the goal location to the starting point), allowed the agent to propagate credit for obtaining the reward along the previous sequence of states, resulting in faster learning. This provided an early computational account for how reverse replay benefits learning, wherein experience replay allows an agent to learn faster from less data. Several years later, experience replay became particularly popular in the ML field, when it was demonstrated to be a crucial component in training a deep neural network (DNN) to learn multiple Atari video games to human performance levels (Mnih et al., 2015). Since then, replay has

---

<sup>1</sup>Of note, the term “replay” seems to have first been used in the ML literature (Lin, 1991) before it entered the neuroscience literature about five years later (Skaggs & McNaughton, 1996), although the phenomenon was described earlier in neuroscience (Pavlidis & Winson, 1989).

been involved in several advances in artificial intelligence (AI) research that were inspired or followed by matching discoveries in neuroscientific replay research (for reviews, see Kumaran et al., 2016; Hassabis et al., 2017). For example, it was shown that replay accelerates learning considerably if those experiences are selected for replay that are most surprising to the agent (and therefore offer the largest gain in learning), rather than selecting experiences at random (Schaul et al., 2015), an idea that has also influenced neuroscience (e.g., Mattar & Daw, 2018; Liu et al., 2021b). Despite the striking overlap, research on biological and artificial replay has unfolded largely in parallel and comprehensive comparisons are still scarce (but see Cazé et al., 2018; Hayes et al., 2021; Roscow et al., 2021). Integrating findings from both research fields promises novel insights into the diversity of replay and its computational benefits for learning and decision-making. These aspects serve as the main motivation for the literature review presented in *Paper I*.

## 1.6 The content of replay

A central objective in deciphering the functional relevance of replay is to understand its content, i.e., what replay events represent (Foster, 2017; Pfeiffer, 2018). The early interpretation that replay reflects the same physical locations in the environment that the animal traversed during previous spatial navigation, is complicated by the diversity of replay found in rodents, the integration of non-spatial replay found in humans, and the definition of replayed states in artificial agents. The variety of these observations has made clear that replay may not just be a veridical recapitulation of the most recent experience, but could instead offer a window into the internal representations and computations that the brain affords to support a wide range of cognitive abilities. In addition to summarizing the main computational benefits of replay, *Paper I* also explores the claim that replay could be involved in building representations that serve learning and decision-making.

Many ideas exist in both neuroscience and ML about how biological and artificial agents might learn and maintain representations of the environment that enable flexible behavior (for reviews, see e.g., Bengio et al., 2012; Niv, 2019). One recent idea from computational neuroscience transforms the concept of the hippocampus as a cognitive map (see O’Keefe & Nadel, 1978, and the beginning of this chapter) and instead proposes that the hippocampus

hosts a compact representation that summarizes future events in a *predictive map* (Stachenfeld et al., 2017). This computational theory proposes that place cells in the hippocampus do not encode the current location of an animal per se, but rather represent each state in terms of the future states that it predicts (up to a certain predictive horizon), formally known in ML as the “successor representation (SR)” (Dayan, 1993). A predictive map in the form of the SR allows agents to adapt rapidly to changing rewards in the environment and anticipate relationships between events that are not necessarily the closest in space or time, but the most predictive of what an agent will experience in the future (Momennejad et al., 2017; Momennejad, 2020). Notably, computational work shows that replay can be used to learn the SR during offline periods which allows efficiently computing predictive relationships before they are retrieved for planning and decision-making later (Russek et al., 2017; Momennejad et al., 2017). Based on this theoretical background, the fMRI study reported in *Paper III* investigated how humans learn an SR of the relationships among sequential stimuli and replay sequences from this internal task representation during on-task intervals to anticipate upcoming events.

# 2

## RESEARCH QUESTIONS

Given the diversity of replay in rodents, humans, and artificial agents, the methodological challenges of studying replay in humans with fMRI, and emerging ideas about the relationship between replay and its underlying internal representations, this dissertation is aimed to address the following three research questions:

**Question 1: What are the major computational benefits of replay for learning and decision-making in biological and artificial agents?**

Advances in the understanding of replay in neuroscience and machine learning (ML) research have progressed largely in parallel and direct comparisons are still scarce (but see Cazé et al., 2018; Hayes et al., 2021; Roscow et al., 2021). Integrating findings from both research fields promises to drive novel interdisciplinary insights. The objective of *Paper I* was therefore to establish links between the neuroscience and ML literature on replay and portray the diversity of computational functions that enhance learning and decision-making in biological and artificial agents alike.

**Question 2: How can we measure fast and sequential replay in humans using functional magnetic resonance imaging (fMRI)?**

Defining characteristics of replay are its speed, sequentiality, and anatomical specificity. Previous fMRI research on replay has left largely unclear how sequential aspects of replay events can be extracted from fMRI signals, and if fMRI is sensitive to varying replay speeds, particularly speeds on the order of milliseconds, as observed in rodents. Furthermore, previous work in humans did not yet take full advantage of the detailed spatial resolution of fMRI that would give insights into the functional role of simultaneous replay events distributed in the brain. This leaves open the question of how fMRI can be used to investigate the speed and sequential nature of replay, which was addressed in *Paper II* and further elaborated in *Paper III*.

**Question 3: Does replay reflect abstract task representations?**

The variety of representations that replay can depict has led to the suggestion that replay could reflect samples from an internal representation of the world (Foster, 2017). As one example of such a representation, previous work has described the hippocampus to host a predictive map in the form of a successor representation (SR) (Stachenfeld et al., 2017). Computational work has indicated that replay could provide a mechanism to update such predictive representations offline (Momennejad et al., 2017; Russek et al., 2017), but how the SR influences replay during short pauses from ongoing behavior remains to be explored. In *Paper III* we asked how humans learn higher-order relationships among consecutive task stimuli and whether they form multi-step predictive representations in the form of an SR. We then asked if online replay during on-task intervals would reflect sampling from an SR to anticipate upcoming stimulus sequences.

Understanding the content of replay as abstract state representations will be beneficial for future replay research, promises to foster closer correspondence between insights from neuroscience and ML, and could achieve a more mechanistic account of the diverse computational functions of replay in both biological and artificial agents (see Question 1). This claim was further explored in *Paper I*.

The present dissertation is publication-based and the research questions listed above are considered in three different papers. *Paper I* addresses Question 1 and also speaks to Question 3. *Paper II* addresses Question 2. *Paper III* mainly addresses Question 3 and provides additional insights into Question 2. The following chapter will summarize the main findings of these three papers.

# 3

## OVERVIEW OF PAPERS

The present dissertation is based on three articles that provide theoretical, methodological, and empirical insights into the role of replay in learning and decision-making and its measurement in humans using fMRI. *Paper I* offers a theoretical perspective on the diversity of putative functions of replay in biological and artificial agents. It integrates findings from the neuroscience and machine learning (ML) literature and distills key computational benefits of replay for decision-making and learning. *Paper II* provides empirical insights into how replay, and fast sequential neural events more generally, can be measured non-invasively in humans using fMRI and demonstrates replay of task-related stimuli during awake post-task rest. *Paper III* builds on the theoretical framework developed in *Paper I* and the methodological advances in *Paper II* and shows how replay during on-task intervals is related to learning of map-like predictive representations.

All papers are based on work that was conducted within the Max Planck Research Group “Neural and Computational Basis of Learning, Decision Making and Memory (NeuroCode)”, led by Dr. Nicolas W. Schuck at the Max Planck Institute for Human Development in Berlin, Germany. The work was funded by an Independent Max Planck Research Group grant awarded to Dr. Nicolas W. Schuck by the Max Planck Society (M.TN.A.BILD0004), and a Starting Grant awarded to Dr. Nicolas W. Schuck by the European Union (ERC-2019-StG REPLAY-852669). The work in *Paper II* was also supported by a DAAD travel stipend awarded by the German Academic Exchange Service.

### **Paper I: A literature review on the computational benefits of replay in biological and artificial agents**

**Wittkuhn, L.**, Chien, S., Hall-McMaster, S., & Schuck, N. W. (2021). Replay in minds and machines. *Neuroscience & Biobehavioral Reviews*, *129*, 367–388. doi: [10.1016/j.neubiorev.2021.08.002](https://doi.org/10.1016/j.neubiorev.2021.08.002)

In this literature review article, we summarized putative computational functions of replay for learning and decision-making in biological and artificial agents. Combining insights from

neuroscience and machine learning (ML), we identified five key computational benefits which included faster learning, less forgetting, the reorganization and augmentation of experiences, planning, and generalization. Finally, we highlighted the benefit of replaying abstract state representations and discussed how replay could provide a mechanism to build such internal representations.

## Summary of Contents

The literature review begins with a short historic perspective on the neuroscientific replay literature that has predominantly interpreted replay as a consolidation process that reactivates previous experience to support spatial navigation and memory (for similar historic perspectives and reviews, see e.g., Redish, 1999; Sutherland & McNaughton, 2000; Foster, 2017). We then argue that the replay phenomenon is much more diverse than initially thought as it occurs during several behavioral states (sleep, wakeful rest, active behavior), at various speeds, and in multiple brain areas, and could be involved in a wide range of cognitive functions. We highlight that insights into this diversity of findings may be found in the ML literature where “experience replay” was introduced in the early 90s (Lin, 1991) and became popular as a crucial ingredient in training DNNs (Mnih et al., 2015). The main body of the review then focuses on key computational benefits of replay that support learning and decision-making in both biological and artificial agents.

First, replay can accelerate learning and improve data efficiency (for reviews, see e.g., Hassabis et al., 2017; Kumaran et al., 2016; Cazé et al., 2018). For example, awake backward replay is thought to retrieve sequences of states that bridge delays between actions and their outcomes (cf. Minsky, 1961). Using simulations of an agent navigating to a goal location in a grid world, we illustrate how backward replay of state sequences allows the agent to reach the goal and receive more reward with fewer learning episodes (see Figure 2 in *Paper I*).

Second, replay can help an agent to forget less and amplify rare experiences. Rooted in the theory of complementary learning systems (CLS) in the brain (McClelland et al., 1995; O’Reilly et al., 2014; Schapiro et al., 2017), replay could interleave past with present experience, thereby addressing the problem of *catastrophic interference* that causes artificial neural networks (ANNs) to forget previously learned information when learning to perform new tasks (McCloskey & Cohen, 1989; Ratcliff, 1990; French, 1999). Replay might even allow to flexibly



adjust the number of learning opportunities from previous experience, for instance by sampling rare but important experiences more frequently (Wang et al., 2016), like paths to a shock zone that should be avoided (e.g., Wu et al., 2017).

Third, replay can arbitrarily adjust the distribution of previous experiences. Experiences might be replayed randomly both in DNNs (Mnih et al., 2015) and rodents (Stella et al., 2019). In ML, replay that prioritizes state transitions leading to important outcomes promotes efficient learning (Schaul et al., 2015; Horgan et al., 2018), and correspondingly, neural replay has been shown to be influenced by rewards (e.g., Foster & Wilson, 2006; Singer & Frank, 2009; Ólafsdóttir et al., 2015; Ambrose et al., 2016; Roscow et al., 2019; Liu et al., 2019; Liu et al., 2021b). In rodents, replayed sequences have been shown to reflect a wide range of trajectories, from paths in the immediate past or near future to entirely remote locations, often depending on the behavioral state (e.g., rest vs. active behavior) of the animal (Ólafsdóttir et al., 2017; Karlsson & Frank, 2009; Gupta et al., 2010).

Fourth, replay can support planning and goal-directed decision-making. Several previous studies demonstrated that rodents at choice points often appear to replay paths in a forward direction from the starting to a goal location, as if deliberating upcoming choices (Johnson & Redish, 2007; Diba & Buzsáki, 2007; Davidson et al., 2009; Karlsson & Frank, 2009; Pfeiffer & Foster, 2013; Singer et al., 2013; Ambrose et al., 2016). These observations could reflect sampling from a learned model, as in the DYNA architecture (Sutton, 1990, 1991). To what extent replayed sequences predict subsequent behavioral trajectories can depend on the learning history (Khamassi & Girard, 2020) as well as the task setting and motivational state of the animal (Carey et al., 2019; Wu et al., 2017). Finally, rodents have been observed to *preplay* place cell sequences consistent with spatial environments that have not been experienced yet (Dragoi & Tonegawa, 2011, 2013a, 2013b). While preplay has been contested (e.g., Silva et al., 2015), potential explanations can be found in the ML literature which include attractor dynamics (Corneil & Gerstner, 2015), and reservoir computing (Cazin et al., 2019; Leibold, 2020).

Fifth, replay can support inference and generalization. Artificial agents can learn generative models that are used to produce replay events for learning (e.g., Stoianov et al., 2020). Further, replay could be used to combine generalized task structure with sensory specifics to recreate sequences in particular orders (Liu et al., 2019). Finally, replay may also interact with predictive

task representations like the successor representation (SR) (Dayan, 1993; Russek et al., 2021).

Sixth, replay may reflect task state representations and could be involved in learning such representations. The efficiency of many ML algorithms depends on how the agent internally represents the environment (Dayan, 1993), but learning useful representations poses a challenge for biological and artificial agents alike, for instance because potentially relevant features of the environment are numerous or only partially observable (Bengio et al., 2013; Niv, 2019). Given that place cells, often thought of as the main neural substrate of replay events in the hippocampus, have been shown to represent non-spatial features of the environment, like sounds (Aronov et al., 2017), time (MacDonald et al., 2011), or SRs (Stachenfeld et al., 2017), individual events in a replay sequence could be more usefully understood as abstract *states* which offers interpreting results from neuroscience and ML in a formal framework.

Finally, despite the broad range of computational functions that replay is involved in, we acknowledge that goal-directed behavior can exist without replay. For instance, in *episodic* reinforcement learning (RL) agents rely on the retrieval of single episodes to control behavior (for reviews, see e.g., Gershman & Daw, 2017; Botvinick et al., 2019). While existing neuroscientific evidence (Bornstein et al., 2017; Lee et al., 2015; Wimmer & Büchel, 2021) suggests that such retrieval could be supported by sequential replay, this question still remains open.

In conclusion, we argue that RL theory and insights from ML provide a useful framework for understanding how biological and artificial agents learn representations of the environment and how replay operates on those internal representations. Replay thereby offers a window into the internal computations and state representations of the brain, fostering our understanding of how it influences learning and decision-making.

## Contributions to Open and Reproducible Science

The version-controlled R source code for the simulations reported in the review is publicly available on GitHub at <https://github.com/lnnrtwttkhn/replaysim-wittkuhn-et al2021>. The Git repository also includes the recipe for a Docker software container that can be retrieved from Docker Hub at <https://hub.docker.com/r/lennartwittkuhn/replaysim-wittkuhn-et al2021> and provides the entire computational environment that is required to run the simulations. The execution of the source code inside the Docker container reproduces the results of the simulations shown in Figure 2c–d of the paper which is demonstrated in a continuous in-

tegration (CI) workflow using GitHub Actions. All figures of the review were additionally released on Figshare at <https://doi.org/10.6084/m9.figshare.14261636.v4> under the terms of a Creative Commons Attribution 4.0 International license (CC-BY 4.0; for details, see <https://creativecommons.org/licenses/by/4.0/>) instead of the more restrictive license of the journal publisher, which facilitates independent reuse. Further bibliographical information and curated lists of the cited references are also available at <https://lennartwittkuhn.com/bibliography>.

## Paper II: An fMRI study to develop analysis methods to measure fast sequential neural replay in humans

**Wittkuhn, L., & Schuck, N. W. (2021).** Dynamics of fMRI patterns reflect sub-second activation sequences and reveal replay in human visual cortex. *Nature Communications*, 12(1795). doi: [10.1038/s41467-021-21970-2](https://doi.org/10.1038/s41467-021-21970-2)

The main goal of this fMRI study was to develop analysis methods based on probabilistic fMRI pattern classifiers that allow studying the sequentiality and speed of fast neural event sequences, like replay, in humans using fMRI, and to demonstrate their sensitivity for detecting sequential replay in post-task fMRI resting-state data.

### Theoretical Background

It is widely thought that the temporal resolution of fMRI is too limited to detect fast neural events because it measures neural activity only indirectly through slow sampling of a delayed blood-oxygen-level dependent (BOLD) response function (Ogawa et al., 1990; Kwong et al., 1992; Heeger & Ress, 2002). fMRI may therefore appear insensitive to the speed and sequentiality of replay and consequently previous fMRI studies in humans have indeed largely reported evidence for non-sequential reactivation (for reviews, see e.g., Zhang et al., 2017; Tambini & Davachi, 2019). Recently, however, Schuck and Niv (2019) leveraged sequential pattern analysis on each volume of fMRI data to demonstrate sequential replay of task states in resting-state fMRI data following a sequential decision-making task. The study reported in *Paper II* aimed to verify and extend the fMRI analysis techniques of Schuck and Niv (2019) and investigate to what extent fMRI is sensitive to the sequentiality and speed of fast neural events, like replay.

## Methods

Probabilistic pattern classifiers were trained on fMRI activation patterns in occipito-temporal cortex that were recorded while participants viewed individual images of everyday objects. The trained classifiers were then applied to fMRI data from trials where the same images were shown in sequences. The presentation speed of the sequences was varied in five steps from only 32 ms to 2048 ms between images (100 ms per image). We investigated probabilistic classifier time courses following image sequences and asked how the sequence speed influenced the detection of images and their temporal order. To quantify sequentiality, we developed a metric based on the slope of a linear regression that related classifier probabilities within each repetition time (TR) to a particular sequential order. The expected speed-dependent temporal dynamics of the classifier time courses were modeled using a sine-based response function that was derived from classifier time courses in single trials. We then applied frequency spectrum analysis to the sequentiality time courses in pre-task resting-state data that was augmented with sequential neural events at fast (32 ms) or slow (2048 ms) speeds to investigate its sensitivity to sub-second vs. supra-second sequences. Finally, we compared the frequency spectra of resting-state data before and after the task for evidence of replay of task stimuli in post-task rest periods.

## Major Findings

We found that probabilistic classifier time courses reflected the content and order of neural event sequences within individual TRs, quantified by our sequentiality metric. The magnitude of the sequentiality metric diminished with faster sequence speed but was still evident at fast sequence speeds of 32 ms, and exhibited a speed-dependent time course that could be predicted from the sine-based model of time-shifted single events. This time course was characterized by forward sequentiality in earlier TRs and backward sequentiality in later TRs, reflecting the dynamics of overlapping hemodynamic response functions (HRFs). A frequency spectrum analysis over the time course of the sequentiality metric in augmented pre-task resting-state data revealed power differences between fast and slow neural sequence data that were sensitive to changes in the number and signal strength of randomly inserted sequence events, mimicking spontaneous replay sequences. Finally, this analysis revealed fast replay of task-related stimuli in post-task compared to pre-task resting-state fMRI data.

## Contributions to Open and Reproducible Science

All data and code used in the study are publicly shared in version-controlled repositories via the G-Node Infrastructure (GIN; <https://gin.g-node.org/lnnrtwttkhn>) and GitHub (<https://github.com/lnnrtwttkhn/>). This includes the magnetic resonance imaging (MRI) and behavioral data organized in the Brain Imaging Data Structure (BIDS) format (cf. Gorgolewski et al., 2016), MRI quality metrics based on MRIQC (cf. Esteban et al., 2017), preprocessed MRI data based on fMRIPrep (cf. Esteban et al., 2018), binarized anatomical masks, results of first-level general linear models (GLMs) and multivariate decoding, as well as unprocessed behavioral data, amounting to roughly 1.5 terabyte of openly shared research data. Code for the main statistical analyses (<https://github.com/lnnrtwttkhn/highspeed-analysis>) and behavioral task (<https://github.com/lnnrtwttkhn/highspeed-task>) were shared via GitHub. Data and code management was performed using DataLad (Halchenko et al., 2019; Halchenko et al., 2021). An overview of all resources and results is available on a project website at <https://wittkuhn.mpib.berlin/highspeed/>. The website is built by retrieving the input data from GIN and executing reproducible code notebooks inside a Docker container environment using CI. Further details can be found in the data and code availability statements of the article. The article was published in the open access journal *Nature Communications*, supported by funding enabled and organized by Projekt DEAL (<https://www.projekt-deal.de/>).

## Paper III: An fMRI study to investigate on-task replay of predictive map-like task representations

**Wittkuhn, L., Krippner, L. M., & Schuck, N. W. (2022).** Statistical learning of successor representations is related to on-task replay. *bioRxiv*. doi: [10.1101/2022.02.02.478787](https://doi.org/10.1101/2022.02.02.478787)

The major aim of this fMRI study was to investigate how humans learn predictive representations of graph-structured statistical relationships between task stimuli from continuous experience and replay sequences sampled from these internal cognitive maps during short pauses from ongoing task performance.

## Theoretical Background

Replay has been described as sampling sequences from a learned model of task transition structure (see e.g., Schuck & Niv, 2019; Wikenheiser & Redish, 2015; Yu et al., 2021, and *Paper I*). These task representations might be predictive and represent each state in terms of the successor states that can follow from it (Garvert et al., 2017; Stachenfeld et al., 2017; Momennejad, 2020; Russek et al., 2017). Predictive task representations could be used to plan and anticipate upcoming event sequences. Replay has been suggested to play a role in planning, where place cell activations sweep ahead of the animal along multiple potential future trajectories (Johnson & Redish, 2007; Kay et al., 2020). So far, it remains unknown to what extent replay reflects predictive representations and reactivates them during on-task pauses to anticipate upcoming stimulus sequences.

## Methods

Participants performed a statistical learning paradigm (cf. Schapiro et al., 2012; Garvert et al., 2017) that exposed them to a fast-paced stream of images. Unbeknownst to the participants, the sequential ordering of images was governed by their arrangement in two ring-like graph structures that resulted in distinct transition probabilities between the images (cf. Garvert et al., 2017; Lynn et al., 2020). We hypothesized that incidental statistical learning would lead participants to extract higher-order predictive relationships among the task stimuli that could be used for replaying upcoming stimulus sequences during on-task pauses. Participants' response times were modeled using SRs (Dayan, 1993) with varying predictive horizons which reflect the expected visitations of future events, discounted up to a certain depth (Momennejad et al., 2017; Russek et al., 2017; Momennejad, 2020). Applying the sequential fMRI pattern analysis methods developed in *Paper II*, we examined the data for evidence of online neural replay during short on-task intervals that were interspersed with ongoing task performance. We used a hidden markov model (HMM) to predict the probability of observing a particular sequence in fMRI data based on the SR and examined our sequentiality metric for the most likely sequences given this predictive representation. Finally, a post-task questionnaire assessed explicit sequence knowledge.

## Major Findings

Behavioral modeling of response times based on an SR model of experienced stimulus transitions, indicated that participants formed multi-step predictive representations of the statistical relationships among the task stimuli. Interestingly, the predictive depth of this representation varied depending on which graph structure participants had learned and in which order, which may suggest that the brain hosts multiple SRs in parallel (Momennejad & Howard, 2018; Brunec & Momennejad, 2021). fMRI signals during short pauses from ongoing behavior, indicated sequential on-task replay particularly of those sequences that were most likely given the SR model. Results from the post-task questionnaire indicated that sequence knowledge remained implicit in half of participants.

# 4

## DISCUSSION

In this chapter, I will summarize the main findings of the two fMRI studies reported in *Paper II* and *Paper III* and interpret the major insights with respect to the existing literature. I will then consider the most significant limitations of the current work and discuss potential avenues for future research based on the conclusions drawn from the presented studies.

### 4.1 Summary and evaluation

#### **fMRI is sensitive to the sequentiality of fast replay events**

Replay is characterized by its sub-second speed, sequential order of reactivated neural representations, and anatomical distribution throughout the brain (Foster, 2017). Capturing all of these aspects using non-invasive neuroimaging in humans is a challenge because existing methods offer either sufficient temporal *or* spatial resolution, but not both. In particular, functional magnetic resonance imaging (fMRI) provides anatomical specificity but is generally considered to lack temporal detail to identify fast and sequential neural events, like replay. Hence, previous fMRI studies almost exclusively reported evidence for non-sequential reactivation and did not take the variability in replay speed into account (for review, see e.g., Tambini & Davachi, 2019). In *Paper II* we examined neural activation patterns following the presentation of image sequences at varying speeds to investigate if we can identify the sequential order and speed of experimentally controlled sequences based on the evoked neural signals alone. We developed a sequentiality metric that was quantified by the slope of a linear regression that related classifier probabilities to a particular sequential order and could distinguish between forward and backward directionality of neural sequences. Analyzing the order of probabilistic classifier time courses for each sequence item within a *single* measurement (i.e., within each volume of fMRI data) provided considerably increased statistical resolution to assess sequential ordering compared to previous methods that focused on the categorical classifier prediction with the highest probability (e.g., Schuck & Niv, 2019). Following known sequences of visual objects, the sequentiality metric indicated forward directionality in earlier TRs (i.e., same order as the true sequence) and backward directionality in later TRs (i.e., reverse order of the true



sequence). The magnitude of the sequentiality metric diminished with faster sequence speeds but the characteristic forward–backward dynamic was still evident when images were separated by only 32 ms. This represents an important insight for the study of replay, given that replay is known to occur in both forward and backward sequences (Foster, 2017). In fMRI data, the same replay sequence will thus elicit both forward and backward directionality in the sequentiality metric relative to the true underlying neural sequence, where the first direction indicates the true order and the second direction the reverse order.

In addition to this aspect, results from both *Paper II* and *Paper III* demonstrated that classifier probabilities were strongly influenced by the most recent item in the preceding neural sequence. In *Paper II* this concerned the last item in a visual sequence presented to participants before the measurement time window and in *Paper III* this related to the last image that was shown shortly before an on-task interval occurred. In *Paper II*, in order to provide a realistic match to spontaneously occurring replay events, image sequences were not masked. The dominance of the last sequence item led to larger backward compared to forward sequentiality in each of the five sequence speed levels that we examined in *Paper II*. In *Paper III*, we accounted for this effect by investigating sequentiality after excluding data related to the image on the current trial. While this appears like a potential limitation, given that replay events in rodents often start or end with goal locations (see e.g., Pfeiffer & Foster, 2013), the dominance of particular sequence items in fMRI data could also reveal information about significant task aspects that are preferentially replayed.

### **fMRI can distinguish between sub-second and supra-second sequences**

Analyzing post-task resting-state fMRI data in *Paper II*, we found evidence for replay at a speed expected for fast sequential events separated by only 32 ms, as in the fastest category of previously experienced stimulus sequences. Of note, the replayed sequences also included stimulus combinations that appeared during slower image sequences with intervals of up to 2048 ms between items. This finding is indicative of fast replay at a characteristic speed of a few tens of milliseconds, which is accelerated independently of the pace of sensory observations during behavior. These results align with the fast replay speed typically observed in rodents (Nádasdy et al., 1999; Lee & Wilson, 2002; Karlsson & Frank, 2009; Davidson et al., 2009), intracranial recordings in humans (Axmacher et al., 2008), as well as previous MEG studies that found

replay with stimulus-to-stimulus lags between 20 and 70 ms (Liu et al., 2019; Wimmer et al., 2020; Liu et al., 2021b). Given that the duration of behavioral episodes can be quite variable in rodents, the constant replay speed that is reported across studies suggests that replay depicts the sequential order of the behavioral experience rather than the exact time course of specific episodes (Foster, 2017). Of note, in *Paper II* we also found evidence for sequential replay that was slower than the slowest stimulus sequences in the task, with a temporal delay of more than 2048 ms between stimuli. In rodents, both reactivation of single locations (Yu et al., 2017), and replay at behavioral timescales (Denovellis et al., 2021) have been observed. It is currently unknown to what extent replay could reflect the reactivation of separate task aspects by different replay speeds. If our observations indicate slower conscious thought about stimulus sequences or reactivation of single stimuli remains speculative. To address these questions, study protocols will be needed that interrogate the interactions between spontaneous replay and voluntary thoughts in more detail (cf. Mildner & Tamir, 2019; Finn, 2021). Together, the analysis methods developed in *Paper II* allow insights into sequential reactivation at fast and slow speeds and could provide the basis for future investigations about the functional relevance of replay speed. Subsequent studies could also further illuminate to what extent the speed of replay is bound to the speed of experience, or whether replay speed is another parameter that the brain may dynamically adjust depending on its functional purpose.

### **Modeling replay sequences informs their detection with fMRI**

Computational and statistical models of replay and their underlying task representations allow quantitative predictions about which and when particular replay sequences are expected to occur (see e.g., Mattar & Daw, 2018; Schuck & Niv, 2019). In addition, these predictions can be combined with knowledge about how the resulting neural signals will be captured by the neuroimaging method used in the given study (e.g., fMRI). In *Paper II*, we demonstrated that a sine-based model of classifier time courses following single image presentations could predict how the speed of image sequences would affect classifier time courses following sequential image presentation that were strikingly congruent with experimental data. These findings therefore provide insights into the dynamics of classifier time courses following replay sequences. Modeling classifier time courses of single and sequential neural events could offer an approach to disentangle stimulus-driven from spontaneous neural activity in future studies.

In *Paper II*, we also found that fMRI is particularly sensitive to the first and last item of a neural event sequence, but that measurement noise could lead to reduced sensitivity for the ordering of intermediate sequence items. Extending our toolkit for detecting replay with fMRI, in *Paper III* we incorporated this knowledge into a hidden markov model (HMM) that yielded the probability of observing a specific sequence in the fMRI signal given that a particular sequence occurred in the brain. In this model, the probability of observing the true sequence item was much higher for items that occurred first or last in the sequence compared to intermediate items, as we found empirically in *Paper II*. The probability that a certain sequence would be replayed was calculated under the assumption that replay would internally sample from a successor representation (SR) model that we had used to analyze participants' response times. Combining assumptions about which sequences are likely to be replayed given a model of the task, and which sequences are likely to be observed in fMRI data offers a promising approach to gain additional sensitivity for detecting the content and sequential order of replayed neural representations with fMRI.

### **fMRI evidence for online replay during on-task intervals**

In *Paper III*, we investigated online replay during short on-task intervals that were interspersed with ongoing task performance. In rodents, replay events during short pauses from active navigation are associated with planning, decision-making, and deliberation (for reviews, see e.g., Redish, 2016; Pezzulo et al., 2019; Drieu & Zugaro, 2019; Zielinski et al., 2020). Several previous fMRI studies have investigated the neural reinstatement of associative memories during instructed retrieval (e.g., Staresina et al., 2012; Ritchey et al., 2012; Gordon et al., 2013; Bosch et al., 2014) but, so far, knowledge about the temporal dynamics of on-task memory recall in humans largely stemmed from electrophysiological recordings using EEG or MEG, and mainly focused on reactivation elicited by explicit external cues (for review, see e.g., Staresina & Wimber, 2019). *Paper III* contributes fMRI evidence for how task-related sequences of stimuli are reactivated during on-task intervals and how their temporal dynamics unfold without explicit cues or instructions. Our main results showed that sequential replay in occipito-temporal and sensorimotor cortices during short on-task intervals was consistent with the probabilistic task structure that participants were exposed to during a statistical learning paradigm. Of note, while participants were informed that intervals would be randomly interleaved with ongoing

task performance, they did not know the exact timing and duration of those breaks. To what extent the replayed sequences during these intervals reflect ongoing maintenance of task representations (e.g., Zielinski et al., 2020), conscious deliberation and planning (e.g., Kurth-Nelson et al., 2016), or anticipatory neural signals of prospective sequences (e.g., Ekman et al., 2017) can not be conclusively clarified.

## 4.2 Limitations and open questions

### Avenues to further improve the detection of replay with fMRI

#### Modeling relative activation differences within TRs

One central methodological contribution of *Paper II* is the sequentiality metric that reflects the sequential ordering of classifier probabilities *within* each TR. The sequentiality metric is quantified by the slope of a linear regression that relates the classifier probabilities of the decoded classes to a particular sequential ordering. While we show in *Paper II* that this approach robustly indicates the true sequential ordering of classifier probabilities, it also has some limitations. First, the assumption that relative differences in classifier probabilities are adequately captured by a linear relationship might not hold true in general. It is likely that additional sensitivity can be gained by making more fine-grained assumptions about the relationship between relative classifier probability differences within each TR. This might also enhance the sensitivity for intermediate sequence items. In addition, because the linear regression is only based on as many data points as there are decoded classes, and the number of classes is usually small (only five and six classes in *Paper II* and *Paper III*, respectively), the fit of the linear model might be heavily influenced by extreme data points. Given that we observed a strong dominance of the last sequence item in both *Paper II* and *Paper III* it seems likely that these data points drive the sequentiality metric. Furthermore, the slope of the linear regression only indicates a general trend in the data but does not indicate differences between successive pairs of items within the sequence. For example, in a hypothetical sequence *ABCDE* where *E* carries a strong signal because it is the last item, the slope will be strongly influenced by *E* but relative differences between e.g., *B* and *C* can not be resolved. Therefore, an important limitation of *Paper II* is that the sequentiality metric can identify the overall directionality of sequentially

ordered classifier probabilities but does not assess pairwise activation differences, in particular of intermediate sequence items. Future studies might therefore consider alternative (e.g., non-linear) models to describe the relative differences between classifier probabilities.

### **Replay detection is only as good as the classifiers**

The replay detection methods we developed in *Paper II* and extended in *Paper III* depend on a strong performance of the classifiers. In order to take full advantage of the high spatial resolution of fMRI and investigate simultaneous replay in separate brain regions, sufficient classifier performance should be obtained in all anatomical regions of interest (ROIs). In both fMRI studies, we could achieve levels of mean decoding accuracy that surpassed the chance-level considerably (69.22% vs. chance level of 20.0% in visual cortex in *Paper II*; 63.08% and 47.05% vs. chance level of 16.67% in visual cortex and motor cortex *Paper III*, respectively). While these levels of classification accuracy are respectable and have yielded sufficiently powerful classifiers to perform replay detection, it is an open empirical question how much more sensitivity could be gained by further improving classifier performance. Furthermore, in both studies, decoding accuracy in the hippocampus – most often studied as the main site of replay in the brain (Foster, 2017) – did not surpass the chance level. The available evidence for the success of decoding stimulus category from hippocampal signals using multi-voxel pattern analysis (MVPA) is mixed, and includes both unsuccessful (e.g., Deuker et al., 2013; Diana et al., 2008) and successful examples (e.g., Hassabis et al., 2009; Schuck & Niv, 2019).

### **The costs of avoiding sequential bias in replay detection**

It is important that classifiers do not have a sequential bias. In both *Paper II* and *Paper III*, the task stimuli used for classifier training were only separated by inter-trial intervals (ITIs) of 2.5 s on average. Given the slow hemodynamic fMRI signal, activation of a preceding stimulus is still present when data for the subsequent stimulus is recorded. Sequentially ordered stimuli will therefore result in similarities between consecutive neural activation patterns which in turn can yield classifiers with a bias to detect a particular sequential order (cf. Cai et al., 2019), which would compromise an unbiased detection of sequential replay. For this reason, we formed a classifier training set in *Paper II* that consisted of all 120 sequential combinations of the five stimulus classes. Similarly, in *Paper III*, the trial procedure was set up in a way that

all pairwise combinations of the six stimulus classes were presented equally often. Notably, trial procedures that are aimed at reducing sequential bias in classifier training sets will result in an exponentially increasing number of required trials which poses practical challenges as they strongly affect study duration. In the cases of the two presented studies, this resulted in experiments with two long MRI sessions of which roughly the equivalent of an entire MRI session of more than 60 minutes (min) was dedicated to the collection of classifier training data. While a larger training set is generally considered beneficial for classification robustness (e.g., Pereira et al., 2009), the interactions between the number of classifier classes, the need for a sufficiently large training set, and the requirement to avoid sequential bias in the training set need to be balanced in future studies. For instance, future studies might consider having more training classes with fewer trials recorded for each. While this might result in a noisier template pattern for a given class, it would allow tasks with more stimuli and also increase the number of data points available for the sequentiality metric at each TR, potentially improving sensitivity. The impact of such tradeoffs could be formally investigated in future methodological work.

### **Cognitive demands of the task may influence classification**

Finally, the cognitive demands of the task that is performed while fMRI data for classifier training are collected may influence decoding success. In *Paper II*, classifier training data was gained from a simple oddball detection paradigm that required participants to respond to occasional upside-down images. This task was intentionally designed to not involve any mnemonic component which might offer an explanation for why the hippocampus was not sufficiently engaged during the task. In an attempt to take this insight into account for the design of the paradigm in *Paper III*, we devised a task that required participants to memorize the associations between visual stimuli and response keys. Participants were additionally instructed to actively remember those associations during task performance. Despite emphasizing mnemonic aspects of the task, we still did not obtain sufficient classifier performance in the hippocampus in *Paper III*. In a previous study, Schuck and Niv (2019) used a task in which participants had to constantly remember information about the previous trial, comprising non-observable task states that were successfully decoded from activity in the hippocampus. Potentially, tasks that draw on memory abilities more strongly will be needed in future studies to yield successful pattern classification that allows investigating sequential replay in the human hippocampus.

### Localizing replay events in time

Replay events during offline periods, like sleep or awake rest, occur spontaneously. Their precise quantity, timing, and temporal separation are usually uncertain to the researcher. In *Paper II*, we have demonstrated that a frequency spectrum analysis over the sequentiality metric is sensitive to the number and signal strength of replay events in artificially augmented resting-state data. Inserting fMRI data resulting from sequential stimulus presentations that were known to contain sequential activation patterns specifically increased power in the frequency spectrum that was expected for the given sequence speed (comparing slow versus fast sequential events). The current work did not examine how the onset of individual replay events could be detected in extended offline periods like rest or sleep. However, our findings provide three key insights that may guide future endeavors to localize replay events in time.

First, our findings suggest that fast replay events can be identified by lower variability of classifier probabilities. In *Paper II*, we have observed that the almost simultaneous activation of neural patterns during fast stimulus sequences leads to reduced classifier probabilities. Following fast neural sequences, the fMRI signal reflects a mixture of activity patterns from all classes, which causes the probabilities of stimulus-specific classifiers to be smaller, closer together, and less extreme. In consequence, the standard deviation of classifier probabilities from fast sequential neural events might be lower compared to slow sequential neural events or resting-state data, as we have shown in *Paper II*. A potential extension of this approach could calculate the standard deviation of classifier probabilities in a running window across TRs of a period of sleep or awake rest in order to identify changes in the variance of classifier probability indicative of fast replay events.

A second insight is that the slow signal changes of the BOLD response related to the rising and falling slope of the HRF cause the fMRI sequentiality metric to reflect the sequential ordering of fast events in the same (forward) order in earlier TRs and in reverse (backward) order during later TRs. Spontaneous replay events could be detected by examining the measurement period for time points where task items are first ordered in one direction and a few TRs later in the opposite direction. In these cases, the ordering during earlier TRs might reflect the true directionality of the underlying neural sequence and the ordering during later TRs the opposite sequential pattern.

Third, in *Paper II*, we had modeled the classifier time courses of single stimulus presentations using a sine-based response function. To derive expectations for the classifier dynamics of sequential events, we calculated the difference between two sine-based response functions that were shifted by a temporal delay  $\delta$ . We then calculated the frequency spectra of the time courses of the sequentiality metric separately for the pre- and post-task resting-state period and concatenated data of fast (32 ms) and slow (2048 ms) sequence trials. As shown in Figure 3d–e of *Paper II*, the time courses of the sequentiality metric derived from the sine-based modeling approach and those found in sequence trial data were clearly correlated. In a future approach, the frequency spectra of concatenated single trial time courses could be calculated for different levels of the time delay  $\delta$ . The resulting frequency spectra could then be compared to the frequency spectra during rest and sleep periods as a function of the  $\delta$  parameter. For example, if the frequency spectra of predicted and observed data show a higher level of congruence when predicted frequency spectra are derived from smaller compared to larger temporal delays ( $\delta$ ) this would suggest that rest or sleep periods contain fast sequential neural events. As decoding analyses are performed within-subject, this approach can be conducted using participant-specific parameters for the sine-based response functions in order to gain additional sensitivity.

Finally, researchers may opt for multi-modal measurement approaches like concurrent EEG-fMRI recordings that combine the advantages of both neuroimaging techniques. For example, using simultaneous EEG and fMRI recordings, researchers may anchor the fMRI analyses to the onset of EEG characteristics known to be associated with replay events like sleep spindles (for an example of this approach, see e.g., Bergmann et al., 2012). While direct evidence for a link between fast sequential replay and EEG oscillations is scarce (for review, see e.g., Schreiner & Staudigl, 2020), several studies in sleeping humans demonstrate a relationship between slow oscillations and sleep spindles in the EEG, behavioral indicators of memory consolidation, and non-sequential reactivation of task-related neural traces (Bergmann et al., 2012; Schönauer et al., 2017; Cairney et al., 2018; Schreiner et al., 2021). In humans, slow oscillations and sleep spindles are coupled to replay-associated sharp wave-ripples (SWRs) (Clemens et al., 2007; Staresina et al., 2015; Helfrich et al., 2019). In summary, several extensions of the presented methods are conceivable that would allow localizing replay events in time using fMRI.



## Investigating the role of replay in updating abstract task representations

### Models of neural task representations

How agents learn efficient representations of tasks that allow to solve them efficiently is a central question in both machine learning (ML) and neuroscience research (Bengio et al., 2012; Niv, 2019). In *Paper III*, we had hypothesized that participants would learn higher-order relationships among consecutively presented stimuli in the form of a successor representation (SR) (Dayan, 1993) and that the sequential order of on-task replay events would be informed by this internal task representation. Our decision to use the SR model was based on previous findings showing that the hippocampus may host a predictive map in the form of SRs (Stachenfeld et al., 2017), that both human behavior and fMRI signals during learning of graph-structured relationships match to the SR (Garvert et al., 2017), and that replay could be used to learn the SR during offline states (Russek et al., 2017). However, the SR is only one of several possible ways how participants might represent sequential relationships between task stimuli. For example, models that only include one-step transition probabilities (e.g., Momennejad et al., 2017) could provide viable alternatives to how participants represented the task in *Paper III*. It is also important to note, that the fMRI analyses in *Paper III* did not probe neural patterns that represented the abstract SR directly, but analyzed reactivated sequences of neural stimulus patterns whose order was determined by the SR. Previous work by Schuck and Niv (2019) has demonstrated that replay can reflect neural representations of abstract task states that are defined by non-observable task-relevant features. In order to investigate replay of abstract task representations, future studies will need to find ways to devise classifier training sets that allow researchers to train neural pattern classifiers on abstract task representations directly in order to probe their reemergence during replay later (cf. Schuck & Niv, 2019). Together, future studies may compare the predictions of different types of abstract task representations and develop novel classifier training schemes to investigate their reactivation during replay.

### Changes in neural representations could influence the detection of replay

In both fMRI studies presented in this dissertation, probabilistic pattern classifiers were trained on fMRI data from single events of the stimulus material (i.e., presentations of individual images in *Paper II* and *Paper III*, as well as motor responses in *Paper III*). The trained classifiers were

then applied to fMRI data known or assumed to contain sequential reactivation of the same task events that the classifiers were trained on. For example, fMRI data following the sequential presentation of images in *Paper II* was *known* to contain sequential neural activation patterns. In contrast, fMRI data from the resting-state scans in *Paper II* and on-task intervals in *Paper III* could only be *assumed* to contain sequential neural activation of task-related events. Neural decoding approaches that are employed in the study of replay often rely on the assumption that individual neural patterns activated in sequence during rest will be the same, or at least sufficiently similar to, neural patterns during previous task performance that classifiers were trained on (Finn, 2021). However, neural representations of task environments might not be as stable and predictable as often assumed. Referring back to the properties of place cells in the hippocampus (see Introduction), a particularly striking observation is that the same hippocampal neurons adapt their firing rate or change their preferred firing location entirely in response to different task environments or modifications to the same spatial context, a phenomenon known as *remapping* (for reviews, see e.g., Colgin et al., 2008; Jeffery, 2011). Broadly, the degree of remapping depends on the similarity between environments and has been shown to be influenced by changes in non-spatial features like the rotation of a visual cue (Bostock et al., 1991), odors (Anderson & Jeffery, 2003), or fear learning (Moita, 2004; Wang et al., 2012). The extent to which remapping could influence the content of replay has not been systematically investigated yet. In one recent example, Tirole et al. (2021) showed that adaptation in the firing rate of place cells in response to two different linear tracks, persisted into subsequent replay of behavioral episodes and allowed researchers to decode the identity of the spatial context from the replayed activity. Remapping has also been investigated in human fMRI studies, where hippocampal activity patterns change in response to distinct task environments (Kyle et al., 2015; Steemers et al., 2016). Neural representations can also change depending on their sequential relationships. For example, representations of visual stimuli in the MTL increase their similarity in response to repeated sequential exposure (Miyashita, 1988; Schapiro et al., 2012), or their proximity within a graph-like task structure (Garvert et al., 2017). How these alterations of neural representations interact with sequential replay remains largely unexplored in humans. In summary, future studies need to carefully consider how neural task representations are transformed in response to changes in the task environment or as a function of experience, and how this affects the content of neural replay.

### 4.3 Conclusion

In three publications, this cumulative dissertation explored the interplay between task representations and replay, its role for learning and decision-making, and its investigation in the human brain using fMRI. *Paper I* summarized the major computational functions of replay that support adaptive behavior in both biological and artificial agents. These benefits include faster learning, less forgetting, the reorganization and augmentation of experience, planning, and generalization. Moreover, the theoretical claim was explored that replay could provide a mechanism to use or even learn the internal representations of the environment that an agent constructs to solve a task. *Paper II* introduced and experimentally validated multivariate analysis methods for fMRI to measure fast neural event sequences, like neural replay, with anatomical precision in humans. The analysis of probabilistic classifier time courses allows studying neural sequences on sub-second timescales and disentangling them from slower supra-sequences during awake rest periods. These methods promise to pave the way for future investigations of the speed and sequentiality of replay events in humans. In addition, this work demonstrated how modern-day software technologies and digital infrastructures can be leveraged to openly share the full scope of analysis code, research data, and computational environments used in a research project in order to enhance the transparency and reproducibility of the scientific work. Finally, *Paper III* combined theoretical and methodological insights of the previous two papers, and related learning of predictive task representations to neural replay during short on-task intervals. In summary, this dissertation provided theoretical and empirical insights into the relationship between internal task representations and neural replay as well as methodological advances for fMRI that improve its investigation in the human brain.

## REFERENCES

- Ambrose, R. E., Pfeiffer, B. E., & Foster, D. J. (2016). Reverse replay of hippocampal place cells is uniquely modulated by changing reward. *Neuron*, *91*(5), 1124–1136. doi: [10.1016/j.neuron.2016.07.047](https://doi.org/10.1016/j.neuron.2016.07.047)
- Anderson, M. I., & Jeffery, K. J. (2003). Heterogeneous modulation of place cell firing by changes in context. *Journal of Neuroscience*, *23*(26), 8827–8835. doi: [10.1523/JNEUROSCI.23-26-08827.2003](https://doi.org/10.1523/JNEUROSCI.23-26-08827.2003)
- Aronov, D., Nevers, R., & Tank, D. W. (2017). Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature*, *543*(7647), 719–722. doi: [10.1038/nature21692](https://doi.org/10.1038/nature21692)
- Axmacher, N., Elger, C. E., & Fell, J. (2008). Ripples in the medial temporal lobe are relevant for human memory consolidation. *Brain*, *131*(7), 1806–1817. doi: [10.1093/brain/awn103](https://doi.org/10.1093/brain/awn103)
- Behrens, T. E., Muller, T. H., Whittington, J. C., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*, *100*(2), 490–509. doi: [10.1016/j.neuron.2018.10.002](https://doi.org/10.1016/j.neuron.2018.10.002)
- Bellmund, J. L. S., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*, *362*(6415), eaat6766. doi: [10.1126/science.aat6766](https://doi.org/10.1126/science.aat6766)
- Bengio, Y., Courville, A., & Vincent, P. (2012). Representation Learning: A Review and New Perspectives. *arXiv e-prints*, arXiv:1206.5538.
- Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *35*(8), 1798–1828. doi: [10.1109/tpami.2013.50](https://doi.org/10.1109/tpami.2013.50)
- Bergmann, T. O., Mölle, M., Diedrichs, J., Born, J., & Siebner, H. R. (2012). Sleep spindle-related reactivation of category-specific cortical regions after learning face-scene associations. *NeuroImage*, *59*(3), 2733–2742. doi: [10.1016/j.neuroimage.2011.10.036](https://doi.org/10.1016/j.neuroimage.2011.10.036)
- Bi, G.-q., & Poo, M.-m. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of Neuroscience*, *18*(24), 10464–10472. doi: [10.1523/jneurosci.18-24-10464.1998](https://doi.org/10.1523/jneurosci.18-24-10464.1998)
- Bornstein, A. M., Khaw, M. W., Shohamy, D., & Daw, N. D. (2017). Reminders of past choices bias decisions for reward in humans. *Nature Communications*, *8*(15958). doi: [10.1038/ncomms15958](https://doi.org/10.1038/ncomms15958)
- Bosch, S. E., Jehee, J. F. M., Fernández, G., & Doeller, C. F. (2014). Reinstatement of associative memories in early visual cortex is signaled by the hippocampus. *Journal of Neuroscience*, *34*(22), 7493–7500. doi: [10.1523/JNEUROSCI.0805-14.2014](https://doi.org/10.1523/JNEUROSCI.0805-14.2014)
- Bostock, E., Muller, R. U., & Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, *1*(2), 193–205. doi: <https://doi.org/10.1002/hipo.450010207>
- Botvinick, M., Ritter, S., Wang, J. X., Kurth-Nelson, Z., Blundell, C., & Hassabis, D. (2019). Reinforcement learning, fast and slow. *Trends in Cognitive Sciences*, *23*(5), 408–422. doi: <https://doi.org/10.1016/j.tics.2019.02.006>

- Brunec, I. K., & Momennejad, I. (2021). Predictive representations in hippocampal and prefrontal hierarchies. *Journal of Neuroscience*. doi: [10.1523/JNEUROSCI.1327-21.2021](https://doi.org/10.1523/JNEUROSCI.1327-21.2021)
- Buckner, R. L. (2010). The role of the hippocampus in prediction and imagination. *Annual Review of Psychology*, *61*(1), 27–48. doi: [10.1146/annurev.psych.60.110707.163508](https://doi.org/10.1146/annurev.psych.60.110707.163508)
- Buzsáki, G. (1989). Two-stage model of memory trace formation: A role for “noisy” brain states. *Neuroscience*, *31*(3), 551–570. doi: [10.1016/0306-4522\(89\)90423-5](https://doi.org/10.1016/0306-4522(89)90423-5)
- Buzsáki, G., & Llinás, R. (2017). Space and time in the brain. *Science*, *358*(6362), 482–485. doi: [10.1126/science.aan8869](https://doi.org/10.1126/science.aan8869)
- Cai, M. B., Schuck, N. W., Pillow, J. W., & Niv, Y. (2019). Representational structure or task structure? Bias in neural representational similarity analysis and a Bayesian method for reducing bias. *PLOS Computational Biology*, *15*(5), e1006299. doi: [10.1371/journal.pcbi.1006299](https://doi.org/10.1371/journal.pcbi.1006299)
- Cairney, S. A., Guttesen, A. á. V., El Marj, N., & Staresina, B. P. (2018). Memory consolidation is linked to spindle-mediated information processing during sleep. *Current Biology*, *28*(6), 948–954.e4. doi: <https://doi.org/10.1016/j.cub.2018.01.087>
- Carey, A. A., Tanaka, Y., & van der Meer, M. A. A. (2019). Reward revaluation biases hippocampal replay content away from the preferred outcome. *Nature Neuroscience*, *22*(9), 1450–1459. doi: [10.1038/s41593-019-0464-6](https://doi.org/10.1038/s41593-019-0464-6)
- Carr, M. F., Jadhav, S. P., & Frank, L. M. (2011). Hippocampal replay in the awake state: A potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, *14*(2), 147–153. doi: [10.1038/nn.2732](https://doi.org/10.1038/nn.2732)
- Cazé, R., Khamassi, M., Aubin, L., & Girard, B. (2018). Hippocampal replays under the scrutiny of reinforcement learning models. *Journal of Neurophysiology*, *120*(6), 2877–2896. doi: [10.1152/jn.00145.2018](https://doi.org/10.1152/jn.00145.2018)
- Cazin, N., Llofriu Alonso, M., Scleidorovich Chiodi, P., Pelc, T., Harland, B., Weitzenfeld, A., ... Dominey, P. F. (2019). Reservoir computing model of prefrontal cortex creates novel combinations of previous navigation sequences from hippocampal place-cell replay with spatial reward propagation. *PLOS Computational Biology*, *15*(7), e1006624. doi: [10.1371/journal.pcbi.1006624](https://doi.org/10.1371/journal.pcbi.1006624)
- Clemens, Z., Molle, M., Eross, L., Barsi, P., Halasz, P., & Born, J. (2007). Temporal coupling of parahippocampal ripples, sleep spindles and slow oscillations in humans. *Brain*, *130*(11), 2868–2878. doi: [10.1093/brain/awm146](https://doi.org/10.1093/brain/awm146)
- Colgin, L. L., Moser, E. I., & Moser, M.-B. (2008). Understanding memory through hippocampal remapping. *Trends in Neurosciences*, *31*(9), 469–477. doi: <https://doi.org/10.1016/j.tins.2008.06.008>
- Corneil, D. S., & Gerstner, W. (2015). Attractor network dynamics enable preplay and rapid path planning in mazelike environments. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, & R. Garnett (Eds.), *Advances in neural information processing systems* (Vol. 28), Curran Associates, Inc.
- Csicsvari, J., O’Neill, J., Allen, K., & Senior, T. (2007). Place-selective firing contributes to the reverse-order reactivation of CA1 pyramidal cells during sharp waves in open-field exploration. *European Journal of Neuroscience*, *26*(3), 704–716. doi: [10.1111/j.1460-9568.2007.05684.x](https://doi.org/10.1111/j.1460-9568.2007.05684.x)

- Davidson, T. J., Kloosterman, F., & Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron*, *63*(4), 497–507. doi: [10.1016/j.neuron.2009.07.027](https://doi.org/10.1016/j.neuron.2009.07.027)
- Dayan, P. (1993). Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, *5*(4), 613–624. doi: [10.1162/neco.1993.5.4.613](https://doi.org/10.1162/neco.1993.5.4.613)
- de Voogd, L. D., Fernández, G., & Hermans, E. J. (2016). Awake reactivation of emotional memory traces through hippocampal–neocortical interactions. *NeuroImage*, *134*, 563–572. doi: [10.1016/j.neuroimage.2016.04.026](https://doi.org/10.1016/j.neuroimage.2016.04.026)
- Deng, X., Chen, S., Sosa, M., Karlsson, M. P., Wei, X.-X., & Frank, L. M. (2020). A variable clock underlies internally generated hippocampal sequences. *bioRxiv*. doi: [10.1101/2020.04.10.035980](https://doi.org/10.1101/2020.04.10.035980)
- Denovellis, E. L., Gillespie, A. K., Coulter, M. E., Sosa, M., Chung, J. E., Eden, U. T., & Frank, L. M. (2021). Hippocampal replay of experience at real-world speeds. *eLife*, *10*, e64505. doi: [10.7554/eLife.64505](https://doi.org/10.7554/eLife.64505)
- Deuker, L., Olligs, J., Fell, J., Kranz, T. A., Mormann, F., Montag, C., ... Axmacher, N. (2013). Memory consolidation by replay of stimulus-specific neural activity. *Journal of Neuroscience*, *33*(49), 19373–19383. doi: [10.1523/jneurosci.0414-13.2013](https://doi.org/10.1523/jneurosci.0414-13.2013)
- Diana, R. A., Yonelinas, A. P., & Ranganath, C. (2008). High-resolution multi-voxel pattern analysis of category selectivity in the medial temporal lobes. *Hippocampus*, *18*(6), 536–541. doi: [10.1002/hipo.20433](https://doi.org/10.1002/hipo.20433)
- Diba, K., & Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, *10*(10), 1241–1242. doi: [10.1038/mm1961](https://doi.org/10.1038/mm1961)
- Dragoi, G. (2020). Cell assemblies, sequences and temporal coding in the hippocampus. *Current Opinion in Neurobiology*, *64*, 111–118. doi: [10.1016/j.conb.2020.03.003](https://doi.org/10.1016/j.conb.2020.03.003)
- Dragoi, G., & Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, *469*(7330), 397–401. doi: [10.1038/nature09633](https://doi.org/10.1038/nature09633)
- Dragoi, G., & Tonegawa, S. (2013a). Development of schemas revealed by prior experience and NMDA receptor knock-out. *eLife*, *2*, e01326. doi: [10.7554/eLife.01326](https://doi.org/10.7554/eLife.01326)
- Dragoi, G., & Tonegawa, S. (2013b). Distinct preplay of multiple novel spatial experiences in the rat. *Proceedings of the National Academy of Sciences*, *110*(22), 9100–9105. doi: [10.1073/pnas.1306031110](https://doi.org/10.1073/pnas.1306031110)
- Drieu, C., & Zugaro, M. (2019). Hippocampal sequences during exploration: Mechanisms and functions. *Frontiers in Cellular Neuroscience*, *13*. doi: [10.3389/fncel.2019.00232](https://doi.org/10.3389/fncel.2019.00232)
- Eckert, M. J., McNaughton, B. L., & Tatsuno, M. (2020). Neural ensemble reactivation in rapid eye movement and slow-wave sleep coordinate with muscle activity to promote rapid motor skill learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *375*(1799), 20190655. doi: [10.1098/rstb.2019.0655](https://doi.org/10.1098/rstb.2019.0655)
- Eichenbaum, H. (2014). Time cells in the hippocampus: A new dimension for mapping memories. *Nature Reviews Neuroscience*, *15*(11), 732–744. doi: [10.1038/nrn3827](https://doi.org/10.1038/nrn3827)
- Eichenbaum, H. (2015). Does the hippocampus preplay memories? *Nature Neuroscience*, *18*(12), 1701–1702. doi: [10.1038/nn.4180](https://doi.org/10.1038/nn.4180)

- Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., & Tanila, H. (1999). The hippocampus, memory, and place cells: Is it spatial memory or a memory space? *Neuron*, *23*(2), 209–226. doi: [https://doi.org/10.1016/S0896-6273\(00\)80773-4](https://doi.org/10.1016/S0896-6273(00)80773-4)
- Ekman, M., Kok, P., & de Lange, F. P. (2017). Time-compressed preplay of anticipated events in human primary visual cortex. *Nature Communications*, *8*(15276), 1–9. doi: [10.1038/ncomms15276](https://doi.org/10.1038/ncomms15276)
- Eldar, E., Lièvre, G., Dayan, P., & Dolan, R. J. (2020). The roles of online and offline replay in planning. *eLife*, *9*. doi: [10.7554/elife.56911](https://doi.org/10.7554/elife.56911)
- Engel, A. K., Moll, C. K. E., Fried, I., & Ojemann, G. A. (2005). Invasive recordings from the human brain: Clinical insights and beyond. *Nature Reviews Neuroscience*, *6*(1), 35–47. doi: [10.1038/nrn1585](https://doi.org/10.1038/nrn1585)
- Esteban, O., Birman, D., Schaer, M., Koyejo, O. O., Poldrack, R. A., & Gorgolewski, K. J. (2017). MRIQC: Advancing the automatic prediction of image quality in MRI from unseen sites. *PLoS ONE*, *12*(9), e0184661. doi: [10.1371/journal.pone.0184661](https://doi.org/10.1371/journal.pone.0184661)
- Esteban, O., Markiewicz, C. J., Blair, R. W., Moodie, C. A., Isik, A. I., Erramuzpe, A., . . . et al. (2018). fMRIPrep: A robust preprocessing pipeline for functional MRI. *Nature Methods*, *16*(1), 111–116. doi: [10.1038/s41592-018-0235-4](https://doi.org/10.1038/s41592-018-0235-4)
- Euston, D. R., Tatsuno, M., & McNaughton, B. L. (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science*, *318*(5853), 1147–1150. doi: [10.1126/science.1148979](https://doi.org/10.1126/science.1148979)
- Farooq, U., Sibille, J., Liu, K., & Dragoi, G. (2019). Strengthened temporal coordination within pre-existing sequential cell assemblies supports trajectory replay. *Neuron*, *103*(4), 719–733.e7. doi: [10.1016/j.neuron.2019.05.040](https://doi.org/10.1016/j.neuron.2019.05.040)
- Finn, E. S. (2021). Is it time to put rest to rest? *Trends in Cognitive Sciences*. doi: <https://doi.org/10.1016/j.tics.2021.09.005>
- Foster, D. J. (2017). Replay comes of age. *Annual Review of Neuroscience*, *40*(1), 581–602. doi: [10.1146/annurev-neuro-072116-031538](https://doi.org/10.1146/annurev-neuro-072116-031538)
- Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, *440*(7084), 680–683. doi: [10.1038/nature04587](https://doi.org/10.1038/nature04587)
- Foster, D. J., & Wilson, M. A. (2007). Hippocampal theta sequences. *Hippocampus*, *17*(11), 1093–1099. doi: [10.1002/hipo.20345](https://doi.org/10.1002/hipo.20345)
- French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, *3*(4), 128–135. doi: [10.1016/s1364-6613\(99\)01294-2](https://doi.org/10.1016/s1364-6613(99)01294-2)
- Garvert, M. M., Dolan, R. J., & Behrens, T. E. (2017). A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *eLife*, *6*. doi: [10.7554/elife.17086](https://doi.org/10.7554/elife.17086)
- Genzel, L., Dragoi, G., Frank, L., Ganguly, K., de la Prida, L., Pfeiffer, B., & Robertson, E. (2020). A consensus statement: Defining terms for reactivation analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *375*(1799), 20200001. doi: [10.1098/rstb.2020.0001](https://doi.org/10.1098/rstb.2020.0001)
- Gerrard, J. L., Kudrimoti, H., McNaughton, B. L., & Barnes, C. A. (2001). Reactivation of hippocampal ensemble activity patterns in the aging rat. *Behavioral Neuroscience*, *115*(6), 1180–1192. doi: [10.1037/0735-7044.115.6.1180](https://doi.org/10.1037/0735-7044.115.6.1180)

- Gershman, S. J., & Daw, N. D. (2017). Reinforcement learning and episodic memory in humans and animals: An integrative framework. *Annual Review of Psychology*, *68*(1), 101–128. doi: [10.1146/annurev-psych-122414-033625](https://doi.org/10.1146/annurev-psych-122414-033625)
- Gordon, A. M., Rissman, J., Kiani, R., & Wagner, A. D. (2013). Cortical reinstatement mediates the relationship between content-specific encoding activity and subsequent recollection decisions. *Cerebral Cortex*, *24*(12), 3350–3364. doi: [10.1093/cercor/bht194](https://doi.org/10.1093/cercor/bht194)
- Gorgolewski, K. J., Auer, T., Calhoun, V. D., Craddock, R. C., Das, S., Duff, E. P., . . . et al. (2016). The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Scientific Data*, *3*(160044). doi: [10.1038/sdata.2016.44](https://doi.org/10.1038/sdata.2016.44)
- Grosmark, A. D., & Buzsáki, G. (2016). Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science*, *351*(6280), 1440–1443. doi: [10.1126/science.aad1935](https://doi.org/10.1126/science.aad1935)
- Gruber, M. J., Ritchey, M., Wang, S.-F., Doss, M. K., & Ranganath, C. (2016). Post-learning hippocampal dynamics promote preferential retention of rewarding events. *Neuron*, *89*(5), 1110–1120. doi: [10.1016/j.neuron.2016.01.017](https://doi.org/10.1016/j.neuron.2016.01.017)
- Gupta, A. S., van der Meer, M. A., Touretzky, D. S., & Redish, A. D. (2010). Hippocampal replay is not a simple function of experience. *Neuron*, *65*(5), 695–705. doi: [10.1016/j.neuron.2010.01.034](https://doi.org/10.1016/j.neuron.2010.01.034)
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–806. doi: [10.1038/nature03721](https://doi.org/10.1038/nature03721)
- Halchenko, Y. O., Hanke, M., Poldrack, B., Meyer, K., Solanky, D. S., Alteva, G., . . . Markiewicz, C. J. (2019). Datalad/datalad 0.11.5. doi: [10.5281/ZENODO.3233911](https://doi.org/10.5281/ZENODO.3233911)
- Halchenko, Y. O., Meyer, K., Poldrack, B., Solanky, D. S., Wagner, A. S., Gors, J., . . . Hanke, M. (2021). DataLad: Distributed system for joint management of code, data, and their relationship. *Journal of Open Source Software*, *6*(63), 3262. doi: [10.21105/joss.03262](https://doi.org/10.21105/joss.03262)
- Hassabis, D., Chu, C., Rees, G., Weiskopf, N., Molyneux, P. D., & Maguire, E. A. (2009). Decoding neuronal ensembles in the human hippocampus. *Current Biology*, *19*(7), 546–554. doi: <https://doi.org/10.1016/j.cub.2009.02.033>
- Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. *Neuron*, *95*(2), 245–258. doi: [10.1016/j.neuron.2017.06.011](https://doi.org/10.1016/j.neuron.2017.06.011)
- Hayes, T. L., Krishnan, G. P., Bazhenov, M., Siegelmann, H. T., Sejnowski, T. J., & Kanan, C. (2021). Replay in deep learning: Current approaches and missing biological elements. *arXiv e-prints*, arXiv:2104.04132.
- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. John Wiley and Sons, Inc.
- Heeger, D. J., & Ress, D. (2002). What does fMRI tell us about neuronal activity? *Nature Reviews Neuroscience*, *3*(2), 142–151. doi: [10.1038/nrn730](https://doi.org/10.1038/nrn730)
- Helfrich, R. F., Lendner, J. D., Mander, B. A., Guillen, H., Paff, M., Mnatsakanyan, L., . . . Knight, R. T. (2019). Bidirectional prefrontal-hippocampal dynamics organize information transfer during sleep in humans. *Nature Communications*, *10*(1). doi: [10.1038/s41467-019-11444-x](https://doi.org/10.1038/s41467-019-11444-x)



- Hermans, E. J., Kanen, J. W., Tambini, A., Fernández, G., Davachi, L., & Phelps, E. A. (2016). Persistence of amygdala–hippocampal connectivity and multi-voxel correlation structures during awake rest after fear learning predicts long-term expression of fear. *Cerebral Cortex*, bhw145. doi: [10.1093/cercor/bhw145](https://doi.org/10.1093/cercor/bhw145)
- Higgins, C., Liu, Y., Vidaurre, D., Kurth-Nelson, Z., Dolan, R., Behrens, T., & Woolrich, M. (2020). Replay bursts in humans coincide with activation of the default mode and parietal alpha networks. *Neuron*. doi: [10.1016/j.neuron.2020.12.007](https://doi.org/10.1016/j.neuron.2020.12.007)
- Horgan, D., Quan, J., Budden, D., Barth-Marion, G., Hessel, M., van Hasselt, H., & Silver, D. (2018). Distributed prioritized experience replay. *arXiv e-prints*, arXiv:1803.00933.
- Jackson, J. C., Johnson, A., & Redish, A. D. (2006). Hippocampal sharp waves and reactivation during awake states depend on repeated sequential experience. *Journal of Neuroscience*, 26(48), 12415–12426. doi: [10.1523/jneurosci.4118-06.2006](https://doi.org/10.1523/jneurosci.4118-06.2006)
- Jadhav, S. P., Rothschild, G., Roumis, D. K., & Frank, L. M. (2016). Coordinated excitation and inhibition of prefrontal ensembles during awake hippocampal sharp-wave ripple events. *Neuron*, 90(1), 113–127. doi: [10.1016/j.neuron.2016.02.010](https://doi.org/10.1016/j.neuron.2016.02.010)
- Jeffery, K. J. (2011). Place cells, grid cells, attractors, and remapping. *Neural Plasticity*, 2011, 182602. doi: [10.1155/2011/182602](https://doi.org/10.1155/2011/182602)
- Ji, D., & Wilson, M. A. (2006). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience*, 10(1), 100–107. doi: [10.1038/mm1825](https://doi.org/10.1038/mm1825)
- Johnson, A., & Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *Journal of Neuroscience*, 27(45), 12176–12189. doi: [10.1523/jneurosci.3761-07.2007](https://doi.org/10.1523/jneurosci.3761-07.2007)
- Kaefer, K., Nardin, M., Blahna, K., & Csicsvari, J. (2020). Replay of behavioral sequences in the medial prefrontal cortex during rule switching. *Neuron*, 106(1), 154–165.e6. doi: [10.1016/j.neuron.2020.01.015](https://doi.org/10.1016/j.neuron.2020.01.015)
- Karlsson, M. P., & Frank, L. M. (2009). Awake replay of remote experiences in the hippocampus. *Nature Neuroscience*, 12(7), 913–918. doi: [10.1038/nm.2344](https://doi.org/10.1038/nm.2344)
- Kay, K., Chung, J. E., Sosa, M., Schor, J. S., Karlsson, M. P., Larkin, M. C., . . . Frank, L. M. (2020). Constant sub-second cycling between representations of possible futures in the hippocampus. *Cell*, 180(3), 552–567.e25. doi: [10.1016/j.cell.2020.01.014](https://doi.org/10.1016/j.cell.2020.01.014)
- Khamassi, M., & Girard, B. (2020). Modeling awake hippocampal reactivations with model-based bidirectional search. *Biological Cybernetics*, 114, 231–248. doi: [10.1007/s00422-020-00817-x](https://doi.org/10.1007/s00422-020-00817-x)
- King, C., Henze, D. A., Leinekugel, X., & Buzsáki, G. (1999). Hebbian modification of a hippocampal population pattern in the rat. *The Journal of Physiology*, 521(1), 159–167. doi: [10.1111/j.1469-7793.1999.00159.x](https://doi.org/10.1111/j.1469-7793.1999.00159.x)
- Kudrimoti, H. S., Barnes, C. A., & McNaughton, B. L. (1999). Reactivation of hippocampal cell assemblies: Effects of behavioral state, experience, and EEG dynamics. *The Journal of Neuroscience*, 19(10), 4090–4101. doi: [10.1523/jneurosci.19-10-04090.1999](https://doi.org/10.1523/jneurosci.19-10-04090.1999)
- Kumaran, D., Hassabis, D., & McClelland, J. L. (2016). What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends in Cognitive Sciences*, 20(7), 512–534. doi: [10.1016/j.tics.2016.05.004](https://doi.org/10.1016/j.tics.2016.05.004)

- Kurth-Nelson, Z., Barnes, G., Sejdinovic, D., Dolan, R., & Dayan, P. (2015). Temporal structure in associative retrieval. *eLife*, *4*. doi: [10.7554/elife.04919](https://doi.org/10.7554/elife.04919)
- Kurth-Nelson, Z., Economides, M., Dolan, R. J., & Dayan, P. (2016). Fast sequences of non-spatial state representations in humans. *Neuron*, *91*(1), 194–204. doi: [10.1016/j.neuron.2016.05.028](https://doi.org/10.1016/j.neuron.2016.05.028)
- Kwong, K. K., Belliveau, J. W., Chesler, D. A., Goldberg, I. E., Weisskoff, R. M., Poncelet, B. P., ... Turner, R. (1992). Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences*, *89*(12), 5675–5679. doi: [10.1073/pnas.89.12.5675](https://doi.org/10.1073/pnas.89.12.5675)
- Kyle, C. T., Stokes, J. D., Lieberman, J. S., Hassan, A. S., & Ekstrom, A. D. (2015). Successful retrieval of competing spatial environments in humans involves hippocampal pattern separation mechanisms. *eLife*, *4*, e10499. doi: [10.7554/eLife.10499](https://doi.org/10.7554/eLife.10499)
- Lansink, C. S., Goltstein, P. M., Lankelma, J. V., Joosten, R. N. J. M. A., McNaughton, B. L., & Pennartz, C. M. A. (2008). Preferential reactivation of motivationally relevant information in the ventral striatum. *Journal of Neuroscience*, *28*(25), 6372–6382. doi: [10.1523/jneurosci.1054-08.2008](https://doi.org/10.1523/jneurosci.1054-08.2008)
- Lansink, C. S., Goltstein, P. M., Lankelma, J. V., McNaughton, B. L., & Pennartz, C. M. A. (2009). Hippocampus leads ventral striatum in replay of place-reward information. *PLoS Biology*, *7*(8), e1000173. doi: [10.1371/journal.pbio.1000173](https://doi.org/10.1371/journal.pbio.1000173)
- Lee, A. K., & Wilson, M. A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*, *36*(6), 1183–1194. doi: [10.1016/s0896-6273\(02\)01096-6](https://doi.org/10.1016/s0896-6273(02)01096-6)
- Lee, S. W., O’Doherty, J. P., & Shimojo, S. (2015). Neural computations mediating one-shot learning in the human brain. *PLOS Biology*, *13*(4), 1–36. doi: [10.1371/journal.pbio.1002137](https://doi.org/10.1371/journal.pbio.1002137)
- Leibold, C. (2020). A model for navigation in unknown environments based on a reservoir of hippocampal sequences. *Neural Networks*, *124*, 328–342. doi: [10.1016/j.neunet.2020.01.014](https://doi.org/10.1016/j.neunet.2020.01.014)
- Lin, L. J. (1991). Programming robots using reinforcement learning and teaching. In *Association for the advancement of artificial intelligence* (pp. 781–786).
- Liu, Y., Dolan, R. J., Kurth-Nelson, Z., & Behrens, T. E. (2019). Human replay spontaneously reorganizes experience. *Cell*, *178*(3), 640–652. doi: [10.1016/j.cell.2019.06.012](https://doi.org/10.1016/j.cell.2019.06.012)
- Liu, Y., Dolan, R. J., Higgins, C., Penagos, H., Woolrich, M. W., Ólafsdóttir, H. F., ... Behrens, T. E. (2021a). Temporally delayed linear modelling (TDLM) measures replay in both animals and humans. *eLife*, *10*, e66917. doi: [10.7554/eLife.66917](https://doi.org/10.7554/eLife.66917)
- Liu, Y., Mattar, M. G., Behrens, T. E. J., Daw, N. D., & Dolan, R. J. (2021b). Experience replay is associated with efficient nonlocal learning. *Science*, *372*(6544). doi: [10.1126/science.abf1357](https://doi.org/10.1126/science.abf1357)
- Louie, K., & Wilson, M. A. (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron*, *29*(1), 145–156. doi: [10.1016/s0896-6273\(01\)00186-6](https://doi.org/10.1016/s0896-6273(01)00186-6)
- Lynn, C. W., Kahn, A. E., Nyema, N., & Bassett, D. S. (2020). Abstract representations of events arise from mental errors in learning and memory. *Nature Communications*, *11*(1). doi: [10.1038/s41467-020-15146-7](https://doi.org/10.1038/s41467-020-15146-7)

- MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron*, *71*(4), 737–749. doi: [10.1016/j.neuron.2011.07.012](https://doi.org/10.1016/j.neuron.2011.07.012)
- Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *262*(841), 23–81. doi: [10.1098/rstb.1971.0078](https://doi.org/10.1098/rstb.1971.0078)
- Mattar, M. G., & Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neuroscience*, *21*(11), 1609–1617. doi: [10.1038/s41593-018-0232-z](https://doi.org/10.1038/s41593-018-0232-z)
- McClelland, J. L., McNaughton, B. L., & O’Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*(3), 419–457. doi: [10.1037/0033-295x.102.3.419](https://doi.org/10.1037/0033-295x.102.3.419)
- McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation*, *24*, 109–165. doi: [10.1016/s0079-7421\(08\)60536-8](https://doi.org/10.1016/s0079-7421(08)60536-8)
- McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M.-B. (2006). Path integration and the neural basis of the “cognitive map”. *Nature Reviews Neuroscience*, *7*(8), 663–678. doi: [10.1038/nrn1932](https://doi.org/10.1038/nrn1932)
- Michelmann, S., Staresina, B. P., Bowman, H., & Hanslmayr, S. (2018). Speed of time-compressed forward replay flexibly changes in human episodic memory. *Nature Human Behaviour*, *3*(2), 143–154. doi: [10.1038/s41562-018-0491-4](https://doi.org/10.1038/s41562-018-0491-4)
- Mildner, J. N., & Tamir, D. I. (2019). Spontaneous thought as an unconstrained memory process. *Trends in Neurosciences*, *42*(11), 763–777. doi: [10.1016/j.tins.2019.09.001](https://doi.org/10.1016/j.tins.2019.09.001)
- Minsky, M. (1961). Steps toward artificial intelligence. *Proceedings of the IRE*, *49*(1), 8–30. doi: [10.1109/JRPROC.1961.287775](https://doi.org/10.1109/JRPROC.1961.287775)
- Miyashita, Y. (1988). Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature*, *335*(6193), 817–820. doi: [10.1038/335817a0](https://doi.org/10.1038/335817a0)
- Mizunuma, M., Norimoto, H., Tao, K., Egawa, T., Hanaoka, K., Sakaguchi, T., . . . et al. (2014). Unbalanced excitability underlies offline reactivation of behaviorally activated neurons. *Nature Neuroscience*, *17*(4), 503–505. doi: [10.1038/nn.3674](https://doi.org/10.1038/nn.3674)
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., . . . et al. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529–533. doi: [10.1038/nature14236](https://doi.org/10.1038/nature14236)
- Moita, M. A. P. (2004). Putting fear in its place: Remapping of hippocampal place cells during fear conditioning. *Journal of Neuroscience*, *24*(31), 7015–7023. doi: [10.1523/jneurosci.5492-03.2004](https://doi.org/10.1523/jneurosci.5492-03.2004)
- Momennejad, I. (2020). Learning structures: Predictive representations, replay, and generalization. *Current Opinion in Behavioral Sciences*, *32*, 155–166. doi: [10.1016/j.cobeha.2020.02.017](https://doi.org/10.1016/j.cobeha.2020.02.017)
- Momennejad, I., & Howard, M. W. (2018). Predicting the future with multi-scale successor representations. *bioRxiv*. doi: [10.1101/449470](https://doi.org/10.1101/449470)

- Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N. D., & Gershman, S. J. (2017). The successor representation in human reinforcement learning. *Nature Human Behaviour*, 1(9), 680–692. doi: [10.1038/s41562-017-0180-8](https://doi.org/10.1038/s41562-017-0180-8)
- Moser, E. I., Kropff, E., & Moser, M.-B. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annual Review of Neuroscience*, 31(1), 69–89. doi: [10.1146/annurev.neuro.31.061307.090723](https://doi.org/10.1146/annurev.neuro.31.061307.090723)
- Murty, V. P., Tompariy, A., Adcock, R. A., & Davachi, L. (2017). Selectivity in postencoding connectivity with high-level visual cortex is associated with reward-motivated memory. *The Journal of Neuroscience*, 37(3), 537–545. doi: [10.1523/jneurosci.4032-15.2016](https://doi.org/10.1523/jneurosci.4032-15.2016)
- Nádasy, Z., Hirase, H., Czurkó, A., Csicsvari, J., & Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *The Journal of Neuroscience*, 19(21), 9497–9507. doi: [10.1523/jneurosci.19-21-09497.1999](https://doi.org/10.1523/jneurosci.19-21-09497.1999)
- Niv, Y. (2019). Learning task-state representations. *Nature Neuroscience*, 22(10), 1544–1553. doi: [10.1038/s41593-019-0470-8](https://doi.org/10.1038/s41593-019-0470-8)
- Nour, M. M., Liu, Y., Arumuham, A., Kurth-Nelson, Z., & Dolan, R. J. (2021). Impaired neural replay of inferred relationships in schizophrenia. *Cell*. doi: [10.1016/j.cell.2021.06.012](https://doi.org/10.1016/j.cell.2021.06.012)
- O'Keefe, J. (1979). A review of the hippocampal place cells. *Progress in Neurobiology*, 13(4), 419–439. doi: [10.1016/0301-0082\(79\)90005-4](https://doi.org/10.1016/0301-0082(79)90005-4)
- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34(1), 171–175. doi: [10.1016/0006-8993\(71\)90358-1](https://doi.org/10.1016/0006-8993(71)90358-1)
- O'Keefe, J., & Krupic, J. (2021). Do hippocampal pyramidal cells respond to non-spatial stimuli? *Physiological Reviews*, 101(3), 1427–1456. doi: [10.1152/physrev.00014.2020](https://doi.org/10.1152/physrev.00014.2020)
- O'Keefe, J., & Nadel, L. (1974). Maps in the brain. *New Scientist*, 62(903), 749–751.
- O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.
- O'Neill, J., Boccara, C. N., Stella, F., Schoenenberger, P., & Csicsvari, J. (2017). Superficial layers of the medial entorhinal cortex replay independently of the hippocampus. *Science*, 355(6321), 184–188. doi: [10.1126/science.aag2787](https://doi.org/10.1126/science.aag2787)
- O'Neill, J., Senior, T., & Csicsvari, J. (2006). Place-selective firing of ca1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. *Neuron*, 49(1), 143–155. doi: <https://doi.org/10.1016/j.neuron.2005.10.037>
- O'Reilly, R. C., Bhattacharyya, R., Howard, M. D., & Ketz, N. (2014). Complementary learning systems. *Cognitive Science*, 38(6), 1229–1248. doi: [10.1111/j.1551-6709.2011.01214.x](https://doi.org/10.1111/j.1551-6709.2011.01214.x)
- Ogawa, S., Lee, T. M., Kay, A. R., & Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences*, 87(24), 9868–9872. doi: [10.1073/pnas.87.24.9868](https://doi.org/10.1073/pnas.87.24.9868)
- Ólafsdóttir, H. F., Barry, C., Saleem, A. B., Hassabis, D., & Spiers, H. J. (2015). Hippocampal place cells construct reward related sequences through unexplored space. *eLife*, 4. doi: [10.7554/elife.06063](https://doi.org/10.7554/elife.06063)
- Ólafsdóttir, H. F., Carpenter, F., & Barry, C. (2016). Coordinated grid and place cell replay during rest. *Nature Neuroscience*, 19(6), 792–794. doi: [10.1038/nn.4291](https://doi.org/10.1038/nn.4291)

- Ólafsdóttir, H. F., Carpenter, F., & Barry, C. (2017). Task demands predict a dynamic switch in the content of awake hippocampal replay. *Neuron*, *96*(4), 925–935.e6. doi: [10.1016/j.neuron.2017.09.035](https://doi.org/10.1016/j.neuron.2017.09.035)
- Pavlidis, C., & Winson, J. (1989). Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *The Journal of Neuroscience*, *9*(8), 2907–2918. doi: [10.1523/jneurosci.09-08-02907.1989](https://doi.org/10.1523/jneurosci.09-08-02907.1989)
- Pereira, F., Mitchell, T., & Botvinick, M. (2009). Machine learning classifiers and fMRI: A tutorial overview. *NeuroImage*, *45*(1), S199–S209. doi: [10.1016/j.neuroimage.2008.11.007](https://doi.org/10.1016/j.neuroimage.2008.11.007)
- Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S. I., & Battaglia, F. P. (2009). Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature Neuroscience*, *12*(7), 919–926. doi: [10.1038/nm.2337](https://doi.org/10.1038/nm.2337)
- Pezzulo, G., Donnarumma, F., Maisto, D., & Stoianov, I. (2019). Planning at decision time and in the background during spatial navigation. *Current Opinion in Behavioral Sciences*, *29*, 69–76. doi: [10.1016/j.cobeha.2019.04.009](https://doi.org/10.1016/j.cobeha.2019.04.009)
- Pfeiffer, B. E. (2018). The content of hippocampal “replay”. *Hippocampus*, 1–13. doi: [10.1002/hipo.22824](https://doi.org/10.1002/hipo.22824)
- Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, *497*(7447), 74–79. doi: [10.1038/nature12112](https://doi.org/10.1038/nature12112)
- Ratcliff, R. (1990). Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psychological Review*, *97*(2), 285–308. doi: [10.1037/0033-295x.97.2.285](https://doi.org/10.1037/0033-295x.97.2.285)
- Redish, A. D. (1999). *Beyond the cognitive map: From place cells to episodic memory*. MIT Press.
- Redish, A. D. (2016). Vicarious trial and error. *Nature Reviews Neuroscience*, *17*(3), 147–159. doi: [10.1038/nrn.2015.30](https://doi.org/10.1038/nrn.2015.30)
- Ritchev, M., Wing, E. A., LaBar, K. S., & Cabeza, R. (2012). Neural similarity between encoding and retrieval is related to memory via hippocampal interactions. *Cerebral Cortex*, *23*(12), 2818–2828. doi: [10.1093/cercor/bhs258](https://doi.org/10.1093/cercor/bhs258)
- Robinson, N. T. M., Descamps, L. A. L., Russell, L. E., Buchholz, M. O., Bicknell, B. A., Antonov, G. K., ... Häusser, M. (2020). Targeted activation of hippocampal place cells drives memory-guided spatial behavior. *Cell*, *183*(6), 1586–1599.e10. doi: [10.1016/j.cell.2020.09.061](https://doi.org/10.1016/j.cell.2020.09.061)
- Roscow, E. L., Chua, R., Costa, R. P., Jones, M. W., & Lepora, N. (2021). Learning offline: Memory replay in biological and artificial reinforcement learning. *Trends in Neurosciences*, *1738*. doi: [10.1016/j.tins.2021.07.007](https://doi.org/10.1016/j.tins.2021.07.007)
- Roscow, E. L., Jones, M. W., & Lepora, N. F. (2019). Behavioural and computational evidence for memory consolidation biased by reward-prediction errors. *bioRxiv*. doi: [10.1101/716290](https://doi.org/10.1101/716290)
- Rothschild, G., Eban, E., & Frank, L. M. (2016). A cortical–hippocampal–cortical loop of information processing during memory consolidation. *Nature Neuroscience*, *20*(2), 251–259. doi: [10.1038/nn.4457](https://doi.org/10.1038/nn.4457)

- Russek, E. M., Momennejad, I., Botvinick, M. M., Gershman, S. J., & Daw, N. D. (2017). Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLoS Computational Biology*, *13*(9), e1005768. doi: [10.1371/journal.pcbi.1005768](https://doi.org/10.1371/journal.pcbi.1005768)
- Russek, E. M., Momennejad, I., Botvinick, M. M., Gershman, S. J., & Daw, N. D. (2021). Neural evidence for the successor representation in choice evaluation. *bioRxiv*. doi: [10.1101/2021.08.29.458114](https://doi.org/10.1101/2021.08.29.458114)
- Sadeh, T., Chen, J., Goshen-Gottstein, Y., & Moscovitch, M. (2019). Overlap between hippocampal pre-encoding and encoding patterns supports episodic memory. *Hippocampus*, *29*(9), 836–847. doi: [10.1002/hipo.23079](https://doi.org/10.1002/hipo.23079)
- Schapiro, A. C., Kustner, L. V., & Turk-Browne, N. B. (2012). Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Current Biology*, *22*(17), 1622–1627. doi: [10.1016/j.cub.2012.06.056](https://doi.org/10.1016/j.cub.2012.06.056)
- Schapiro, A. C., McDevitt, E. A., Rogers, T. T., Mednick, S. C., & Norman, K. A. (2018). Human hippocampal replay during rest prioritizes weakly learned information and predicts memory performance. *Nature Communications*, *9*(1). doi: [10.1038/s41467-018-06213-1](https://doi.org/10.1038/s41467-018-06213-1)
- Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary learning systems within the hippocampus: A neural network modelling approach to reconciling episodic memory with statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*(1711), 20160049. doi: [10.1098/rstb.2016.0049](https://doi.org/10.1098/rstb.2016.0049)
- Schaul, T., Quan, J., Antonoglou, I., & Silver, D. (2015). Prioritized experience replay. *arXiv e-prints*, arXiv:1511.05952.
- Schlichting, M. L., & Preston, A. R. (2014). Memory reactivation during rest supports upcoming learning of related content. *Proceedings of the National Academy of Sciences*, *111*(44), 15845–15850. doi: [10.1073/pnas.1404396111](https://doi.org/10.1073/pnas.1404396111)
- Schönauer, M., Alizadeh, S., Jamalabadi, H., Abraham, A., Pawlizki, A., & Gais, S. (2017). Decoding material-specific memory reprocessing during sleep in humans. *Nature Communications*, *8*(1). doi: [10.1038/ncomms15404](https://doi.org/10.1038/ncomms15404)
- Schreiner, T., Petzka, M., Staudigl, T., & Staresina, B. P. (2021). Endogenous memory reactivation during sleep in humans is clocked by slow oscillation-spindle complexes. *Nature Communications*, *12*(1), 3112. doi: [10.1038/s41467-021-23520-2](https://doi.org/10.1038/s41467-021-23520-2)
- Schreiner, T., & Staudigl, T. (2020). Electrophysiological signatures of memory reactivation in humans. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *375*(1799), 20190293. doi: [10.1098/rstb.2019.0293](https://doi.org/10.1098/rstb.2019.0293)
- Schuck, N. W., & Niv, Y. (2019). Sequential replay of nonspatial task states in the human hippocampus. *Science*, *364*(6447), eaaw5181. doi: [10.1126/science.aaw5181](https://doi.org/10.1126/science.aaw5181)
- Schwartenbeck, P., Baram, A., Liu, Y., Mark, S., Muller, T., Dolan, R., . . . Behrens, T. (2021). Generative replay for compositional visual understanding in the prefrontal-hippocampal circuit. *bioRxiv*. doi: [10.1101/2021.06.06.447249](https://doi.org/10.1101/2021.06.06.447249)
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of neurology, neurosurgery, and psychiatry*, *20*(1), 11–21. doi: [10.1136/jnnp.20.1.11](https://doi.org/10.1136/jnnp.20.1.11)

- Shin, J. D., Tang, W., & Jadhav, S. P. (2019). Dynamics of awake hippocampal-prefrontal replay for spatial learning and memory-guided decision making. *Neuron*, *104*(6), 1110–1125. doi: [10.1016/j.neuron.2019.09.012](https://doi.org/10.1016/j.neuron.2019.09.012)
- Silva, D., Feng, T., & Foster, D. J. (2015). Trajectory events across hippocampal place cells require previous experience. *Nature Neuroscience*, *18*(12), 1772–1779. doi: [10.1038/nn.4151](https://doi.org/10.1038/nn.4151)
- Singer, A. C., Carr, M. F., Karlsson, M. P., & Frank, L. M. (2013). Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. *Neuron*, *77*(6), 1163–1173. doi: [10.1016/j.neuron.2013.01.027](https://doi.org/10.1016/j.neuron.2013.01.027)
- Singer, A. C., & Frank, L. M. (2009). Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron*, *64*(6), 910–921. doi: [10.1016/j.neuron.2009.11.016](https://doi.org/10.1016/j.neuron.2009.11.016)
- Skaggs, W. E., & McNaughton, B. L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, *271*(5257), 1870–1873. doi: [10.1126/science.271.5257.1870](https://doi.org/10.1126/science.271.5257.1870)
- Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*, *99*(2), 195–231. doi: [10.1037/0033-295x.99.2.195](https://doi.org/10.1037/0033-295x.99.2.195)
- Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, *20*(11), 1643–1653. doi: [10.1038/nm.4650](https://doi.org/10.1038/nm.4650)
- Staresina, B. P., Alink, A., Kriegeskorte, N., & Henson, R. N. (2013). Awake reactivation predicts memory in humans. *Proceedings of the National Academy of Sciences*, *110*(52), 21159–21164. doi: [10.1073/pnas.1311989110](https://doi.org/10.1073/pnas.1311989110)
- Staresina, B. P., Henson, R. N. A., Kriegeskorte, N., & Alink, A. (2012). Episodic reinstatement in the medial temporal lobe. *Journal of Neuroscience*, *32*(50), 18150–18156. doi: [10.1523/JNEUROSCI.4156-12.2012](https://doi.org/10.1523/JNEUROSCI.4156-12.2012)
- Staresina, B. P., & Wimber, M. (2019). A neural chronometry of memory recall. *Trends in Cognitive Sciences*, *23*(12), 1071–1085. doi: [10.1016/j.tics.2019.09.011](https://doi.org/10.1016/j.tics.2019.09.011)
- Staresina, B. P., Bergmann, T. O., Bonnefond, M., van der Meij, R., Jensen, O., Deuker, L., ... Fell, J. (2015). Hierarchical nesting of slow oscillations, spindles and ripples in the human hippocampus during sleep. *Nature Neuroscience*, *18*(11), 1679–1686. doi: [10.1038/nm.4119](https://doi.org/10.1038/nm.4119)
- Steemers, B., Vicente-Grabovetsky, A., Barry, C., Smulders, P., Schröder, T. N., Burgess, N., & Doeller, C. F. (2016). Hippocampal attractor dynamics predict memory-based decision making. *Current Biology*, *26*(13), 1750–1757. doi: <https://doi.org/10.1016/j.cub.2016.04.063>
- Stella, F., Baracska, P., O'Neill, J., & Csicsvari, J. (2019). Hippocampal reactivation of random trajectories resembling brownian diffusion. *Neuron*, *102*(2), 450–461. doi: [10.1016/j.neuron.2019.01.052](https://doi.org/10.1016/j.neuron.2019.01.052)
- Stoianov, I., Maisto, D., & Pezzulo, G. (2020). The hippocampal formation as a hierarchical generative model supporting generative replay and continual learning. *bioRxiv*. doi: [10.1101/2020.01.16.908889](https://doi.org/10.1101/2020.01.16.908889)
- Sutherland, G. R., & McNaughton, B. (2000). Memory trace reactivation in hippocampal and neocortical neuronal ensembles. *Current Opinion in Neurobiology*, *10*(2), 180–186. doi: [10.1016/S0959-4388\(00\)00079-9](https://doi.org/10.1016/S0959-4388(00)00079-9)

- Sutton, R. S. (1990). Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In B. Porter & R. Mooney (Eds.), *Machine learning proceedings 1990* (pp. 216–224). doi: [10.1016/B978-1-55860-141-3.50030-4](https://doi.org/10.1016/B978-1-55860-141-3.50030-4)
- Sutton, R. S. (1991). Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART Bulletin*, *2*(4), 160–163. doi: [10.1145/122344.122377](https://doi.org/10.1145/122344.122377)
- Tambini, A., & Davachi, L. (2013). Persistence of hippocampal multivoxel patterns into post-encoding rest is related to memory. *Proceedings of the National Academy of Sciences*, *110*(48), 19591–19596. doi: [10.1073/pnas.1308499110](https://doi.org/10.1073/pnas.1308499110)
- Tambini, A., & Davachi, L. (2019). Awake reactivation of prior experiences consolidates memories and biases cognition. *Trends in Cognitive Sciences*, *23*(10), 876–890. doi: [10.1016/j.tics.2019.07.008](https://doi.org/10.1016/j.tics.2019.07.008)
- Tambini, A., Ketz, N., & Davachi, L. (2010). Enhanced brain correlations during rest are related to memory for recent experiences. *Neuron*, *65*(2), 280–290. doi: [10.1016/j.neuron.2010.01.001](https://doi.org/10.1016/j.neuron.2010.01.001)
- Taube, J., Muller, R., & Ranck, J. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. ii. effects of environmental manipulations. *Journal of Neuroscience*, *10*(2), 436–447. doi: [10.1523/JNEUROSCI.10-02-00436.1990](https://doi.org/10.1523/JNEUROSCI.10-02-00436.1990)
- Tirole, M., Gorriz, M. H., Takigawa, M., Kukovska, L., & Bendor, D. (2021). Experience-driven rate modulation is reinstated during hippocampal replay. *bioRxiv*. doi: [10.1101/2021.07.15.452506](https://doi.org/10.1101/2021.07.15.452506)
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*(4), 189–208. doi: [10.1037/h0061626](https://doi.org/10.1037/h0061626)
- Tompary, A., Duncan, K., & Davachi, L. (2015). Consolidation of associative and item memory is related to post-encoding functional connectivity between the ventral tegmental area and different medial temporal lobe subregions during an unrelated task. *Journal of Neuroscience*, *35*(19), 7326–7331. doi: [10.1523/jneurosci.4816-14.2015](https://doi.org/10.1523/jneurosci.4816-14.2015)
- Trettel, S. G., Trimper, J. B., Hwaun, E., Fiete, I. R., & Colgin, L. L. (2019). Grid cell co-activity patterns during sleep reflect spatial overlap of grid fields during active behaviors. *Nature Neuroscience*, *22*(4), 609–617. doi: [10.1038/s41593-019-0359-6](https://doi.org/10.1038/s41593-019-0359-6)
- Tulving, E. (2002). Episodic memory: From mind to brain. *Annual Review of Psychology*, *53*(1), 1–25. doi: [10.1146/annurev.psych.53.100901.135114](https://doi.org/10.1146/annurev.psych.53.100901.135114)
- Wang, M. E., Wann, E. G., Yuan, R. K., Ramos Alvarez, M. M., Stead, S. M., & Muzzio, I. A. (2012). Long-term stabilization of place cell remapping produced by a fearful experience. *Journal of Neuroscience*, *32*(45), 15802–15814. doi: [10.1523/jneurosci.0480-12.2012](https://doi.org/10.1523/jneurosci.0480-12.2012)
- Wang, Z., Bapst, V., Heess, N., Mnih, V., Munos, R., Kavukcuoglu, K., & de Freitas, N. (2016). Sample Efficient Actor-Critic with Experience Replay. *arXiv e-prints*, arXiv:1611.01224.
- Wikenheiser, A. M., & Redish, A. D. (2015). Decoding the cognitive map: Ensemble hippocampal sequences and decision making. *Current Opinion in Neurobiology*, *32*, 8–15. doi: [10.1016/j.conb.2014.10.002](https://doi.org/10.1016/j.conb.2014.10.002)
- Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, *265*(5172), 676–679. doi: [10.1126/science.8036517](https://doi.org/10.1126/science.8036517)



- Wimmer, G. E., & Büchel, C. (2021). Reactivation of single-episode pain patterns in the hippocampus and decision making. *Journal of Neuroscience*. doi: [10.1523/JNEUROSCI.1350-20.2021](https://doi.org/10.1523/JNEUROSCI.1350-20.2021)
- Wimmer, G. E., Liu, Y., McNamee, D., & Dolan, R. (2021). Distinct replay signatures for planning and memory maintenance. *bioRxiv*. doi: [10.1101/2021.11.08.467745](https://doi.org/10.1101/2021.11.08.467745)
- Wimmer, G. E., Liu, Y., Vehar, N., Behrens, T. E. J., & Dolan, R. J. (2020). Episodic memory retrieval success is associated with rapid replay of episode content. *Nature Neuroscience*. doi: [10.1038/s41593-020-0649-z](https://doi.org/10.1038/s41593-020-0649-z)
- Winocur, G., & Moscovitch, M. (2011). Memory transformation and systems consolidation. *Journal of the International Neuropsychological Society*, *17*(5), 766–780. doi: [10.1017/S1355617711000683](https://doi.org/10.1017/S1355617711000683)
- Wise, T., Liu, Y., Chowdhury, F., & Dolan, R. J. (2021). Model-based aversive learning in humans is supported by preferential task state reactivation. *Science Advances*, *7*(31). doi: [10.1126/sciadv.abf9616](https://doi.org/10.1126/sciadv.abf9616)
- Wu, C.-T., Haggerty, D., Kemere, C., & Ji, D. (2017). Hippocampal awake replay in fear memory retrieval. *Nature Neuroscience*, *20*(4), 571–580. doi: [10.1038/nn.4507](https://doi.org/10.1038/nn.4507)
- Xu, W., de Carvalho, F., & Jackson, A. (2019). Sequential neural activity in primary motor cortex during sleep. *Journal of Neuroscience*, *39*(19), 3698–3712. doi: [10.1523/JNEUROSCI.1408-18.2019](https://doi.org/10.1523/JNEUROSCI.1408-18.2019)
- Yu, J. Y., Kay, K., Liu, D. F., Grossrubatscher, I., Loback, A., Sosa, M., . . . Frank, L. M. (2017). Distinct hippocampal-cortical memory representations for experiences associated with movement versus immobility. *eLife*, *6*. doi: [10.7554/elife.27621](https://doi.org/10.7554/elife.27621)
- Yu, L. Q., Wilson, R. C., & Nassar, M. R. (2021). Adaptive learning is structure learning in time. *Neuroscience & Biobehavioral Reviews*, *128*, 270–281. doi: <https://doi.org/10.1016/j.neubiorev.2021.06.024>
- Zhang, H., Deuker, L., & Axmacher, N. (2017). Replay in humans - first evidence and open questions. In N. Axmacher & B. Rasch (Eds.), *Cognitive neuroscience of memory consolidation* (pp. 251–263). doi: [10.1007/978-3-319-45066-7\\_15](https://doi.org/10.1007/978-3-319-45066-7_15)
- Zielinski, M. C., Tang, W., & Jadhav, S. P. (2020). The role of replay and theta sequences in mediating hippocampal-prefrontal interactions for memory and cognition. *Hippocampus*, *30*(1), 60–72. doi: [10.1002/hipo.22821](https://doi.org/10.1002/hipo.22821)

# Appendices

## A Declaration of own share

Declaration pursuant to Sec. 7 (3), fourth sentence, of the Doctoral Study Regulations regarding my own share of the submitted scientific or scholarly work that has been published or is intended for publication within the scope of my publication-based work.

- I. Last name, first name:** Wittkuhn, John Lennart  
*Institute:* Max Planck Institute for Human Development, Berlin  
*Doctoral study subject:* Psychology  
*Title:* Investigating neural replay of task representations in the human brain using fMRI

**II. Numbered listing of works submitted (title, authors, where and when published and/or submitted):**

1. **Wittkuhn, L.**, Chien, S., Hall-McMaster, S., & Schuck, N. W. (2021). Replay in minds and machines. *Neuroscience & Biobehavioral Reviews*, 129, 367–388. doi: [10.1016/j.neubiorev.2021.08.002](https://doi.org/10.1016/j.neubiorev.2021.08.002)
2. **Wittkuhn, L.**, & Schuck, N. W. (2021). Dynamics of fMRI patterns reflect sub-second activation sequences and reveal replay in human visual cortex. *Nature Communications*, 12(1795). doi: [10.1038/s41467-021-21970-2](https://doi.org/10.1038/s41467-021-21970-2)
3. **Wittkuhn, L.**, Krippner, L. M., & Schuck, N. W. (2022). Statistical learning of successor representations is related to on-task replay. *bioRxiv*. doi: [10.1101/2022.02.02.478787](https://doi.org/10.1101/2022.02.02.478787)

**III. Explanation of own share of these works:**

The own share is assessed on the scale: *all – the vast majority – most – part*

Regarding II. 1.: development of the concept (most), literature research (the vast majority), preparing the manuscript (most), programming (part), publication of data/code (all).

Regarding II. 2.: development of the concept (most), literature research (the vast majority), method development (most), experiment design (the vast majority), data collection (all), data analysis/interpretation (the vast majority), discussion of results (most), preparing the manuscript (most), programming (all), publication of data/code (all).

Regarding II. 3.: development of the concept (the vast majority), literature research (all), method development (most), experiment design (the vast majority), data collection (most), data analysis/interpretation (the vast majority), discussion of results (most), preparing the manuscript (the vast majority), programming (the vast majority).

**IV. Names and e-mail addresses for the relevant co-authors:**

Regarding II. 1.: Nicolas W. Schuck ([schuck@mpib-berlin.mpg.de](mailto:schuck@mpib-berlin.mpg.de))

Regarding II. 2.: Samson Chien ([chiensam@gmail.com](mailto:chiensam@gmail.com)), Sam Hall-McMaster ([hall-mcmaster@mpib-berlin.mpg.de](mailto:hall-mcmaster@mpib-berlin.mpg.de)), Nicolas W. Schuck (see above)

Regarding II. 3.: Lena M. Krippner ([lena.m.krippner@gmail.com](mailto:lena.m.krippner@gmail.com)), Nicolas W. Schuck (see above)

## B Declaration of independent work

I hereby declare that:

- I completed this doctoral thesis independently. Except where otherwise stated, I confirm that the work presented in this thesis is my own.
- Where information has been derived from other sources, I confirm that this has been indicated in the thesis.
- I have not applied for a doctoral degree elsewhere and do not have a corresponding doctoral degree.
- I have acknowledged the Doctoral Degree Regulations which underlie the procedure of the Department of Education and Psychology of Freie Universität Berlin, as amended on August 8<sup>th</sup> 2016.
- The principles of Freie Universität Berlin for ensuring good academic practice have been complied with.

John Lennart Wittkuhn

Berlin, 08.02.2022

## C Paper I

**Wittkuhn, L.**, Chien, S., Hall-McMaster, S., & Schuck, N. W. (2021). Replay in minds and machines. *Neuroscience & Biobehavioral Reviews*, *129*, 367–388. doi: [10.1016/j.neubiorev.2021.08.002](https://doi.org/10.1016/j.neubiorev.2021.08.002)



Please note that the original article has been removed from the online version of this document to avoid copyright infringements. Please refer to the journal's website for access to the original publication.

## D Paper II

**Wittkuhn, L., & Schuck, N. W. (2021).** Dynamics of fMRI patterns reflect sub-second activation sequences and reveal replay in human visual cortex. *Nature Communications*, *12*(1795). doi: [10.1038/s41467-021-21970-2](https://doi.org/10.1038/s41467-021-21970-2)

This article is licensed under a Creative Commons Attribution 4.0 International (CC BY 4.0) license. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

# Dynamics of fMRI patterns reflect sub-second activation sequences and reveal replay in human visual cortex

Lennart Wittkuhn <sup>1,2</sup>✉ & Nicolas W. Schuck <sup>1,2</sup>✉

Neural computations are often fast and anatomically localized. Yet, investigating such computations in humans is challenging because non-invasive methods have either high temporal or spatial resolution, but not both. Of particular relevance, fast neural replay is known to occur throughout the brain in a coordinated fashion about which little is known. We develop a multivariate analysis method for functional magnetic resonance imaging that makes it possible to study sequentially activated neural patterns separated by less than 100 ms with precise spatial resolution. Human participants viewed five images individually and sequentially with speeds up to 32 ms between items. Probabilistic pattern classifiers were trained on activation patterns in visual and ventrotemporal cortex during individual image trials. Applied to sequence trials, probabilistic classifier time courses allow the detection of neural representations and their order. Order detection remains possible at speeds up to 32 ms between items (plus 100 ms per item). The frequency spectrum of the sequentiality metric distinguishes between sub- versus supra-second sequences. Importantly, applied to resting-state data our method reveals fast replay of task-related stimuli in visual cortex. This indicates that non-hippocampal replay occurs even after tasks without memory requirements and shows that our method can be used to detect such spontaneously occurring replay.

<sup>1</sup>Max Planck Research Group NeuroCode, Max Planck Institute for Human Development, Berlin, Germany. <sup>2</sup>Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Berlin, Germany. ✉email: [wittkuhn@mpib-berlin.mpg.de](mailto:wittkuhn@mpib-berlin.mpg.de); [schuck@mpib-berlin.mpg.de](mailto:schuck@mpib-berlin.mpg.de)

Many cognitive processes are underpinned by rapidly changing neural activation patterns. Most famously, memory and planning have been linked to fast replay of representation sequences in the hippocampus, happening approximately within 200–300 milliseconds (ms) while the animal is resting or sleeping, e.g.<sup>1–9</sup>. Similar events have been observed during behavior<sup>10,11</sup>, as well as outside of the hippocampus<sup>12–17</sup>. Likewise, internal deliberations during choice are reflected in alternations between orbitofrontal value representations that last less than 100 ms<sup>18</sup>, while perceptual learning has been shown to result in sub-second anticipatory activation sequences in visual cortex<sup>19–21</sup>. Investigating fast-paced representational dynamics within specific brain areas therefore promises important insights into a variety of cognitive processes. Such investigations could be crucial for understanding replay, which is characterized by a widespread co-occurrence of neural reactivation events throughout the brain of mostly unknown functional significance, in particular outside of the hippocampus, see, e.g.<sup>17,22</sup>. These aspects are still understudied in humans.

Studying fast neural dynamics is particularly difficult in humans because signal recording must mainly occur non-invasively. How fast and anatomically localized neural dynamics can be investigated using non-invasive neuroimaging techniques is therefore a major challenge for human neuroscience, see, e.g.<sup>23,24</sup>. The main concern related to functional magnetic resonance imaging (fMRI) is that this technique measures neural activity indirectly through slow sampling of an extended and delayed blood-oxygen-level-dependent (BOLD) response function<sup>25–27</sup> that can obscure temporal detail. Yet, the problems arising in BOLD fMRI might not be as insurmountable as they seem. First, BOLD signals from the same participant and brain region show reliable timing and last for several seconds. Miezin et al.<sup>28</sup>, for instance, reported a between-session reliability of hemodynamic peak times in visual cortex of  $r^2 = 0.95$ , see also<sup>29,30</sup>. Even for closely timed events, the sequential order can therefore result in systematic differences in activation strength<sup>31</sup> that remain in the signal long after the fast sequence event is over, effectively mitigating the problems that arise from slow sampling. Moreover, Misaki et al.<sup>32</sup> were able to decode onset differences in visual stimulation of only 100 ms when two stimuli were shown to one eye before the other. Interestingly, Misaki et al.<sup>32</sup> indicated that timing differences become most apparent in peak activation strength, rather than temporal aspects of the hemodynamic response function (HRF). A second reason that makes the investigation of fast neural dynamics feasible is that some fast sequence events have properties that make it easier to detect them. Replay events, in particular, involve reactivation of spatially tuned cells in the order of a previously traveled path. But these reactivated paths do not typically span the entire spatial environment and only involve a local subset of all possible places the animal could occupy<sup>7,8</sup>. This locality means that even when measurement noise causes some elements of a fast sequence to remain undetected, or leads to partially re-ordered detection, the set of detected representations will still reflect positions nearby in space. In this case, successive detection of elements nearby in space or time would still identify the fast process under investigation even under noisy conditions.

If fMRI analyses can capitalize on such effects, this could allow the investigation of fast sequential activations. As mentioned above, one important application of such methods would be hippocampal replay, a topic of intense recent interest, for reviews, see, e.g.<sup>24,33–37</sup>. To date, most replay research has studied the phenomenon in rodents because investigations in humans and other primates either required invasive recordings from the hippocampus<sup>38–42</sup>, used techniques with reduced hippocampal sensitivity and spatial resolution<sup>43–48</sup>, or investigated non-

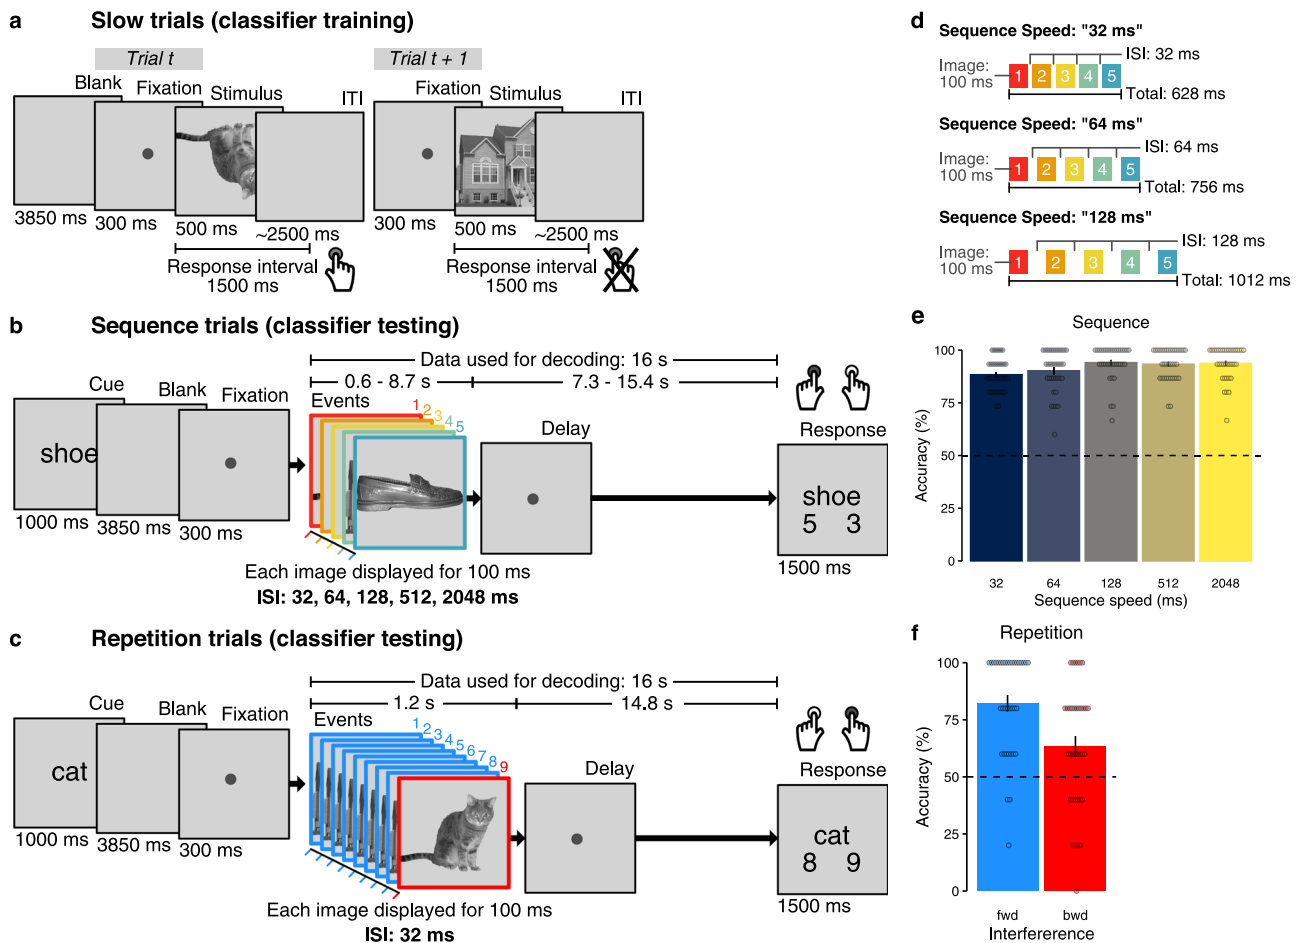
sequential fMRI activation patterns over seconds or minutes<sup>49–53</sup>. Recently, we have hypothesized that the properties of BOLD signals mentioned above should enable the investigation of rapid neural dynamics. Indeed, using fMRI, we identified fast sequential hippocampal pattern reactivation in resting humans<sup>54</sup>. However, Schuck and Niv<sup>54</sup> did not yet answer questions about how fMRI could be used to measure the speed of replay. One additional exploratory question is whether replay occurs outside of the hippocampus, and even following simple visual detection tasks.

Here, we provide and experimentally validate a multivariate analysis approach for fMRI that addresses the challenges and questions outlined above. The main idea of our approach is that fast neural event sequences will cause characteristic time courses of overlapping activation patterns. While the effects of co-occurring activations on individual voxels is complex, we reason that characteristic overlap will nevertheless lead to predictable and simple fluctuations in the time courses of pattern classifiers. The present experiment tests this idea and our results confirm that logistic regression classifier time courses reveal the content and order of fast sequential neural events using fMRI. Importantly, we use this method to ask whether sequential reactivations of sensory events occur outside of the hippocampus, even if task experiences did not require memorization or involve repeated sequential structure. Our study extends our previous work in several ways. First, our controlled experimental design provides evidence for the decodability of fast sequential neural events in a setting where the speed and order of fast neural event sequences are known. We also show that sequence detection can be achieved in the presence of high levels of signal noise and timing uncertainty, and is specific enough to differentiate fast sequences from activation patterns that could reflect slow conscious thinking. Second, we develop a modeling approach of multivariate fMRI pattern classification time courses that validates our experimental results and allows inference of the speed of fast sequential neural processes from the frequency spectra of our fMRI sequentiality metric. Third, we report that our task induced fast sequential replay in sensory brain areas during post-task rest, although it did not require any memorization, did not feature strong sequential structure, and did not elicit systematic hippocampal responses. Finally, our results have implications for the interpretation of our own previous results in Schuck and Niv<sup>54</sup> and future fMRI studies investigating fast neural event sequences, like hippocampal replay.

## Results

As discussed above, we investigated the possibility that fMRI can be used to address two cornerstones of understanding signals resulting from fast activation sequences: *order detection* and *element detection*. The first effect, order detection, pertains to the presence of order structure in the signal that is caused by the sequential order of fast neural events. We evaluated this effect by investigating the impact of item order on (a) the relative strength of activations within a single measurement, and (b) the order of decoded patterns across successive measurements. The second effect, element detection, quantifies to what extent fMRI allows detection of elements that were part of a sequence versus those that were not. While event detection is a standard problem in fMRI, we focused on the special case relevant to our question: detecting neural patterns of brief events that are affected by patterns from other sequence elements occurring only tens of milliseconds before or afterwards, causing backward and forward interference, respectively. Using full sequences of all possible elements in our experimental setup that tested sequence ordering, our design ensured that the two effects tested can be demonstrated





**Fig. 1 Task design and behavioral performance.** **a** On slow trials, individual images were presented and inter-trial intervals (ITIs) were 2.5 s on average. Participants were instructed to detect upside-down visual stimuli (20% of trials) but not respond to upright pictures. Classifier training was performed on fMRI data from correct upright trials only. **b** Sequence trials contained five unique visual images, separated by five levels of inter-stimulus intervals (ISIs) between 32 and 2048 ms. **c** Repetition trials were always fast (32 ms ISI) and contained two visual images of which either the first or the second was repeated eight times (causing backward and forward interference, respectively). In both task conditions, participants were asked to detect the serial position of a cued target stimulus in a sequence and select the correct answer after a delay period without visual input. One sequence or repetition trial came after five slow trials. fMRI analyses focused on the time from sequence onset to the end of the delay period (16 s  $\approx$  13 TRs, 1 TR = 1.25 s). **d** Illustration of the three fastest sequence speed conditions of 32, 64, and 128 ms ISI between images. **e** Mean behavioral accuracy in sequence trials (in %) as a function of sequence speed (ISI, in ms;  $N = 36$ ,  $t_s \geq 23.78$ ,  $p_s < 0.001$ ,  $d_s \geq 3.96$ , linear mixed effects (LME) model and five one-sided one-sample  $t$ -tests against chance (50%), false discovery rate (FDR) correction). **f** Mean behavioral accuracy in repetition trials (in %), as a function of which sequence item was repeated (fwd = forward, bwd = backward condition;  $N = 36$ ,  $t_s \geq 2.94$ ,  $p_s \leq 0.003$ ,  $d_s \geq 0.49$ , two one-sided one-sample  $t$ -tests against chance (50%) with FDR-correction). All error bars represent  $\pm 1$  standard error of the mean (SEM). All statistics have been derived from data of  $N = 36$  human participants who participated in one experiment. The horizontal dashed lines in (e) and (f) indicate 50% chance level. The original authors of Haxby et al.<sup>55</sup> hold the copyright of the stimulus material (individual images of a cat, chair, face, house, and shoe) shown in (a), (b), and (c) and made it available under the terms of the Creative Commons Attribution-Share Alike 3.0 license (see <http://data.pymvpa.org/datasets/haxby2001/> and <http://creativecommons.org/licenses/by-sa/3.0/> for details). Source data are provided as a Source Data file.

independently, i.e., the order effect could not have been a side effect of element detection.

Participants viewed images of five different objects. During *slow trials* (Fig. 1a, 600 trials in total), individual images were shown with inter-trial intervals (ITIs) of approximately 2.5 s, as is common in fMRI decision-making experiments (cf.<sup>44,47,52,54</sup>). In fast trials (120 trials in total), the same images were shown as either a random sequence of all five objects (*sequence trials*, 75 trials, Fig. 1b), or two objects were repeated several times (*repetition trials*, 45 trials, Fig. 1c). Importantly, image presentation rate was greatly increased in sequence and repetition trials, with as little as 32 ms between stimuli and a presentation time of 100 ms per stimulus. Logistic regression classifiers were trained on data from slow trials and applied to sequence and repetition trials,

as well as to resting-state data. We then asked whether the order and the elements of fast sequences are detectable from fMRI signals, depending on sequence speed, number of repetitions, level of background noise, and timing uncertainty. To this end, visual stimuli in sequence and repetition trials were presented in a precisely timed and ordered manner, as detailed below. Since activation patterns were primarily visual in nature, only data from visual and ventral temporal cortex were considered. A corresponding analysis using hippocampal data did not yield comparable results, see below. The analyses included  $N = 36$  human participants who underwent two fMRI sessions with four task runs each, i.e., eight runs in total. Four additional participants were excluded from analyses due to insufficient performance, see Methods and Supplementary Information (SI) (Supplementary

Fig. 1a). Sessions were separated by 9 days on average (SD = 6 days, range: 1–24 days).

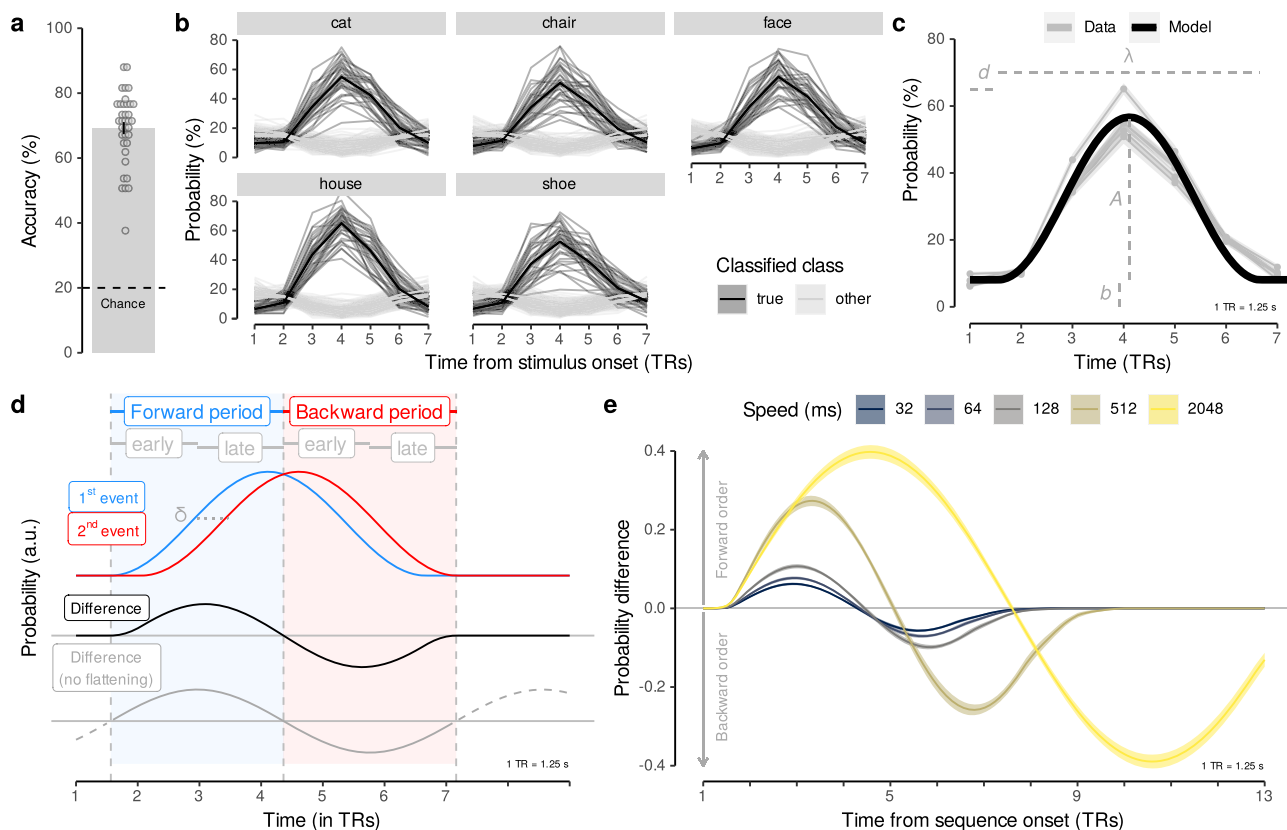
**Training fMRI pattern classifiers on slow events.** In slow trials, participants repeatedly viewed the same five images individually for 500 ms (images showed a cat, chair, face, house, and shoe, taken from<sup>55</sup>). Temporal delays between images were set to 2.5 s on average, as typical for task-based fMRI experiments<sup>56</sup>. To ensure that image ordering did not yield biased classifiers through biased pattern similarities (cf.<sup>57</sup>), each possible order permutation of the five images was presented exactly once (120 sets of 5 images each). Participants were kept attentive by a cover task that required them to press a button whenever a picture was shown upside-down (20% of trials; mean accuracy = 99.44%;  $t_{(35)} = 263.27$ , 95% CI [99.13, +∞];  $p < 0.001$ , compared to chance (50%);  $d = 43.88$ ; Supplementary Fig. 1a–c). Using data from correct upright slow trials, we trained five separate multinomial logistic regression classifiers, one for each image category (one-vs.-rest; see Methods for details; cf.<sup>55</sup>). fMRI data were masked by a gray-matter-restricted region of interest (ROI) of occipito-temporal cortex, known to be related to visual object processing (11,162 voxels in the masks on average; cf.<sup>55,58–60</sup>). Spatial patterns associated with image categories indicated a mix of overlapping and non-overlapping sets of voxels, and average correlations between the mean voxel patterns were negative (see SI). We accounted for hemodynamic lag by extracting fMRI data acquired 3.75–5 s after stimulus onset (corresponding to the fourth repetition time (TR), see Methods). Cross-validated (leave-one-run-out) classification accuracy was on average 69.22% (SD = 11.18%;  $t_{(35)} = 26.41$ , 95% CI [66.07, +∞],  $p < 0.001$ , compared to chance (20%);  $d = 4.40$ ; Fig. 2a). In order to examine the sensitivity of the classifiers to pattern activation time courses, we applied them to seven TRs following stimulus onset on each trial. This analysis confirmed delayed and distinct increases in the estimated probability of the true stimulus class given the data, peaking at the fourth TR after stimulus onset, as expected, given that the classifiers were trained on data from the fourth TR following stimulus onset (Fig. 2b). The peak in probability for the true stimulus shown on the corresponding trial was significantly higher than the mean probability of all other stimuli at that time point ( $ts \geq 17.95$ ,  $ps < 0.001$ ,  $ds \geq 2.99$ ; Bonferroni-corrected). Decoding in an anatomical ROI of the hippocampus did not surpass the chance level (decoding accuracy: mean (M) = 20.52%, SD = 1.49%;  $t_{35} = 2.10$ , 95% CI [20.02, 21.03],  $p = 0.05$ , compared to chance (20%),  $d = 0.35$ ; using the same decoding approach, see SI for details).

**Single event and event sequence modeling.** The data shown in Fig. 2b highlight that multivariate decoding time courses are delayed and sustained, similar to single-voxel hemodynamics. We captured these dynamics elicited by single events by fitting a sine-based response function to the time courses on slow trials (a single sine wave flattened after one cycle, with parameters for amplitude  $A$ , response duration  $\lambda$ , onset delay  $d$ , and baseline  $b$ ; Fig. 2c and Supplementary Fig. 4; see Methods). Based on this fit to single events, we derived expectations for probabilistic time courses during sequential events. The sequentiality analyses reported below essentially quantify how well successive activation patterns can be differentiated from one another depending on the speed of stimulus sequences. We therefore considered two time-shifted response functions and derived the magnitude and time course of differences between them. Based on the sinusoidal nature of the response function, the time course of this difference can be approximated by a single sine wave with duration  $\lambda_\delta = \lambda + \delta$ , where  $\delta$  is the time between events and  $\lambda$  is the average fitted

single event duration, here  $\lambda = 5.24$  TRs (see Eqs. (4) and (5), Methods). This average parameter was used for all further analyses (Fig. 2c, d; see Methods). In this model, the amplitude is proportional to the time shift between events (until time shifts become larger than the time-to-peak of the response function). Consequently, after an onset delay ( $d = 0.56$  TRs), the difference in probability of two time-shifted events is expected to be positive for the duration of half a cycle, i.e.,  $0.5\lambda_\delta = 0.5(5.24 + \delta)$  TRs, and negative for the same period thereafter. Simply put, this means that the strength of overlapping activations will initially be ordered forward, in the same way as the sequence, i.e., earlier items will be activated stronger. In a later period, however, this will reverse and result in backwards ordering, i.e., earlier items will be activated less. In summary, three predictions therefore arise from this model: (1) the first event will dominate the signal in earlier TRs, and activation strengths will be proportional to the true event order during the sequential process; (2) in later TRs, the last sequence element will dominate the signal, and the activation strengths will be ordered backwards; and (3) the duration and strength of these two effects will depend on the fitted response duration and the timing of the stimuli as specified above (Fig. 2e and Eqs. (1)–(5); see Methods). For sequences with more than two items (as in sequence trials, see below),  $\delta$  is defined as the interval between the onsets of the first and last sequence item. To reflect the relation between the true order and the activation strength, we henceforth term the above-mentioned early and late TRs as the forward and backward periods, and consider all results below either separately for these phases, or for both relevant periods combined (calculating periods depending on the timings of image sequences and rounding TRs, see Methods).

**Detecting sequentiality in fMRI patterns following fast and slow neural event sequences.** Our first major aim was to test detection of sequential order of fast neural events with fMRI. We therefore investigated the above-mentioned sequence trials in which participants viewed a series of five unique images at different speeds (Fig. 1b). Sequence speed was manipulated by leaving either 32, 64, 128, 512, or 2048 ms between pictures, while images were always presented briefly (100 ms per image, total sequence duration 0.628–8.692 s). Note, that we refer to the inter-stimulus interval (ISI) as “sequence speed” (see Fig. 1d). Sequences always contained each image exactly once. Every participant experienced 15 randomly selected image orders that ensured that each image appeared equally often at the first and last position of the sequence (all 120 possible orders counterbalanced across participants). The task required participants to indicate the serial position of a verbally cued image 16 s after the first image was presented. This delay between visual events and response (roughly spanning 13 TRs; see  $x$ -axes in Fig. 3a, b) allowed us to measure sequence-related fMRI signals without interference from following trials, while the upcoming question did not necessitate memorization of the sequence during the delay period. Performance was high even in the fastest sequence trials (32 ms: M = 88.33%, SD = 7.70,  $t_{35} = 29.85$ , 95% CI [86.16, +∞],  $p < 0.001$  compared to chance (50%),  $d = 4.98$ ), and only slightly reduced compared to the slowest condition (2048 ms: M = 93.70%, SD = 7.96,  $t_{35} = 32.95$ , 95% CI [91.46, +∞],  $p < 0.001$  compared to chance (50%),  $d = 5.49$ ; Fig. 1e and Supplementary Fig. 1d).

We investigated whether sequence order was detectable from the relative pattern activation strength within a single measurement. Examining the time courses of probabilistic classifier evidence during sequence trials (Fig. 3a) showed that the time delay between events was indeed reflected in sustained within-TR ordering of probabilities in all speed conditions. Specifically, immediately after sequence onset, the first element (red line) had

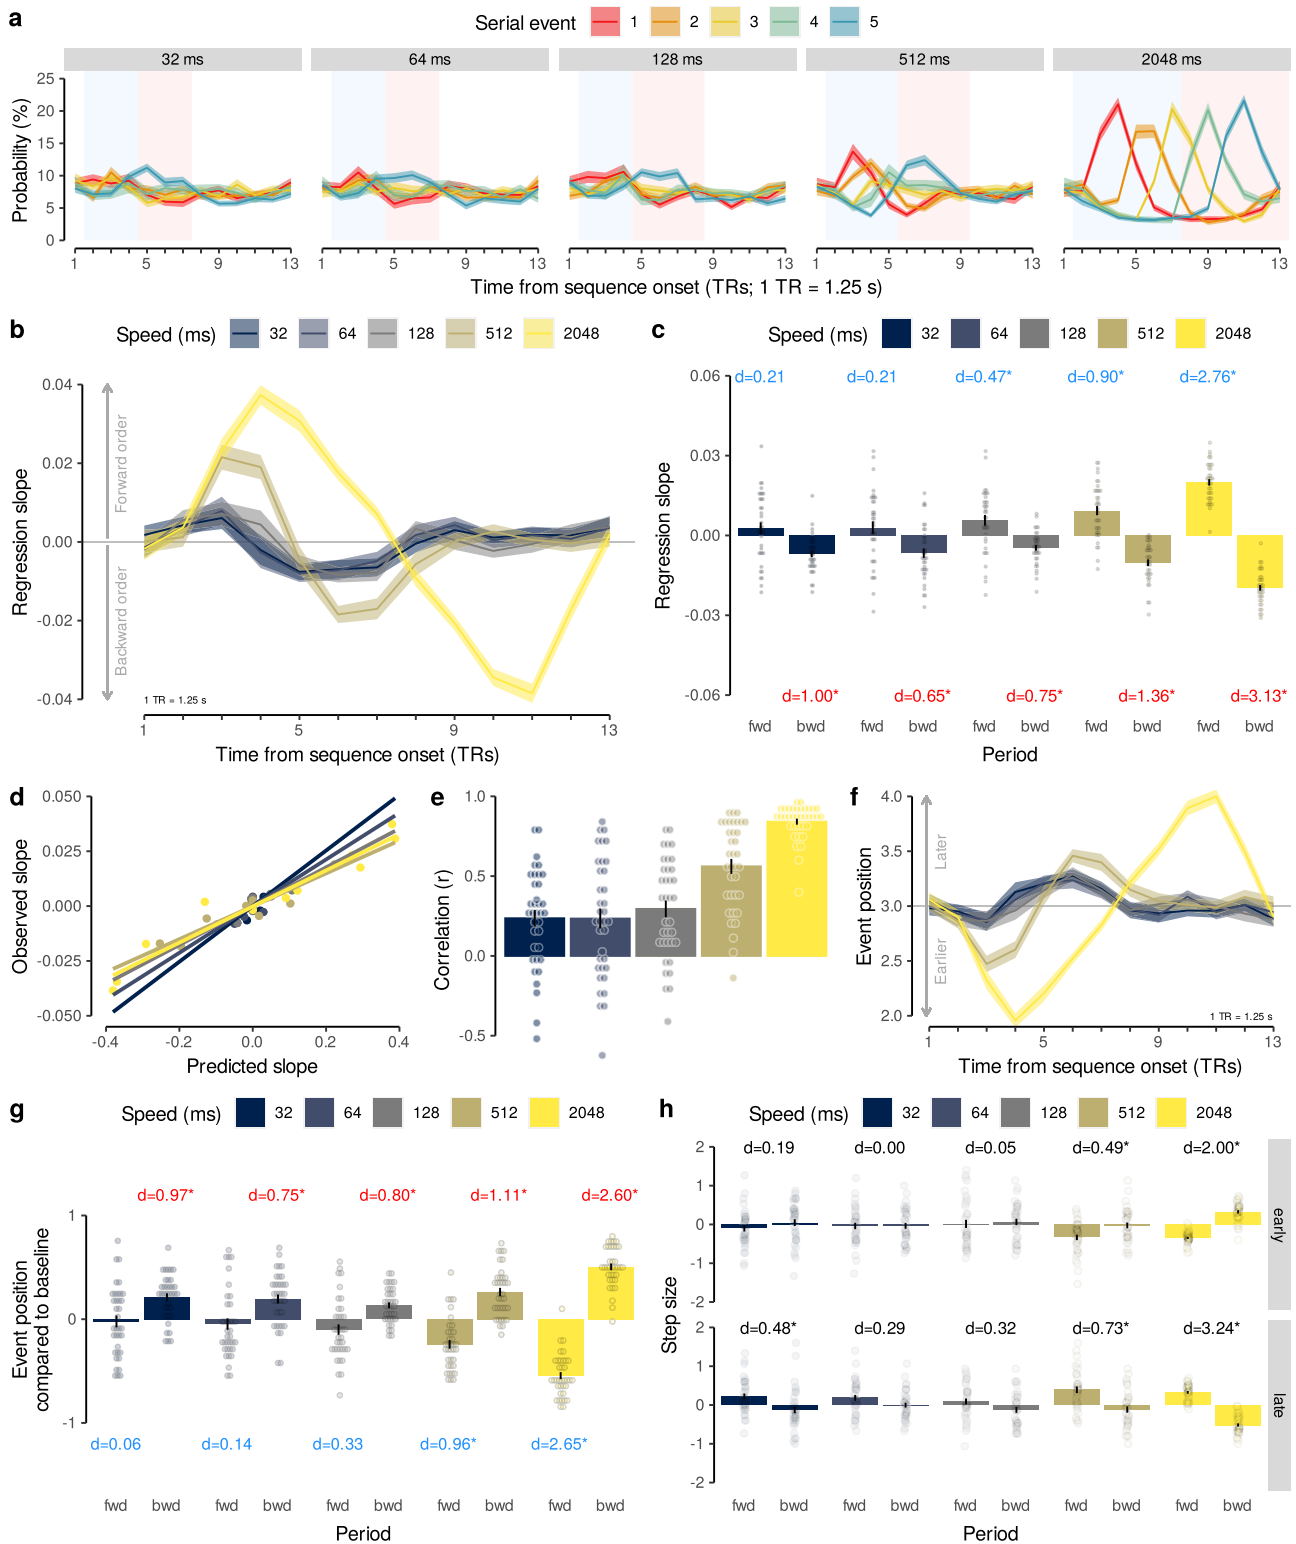


**Fig. 2 Classification accuracy and multivariate response functions.** **a** Cross-validated classification accuracy in decoding the five unique visual objects in occipito-temporal data during task performance (in %;  $N = 36$ ,  $t_{(35)} = 26.41$ , 95% CI [66.07,  $+\infty$ ],  $p < 0.001$ ,  $d = 4.40$ , one one-sided one-sample  $t$ -test, no multiple comparisons). Chance level is 20% (dashed line). Each dot corresponds to averaged data from one participant. Error bar represents  $\pm 1$  SEM. **b** Time courses (in TRs from stimulus onset) of probabilistic classification evidence (in %) for all five stimulus classes. Substantial delayed and extended probability increases for the stimulus presented (black lines) on a given trial (gray panels) were found. Each line represents one participant ( $N = 36$ ,  $t_s \geq 17.95$ ,  $p_s < 0.001$ ,  $d_s \geq 2.99$ , 35 two-sided two-sample  $t$ -tests, Bonferroni-corrected). **c** Average probabilistic classifier response for the five stimulus classes (gray lines) and fitted sine-wave response model using averaged parameters (black line). **d** Illustration of sinusoidal response functions following two neural events (blue and red lines) time-shifted by delta seconds (dashed horizontal line). The resulting difference between event probabilities (black line) establishes a forward (blue area) and backward (red area) time period, split into early and late phases. The sine-wave approximation without flattened tails is shown in gray. **e** Probability differences between two time-shifted events predicted by the sinusoidal response functions depending on the event delays ( $\delta$ ) as they occurred in the five different sequence speed conditions (colors), based on Eq. (6). All statistics have been derived from data of  $N = 36$  human participants who participated in one experiment. Source data are provided as a Source Data file.

the highest probability and the last element (blue line) had the lowest probability. This pattern reversed afterwards, following the forward and backward dynamics that were predicted by the time-shifted response functions (Fig. 2d; forward and backward periods adjusted to sequence speed, see above and Methods). A TR-wise linear regression between the serial positions of the images and their probabilities confirmed this impression. In all speed conditions, the mean slope coefficients initially increased above zero (reflecting higher probabilities of earlier compared to later items) and decreased below zero afterwards (Fig. 3b and Supplementary Fig. 6a). Considering mean regression coefficients during the predicted forward and backward periods, we found significant forward ordering in the forward period at ISIs of 128, 512, and 2048 ms ( $t_s \geq 2.85$ ,  $p_s \leq 0.009$ ,  $d_s \geq 0.47$ ) and significant backward ordering in the backward period in all speed conditions ( $t_s \geq 3.89$ ,  $p_s < 0.001$ ,  $d_s \geq 0.65$ , FDR-corrected; Fig. 3c). Notably, the observed time course of regression slopes on sequence trials (Fig. 3b) closely matched the time course predicted by our modeling approach (Fig. 2d), as indicated by strong correlations for all speed conditions between model predictions and the averaged time courses (Fig. 3d; Pearson’s  $r_s \geq 0.81$ ,  $p_s < 0.001$ ) as well as significant within-participant correlations (Fig. 3e; mean

Pearson’s  $r_s \geq 0.23$ ,  $t_s \geq 3.76$ ,  $p_s < 0.001$ , compared to zero,  $d_s \geq 0.63$ , FDR-corrected).

Choosing a different index of association like rank correlation coefficients (Supplementary Figs. 5a, b and 6c) or the mean step size between probability-ordered events within TRs (Supplementary Figs. 5c, d and 6d) produced qualitatively similar results (for details, see SI). Removing the sequence item with the highest probability at every TR also resulted in similar effects, with backward sequentiality remaining significant at all speeds ( $p \leq 0.002$ ) except the 32 and 128 ms conditions ( $p \geq 0.20$ ), and forward sequentiality still being evident at speeds of 512 and 2048 ms ( $p \leq 0.004$ ; Supplementary Fig. 7a, b). To identify the drivers of the apparent asymmetry in detecting forward and backward sequentiality, we ran two additional control analyses and either removed the probability of the first or the last sequence item (forward and backward periods adjusted accordingly). Removal of the first sequence item had little impact on sequentiality detection (Supplementary Fig. 7c, d and SI), but removing the last sequence item markedly affected the results such that significant forward and backward sequentiality was only evident at speeds of 512 and 2048 ms (Supplementary Fig. 7e, f and SI).



Next, we investigated evidence of pattern sequentiality across successive measurements, similar to Schuck and Niv<sup>54</sup>. Specifically, for each TR we only considered the decoded image with the highest probability and asked whether earlier images were decoded primarily in earlier TRs, and whether later images were primarily decoded in later TRs. In line with this prediction, the average serial position fluctuated in a similar manner as the regression coefficients, with a tendency of early positions to be decoded in early TRs, and later positions in later TRs (Fig. 3f).

The average serial position of the decoded images was therefore significantly different between the predicted forward and backward period at all sequence speeds (all  $ps < 0.001$ , Fig. 3g, Supplementary Fig. 6d). Compared to baseline (mean serial position of 3), the average serial position during the forward period was significantly lower for speeds of 512 and 2048 ms (all  $ps < 0.001$ ). The average decoded serial position at later time points was significantly higher compared to baseline in all speed conditions, including the 32 ms condition (all  $ps < 0.001$ ). Thus,

**Fig. 3 Sequence order is reflected in probability time courses.** **a** Time courses (TRs from sequence onset) of classifier probabilities (%) per event (colors) and sequence speed (panels). Forward (blue) and backward (red) periods shaded as in Fig. 2d. **b** Time courses of mean regression slopes between event position and probability for each speed (colors). Positive/negative values indicate forward/backward sequentiality, respectively. **c** Mean slope coefficients for each speed (colors) and period (forward vs. backward;  $N = 36$ ,  $t_s \geq 2.85$ ,  $p_s \leq 0.009$ ,  $d_s \geq 0.47$  (significant tests only), ten two-sided one-sample  $t$ -tests against zero, FDR-corrected). Asterisks indicate significant differences from baseline. **d** Between-participant correlation between predicted (Fig. 2e, Eq. (6)) and observed (b) time courses of mean regression slopes (13 TRs per correlation, Pearson's  $r_s \geq 0.81$ ,  $p_s < 0.001$ ). Each dot represents one TR. **e** Mean within-participant correlations between predicted and observed slopes as in (d) ( $N = 36$ , mean Pearson's  $r_s \geq 0.23$ ,  $t_s \geq 3.76$ ,  $p_s \leq 0.001$ , compared to zero,  $d_s \geq 0.63$ , FDR-corrected). **f** Time courses of mean event position for each speed, as in (b). **g** Mean event position for each period and speed, as in (c) ( $N = 36$ ,  $t_s \geq 4.78$ ,  $p_s < 0.001$ ,  $d_s \geq 0.75$  (significant tests only), ten two-sided one-sample  $t$ -tests against baseline, FDR-corrected). **h** Mean step sizes of early and late transitions for each period and speed ( $N = 36$ ,  $t_s \geq 2.88$ ,  $p_s \leq 0.006$ ,  $d_s \geq 0.48$  (significant tests only), ten two-sided one-sample  $t$ -tests against zero, FDR-corrected). Asterisks indicate differences between periods, otherwise as in (c). Each dot represents data of one participant. Error bars/shaded areas represent  $\pm 1$  SEM. All statistics have been derived from data of  $N = 36$  human participants who participated in one experiment. Effect sizes indicated by Cohen's  $d$ . Asterisks indicate  $p < 0.05$ , FDR-corrected. 1 TR = 1.25 s. Source data are provided as a Source Data file.

earlier images were decoded earlier after sequence onset and later images later, as expected.

This sequential progression through the involved sequence elements had implications for transitions between consecutively decoded events. The transitions will be a direct function of the slope of the average decoded position shown in Fig. 3f. When the slope is negative, the steps between successive sequence items are backward and reflect the transition from a later position to an earlier position. When the slope is positive, the steps are forward, reflecting a progression from an earlier event position to a later event position. This can be verified by computing the step sizes between consecutively decoded serial events as in Schuck and Niv<sup>54</sup>. For example, observing a  $2 \rightarrow 4$  transition of decoded events in consecutive TRs would correspond to a forward step of size  $+2$ , while a  $3 \rightarrow 2$  transition would reflect a backward step of size  $-1$ . As can be seen from Fig. 3f, both the early and late phase of the response (see phases in Fig. 2d) included periods with a negative and a positive slope, in line with our predictions (formally, the prediction can be obtained by taking the derivative with respect to time of Eq. (6), see Methods, i.e., the function shown in Fig. 2e). We therefore considered the periods with a positive and negative position slope separately for the early and late phase. As expected, the early transitions were mainly forward during the period of a positive slope as compared to the negative slope periods for speed conditions of 512 and 2048 ms ( $p_s \leq 0.01$ , Fig. 3h). Similarly, the late transitions were also forward and backward during the positive and negative slope periods, respectively, and differed in all speed conditions ( $p_s \leq 0.01$ , Fig. 3h), except the 64 and 128 ms conditions ( $p = 0.12$  and  $p = 0.10$ ; FDR-corrected). This analysis suggests that transitions between decoded items reflect the ordered progression from early to late and then from late to early sequence events, even when events were separated only by tens of milliseconds.

#### Detecting sequence elements: asymmetries and interference effects.

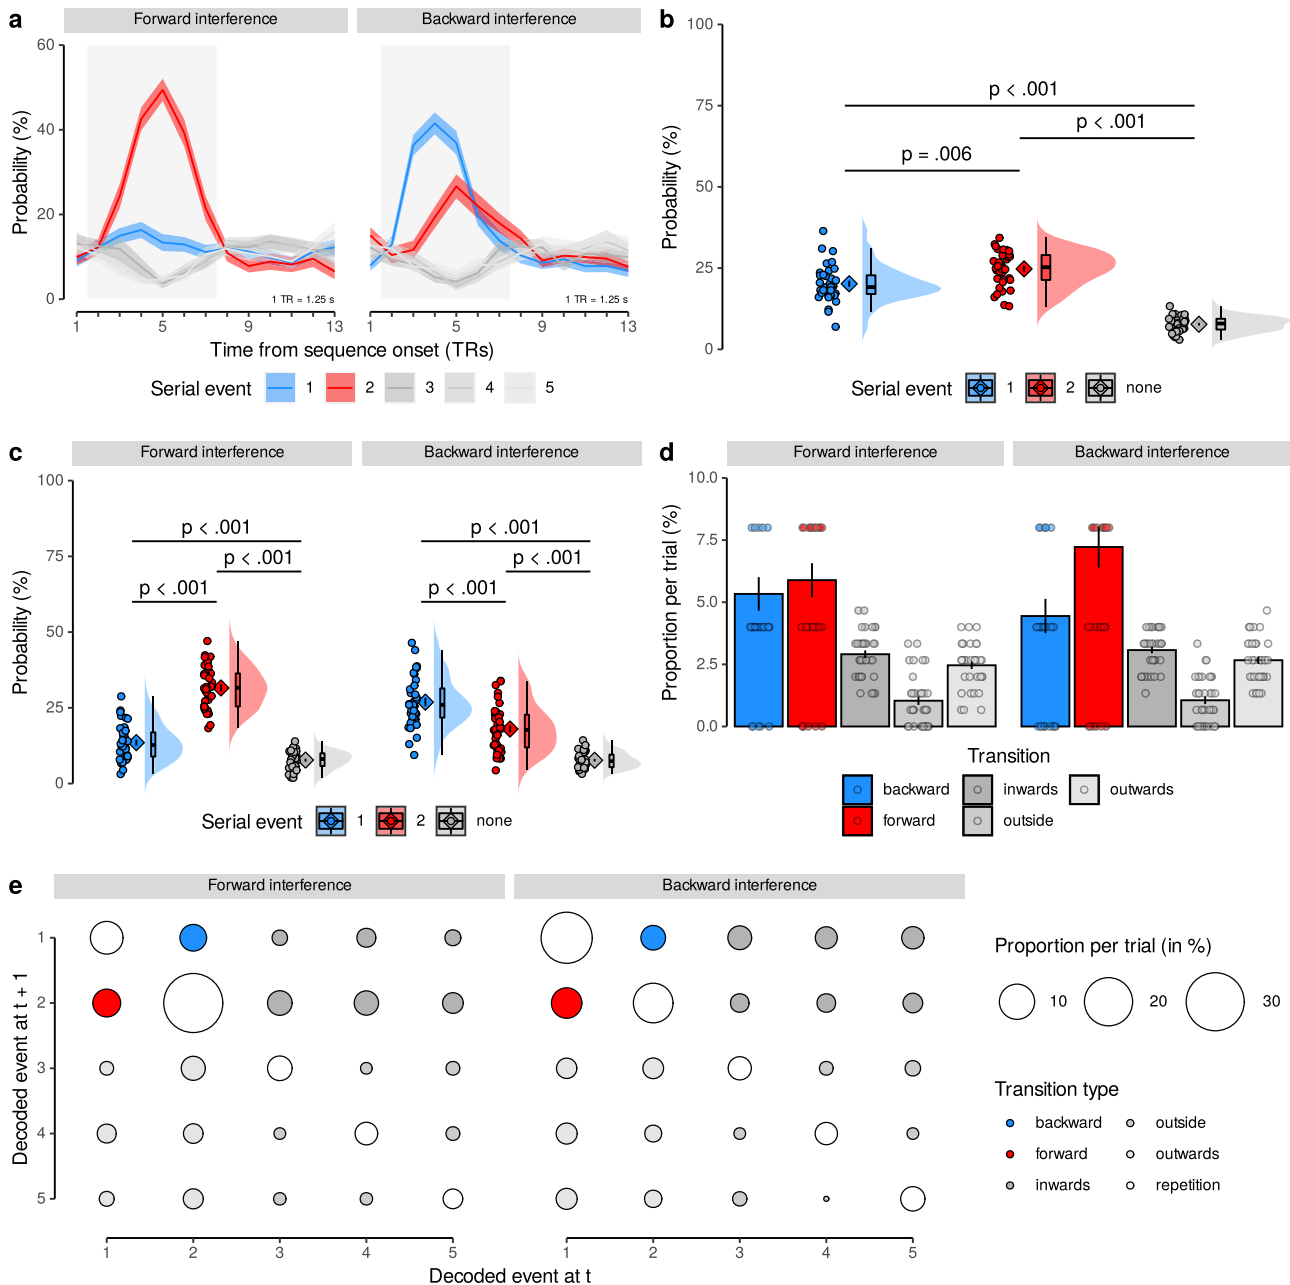
We next turned to our second main question, asking whether we can detect which patterns were part of a fast sequence and which were not. One important reason why detecting which patterns were activated during sequence events might be more difficult than in a standard setting is that co-activation of multiple patterns close in time could lead to interference. We therefore investigate such interference in detail below.

We analyzed classification time courses in repetition trials, in which only two out of the five possible images were shown. One of the two images was repeated, while the other one was shown only once. This setup allowed us to study to what extent another activation (the repeated image) can interfere with the detection of a brief activation pattern of interest (the image shown only once). The repeating image was shown eight times, which created maximally adverse effects for the detection of the single image. To

ask if detection of brief activations is differently affected by events occurring before versus after the single event, we varied whether the single item was preceded or followed by the repeated item. We pose this question because the backward effects were consistently larger than forward effects in our sequentiality analyses reported above (Fig. 3c), suggesting asymmetric detection sensitivity. This implies that one briefly presented item at the end of a sequence will be easier to detect than a briefly presented item at the beginning of a sequence, even though both were equally close in time to another strong activation signal. To test this idea, we considered the two order conditions described above. We will term the case in which the first image was shown briefly once and followed immediately by eight repetitions of a second image the *forward interference* condition, because the forward phase of the sequential responses suffers from interference. Correspondingly, trials in which the first image was repeated eight times and the second image was shown once will be termed the *backward interference* condition. In all cases, images were separated by only 32 ms. Participants were kept attentive by the same cover task used in sequence trials (Fig. 1c). Average behavioral accuracy was high on repetition trials ( $M = 73.46\%$ ,  $SD = 9.71\%$ ; Fig. 1f and Supplementary Fig. 1a) and clearly differed from a 50% chance level ( $t_{(35)} = 14.50$ , 95% CI [70.72,  $+\infty$ ],  $p < 0.001$ ,  $d = 2.42$ ). Splitting up performance into forward and backward interference trials showed performance above chance level in both conditions ( $M = 82.22\%$  and  $M = 63.33\%$ , respectively,  $t_s \geq 2.94$ ,  $p_s \leq 0.003$ ,  $d_s \geq 0.49$ , Fig. 1f).

As before, we applied the classifiers trained on slow trials to the data acquired in repetition trials and obtained the estimated probability of every class given the data for each TR (Fig. 4a and Supplementary Fig. 9). The expected relevant time period was determined to be from TRs 2 to 7 and used in all analyses (see rectangular areas in Fig. 4a).

We first asked whether our classifiers indicated that the two events that were part of the sequence were more likely decoded than items that were not part of the sequence. Indeed, the event types (first, second, non-sequence) had significantly different mean decoding probabilities, with sequence items having a higher probability (first:  $M = 20.19\%$ ; second:  $M = 24.78\%$ ) compared to non-sequence items ( $M = 7.72\%$ ; both  $p_s < 0.001$ , corrected; main effect:  $F_{2,57.78} = 110.13$ ,  $p < 0.001$ ; Fig. 4b). Moreover, the probability of decoding within-sequence items depended on the condition and whether the item was repeated or not. Considering both interference conditions (forward/backward) in the same analysis revealed a main effect of condition,  $F_{2,41.64} = 146.15$ ,  $p < 0.001$ , as well as an interaction between condition and whether the item was repeated,  $F_{2,140.00} = 122.59$ ,  $p < 0.001$ . This indicated that the forward phase suffered from much stronger interference than the backward phase. In the *forward interference* condition, the repeated second event had an approximately 18% higher



**Fig. 4** Ordering of two-item sequences on repetition trials. **a** Time courses (in TRs from sequence onset) of probabilistic classifier evidence (in %) in repetition trials, color-coded by event type (first, second and the three remaining non-sequence items, see legend). Data shown separately for forward (left) and backward (right) interference conditions. Gray background indicates relevant time period independently inferred from response functions (Fig. 2d). Shaded areas represent  $\pm 1$  SEM. 1 TR = 1.25 s. **b** Mean probability of event types averaged across all TRs in the relevant time period, as in (a). Each dot represents one participant, the probability density of the data is shown as rain cloud plots (cf.<sup>141</sup>). Boxplots indicate the median and interquartile range (IQR, i.e., distance between the first and third quartiles). The lower and upper hinges correspond to the first and third quartiles (the 25th and 75th percentiles). The upper whisker extends from the hinge to the largest value no further than  $1.5 \times$  IQR from the hinge. The lower whisker extends from the hinge to the smallest value at most  $1.5 \times$  IQR of the hinge. The diamond shapes show the sample mean and error bars indicate  $\pm 1$  SEM ( $N = 36$ ,  $t_s \geq 3.31$ ,  $p_s \leq 0.006$ , LME model with post hoc Tukey’s honest significant difference (HSD) tests). **c** Average probability of event types, separately for forward/backward conditions as in (a), plots as in (b) ( $N = 36$ ,  $t_s \geq 4.14$ ,  $p_s < 0.001$ , LME model with post hoc Tukey’s HSD tests). **d** Mean trial-wise proportion of each transition type, separately for forward/backward conditions, as in (a) ( $N = 36$ ,  $t_s \geq 4.64$ ,  $p_s < 0.001$ , four two-sided paired  $t$ -tests, Bonferroni-corrected). **e** Transition matrix of decoded images indicating mean proportions per trial, separately for forward/backward conditions, as in (a). Transition types highlighted in colors (see legend). All statistics have been derived from data of  $N = 36$  human participants who participated in one experiment. Source data are provided as a Source Data file.

probability than the single first event (31.55% vs. 13.50%,  $p < 0.001$ ). In the *backward interference* condition, the repeated first event had only 9% higher probability than the single second event (26.87% vs. 18.00%,  $p < 0.001$ , corrected). This means that the

item shown only once was easier to detect when it followed a sustained activation of a different pattern, compared to when it preceded an interfering activation (Fig. 4c). We found no main effect of repetition,  $p = 0.91$  (Fig. 4c).

Importantly, however, both sequence elements still differed from non-sequence items even under conditions of interference (forward: 7.75% and backward: 7.69%, respectively, all  $p$ s < 0.001, corrected), indicating that sequence element detection remains possible under such circumstances. Using data from all TRs revealed qualitatively similar significant effects ( $p \leq 0.04$  for all but one test after correction, see SI). Repeating all analyses using proportions of decoded classes (the class with the maximum probability was considered decoded at every TR), or considering all repetition trial conditions, also revealed qualitatively similar results. Thus, brief events can be detected despite significant interference.

We next asked which implications these findings have for the observed pattern transitions (cf.<sup>54</sup>). To this end, we analyzed the trial-wise proportions of transitions between consecutively decoded events, and asked whether forward transitions between sequence items were more likely than transitions between a sequence and a non-sequence item (*outward transitions*) or between two non-sequence items (*outside transitions*; for details, see Methods). This analysis revealed that forward transitions (5.89%) were more frequent than both outward transitions (2.46%), and outside transitions (1.04%, both  $p$ s < 0.001,  $t$ s  $\geq 4.64$ , Bonferroni-corrected; Fig. 4d) in the forward interference condition. The same was true in the backward interference condition (forward transitions: 7.22%; outward transitions: 2.67%; outside transitions: 1.06%, all  $p$ s < 0.001,  $t$ s  $\geq 5.14$ ). The full transition matrix is shown in Fig. 4e. Repetitions of the first or second item are shown on the upper two diagonal elements (with all consecutive repetitions of items labeled *repetition* in Fig. 4e), and were not considered in this analysis.

Together, the results from repetition trials indicated that: (1) within-sequence items could be clearly detected despite interference from other sequence items; (2) event detection was asymmetric, such that items occurring at the end of sequences can be detected more easily than those occurring at the beginning; and (3) the detection of sequence items made it possible to observe within-sequence transitions between decoded items.

Note that our analyses focused on the two extreme cases of repetition trials with one versus eight repetitions of the first or second item while the experiment also included repetition trials with intermediate levels of repetitions (see SI). Specifically, other repetition trials included cases in which the second item began to appear at each possible position from 2 to 9. The other repetition trials could therefore include, for instance, three repetitions of the first and six repetitions of the second image, or four repetitions of the first and five repetitions of the second item, etc. The results reported in the SI indicate that effects in these trials show smooth transition between the extremes shown in the main manuscript.

*Detecting sparse sequence events with lower signal-to-noise ratio (SNR).* The results above indicate that detection of fast sequences is possible if they are under experimental control. In most applications of our method, however, this will not be the case. When detecting replay, for instance, sequential events will occur spontaneously during a period of noise. We therefore next assessed the usefulness of our method under such circumstances.

We first characterized the behavior of sequence detection metrics during periods of noise. To this end, we applied the logistic regression classifiers to fMRI data acquired from the same participants ( $n = 32$  out of 36) during a 5-min (233 TRs) resting period before any task exposure. Classifier probabilities during rest fluctuated wildly, often with a single category having a high probability, while all other categories had probabilities close to zero. During fast sequence periods, in contrast, the near-simultaneous activation of stimulus-driven activity led to reduced

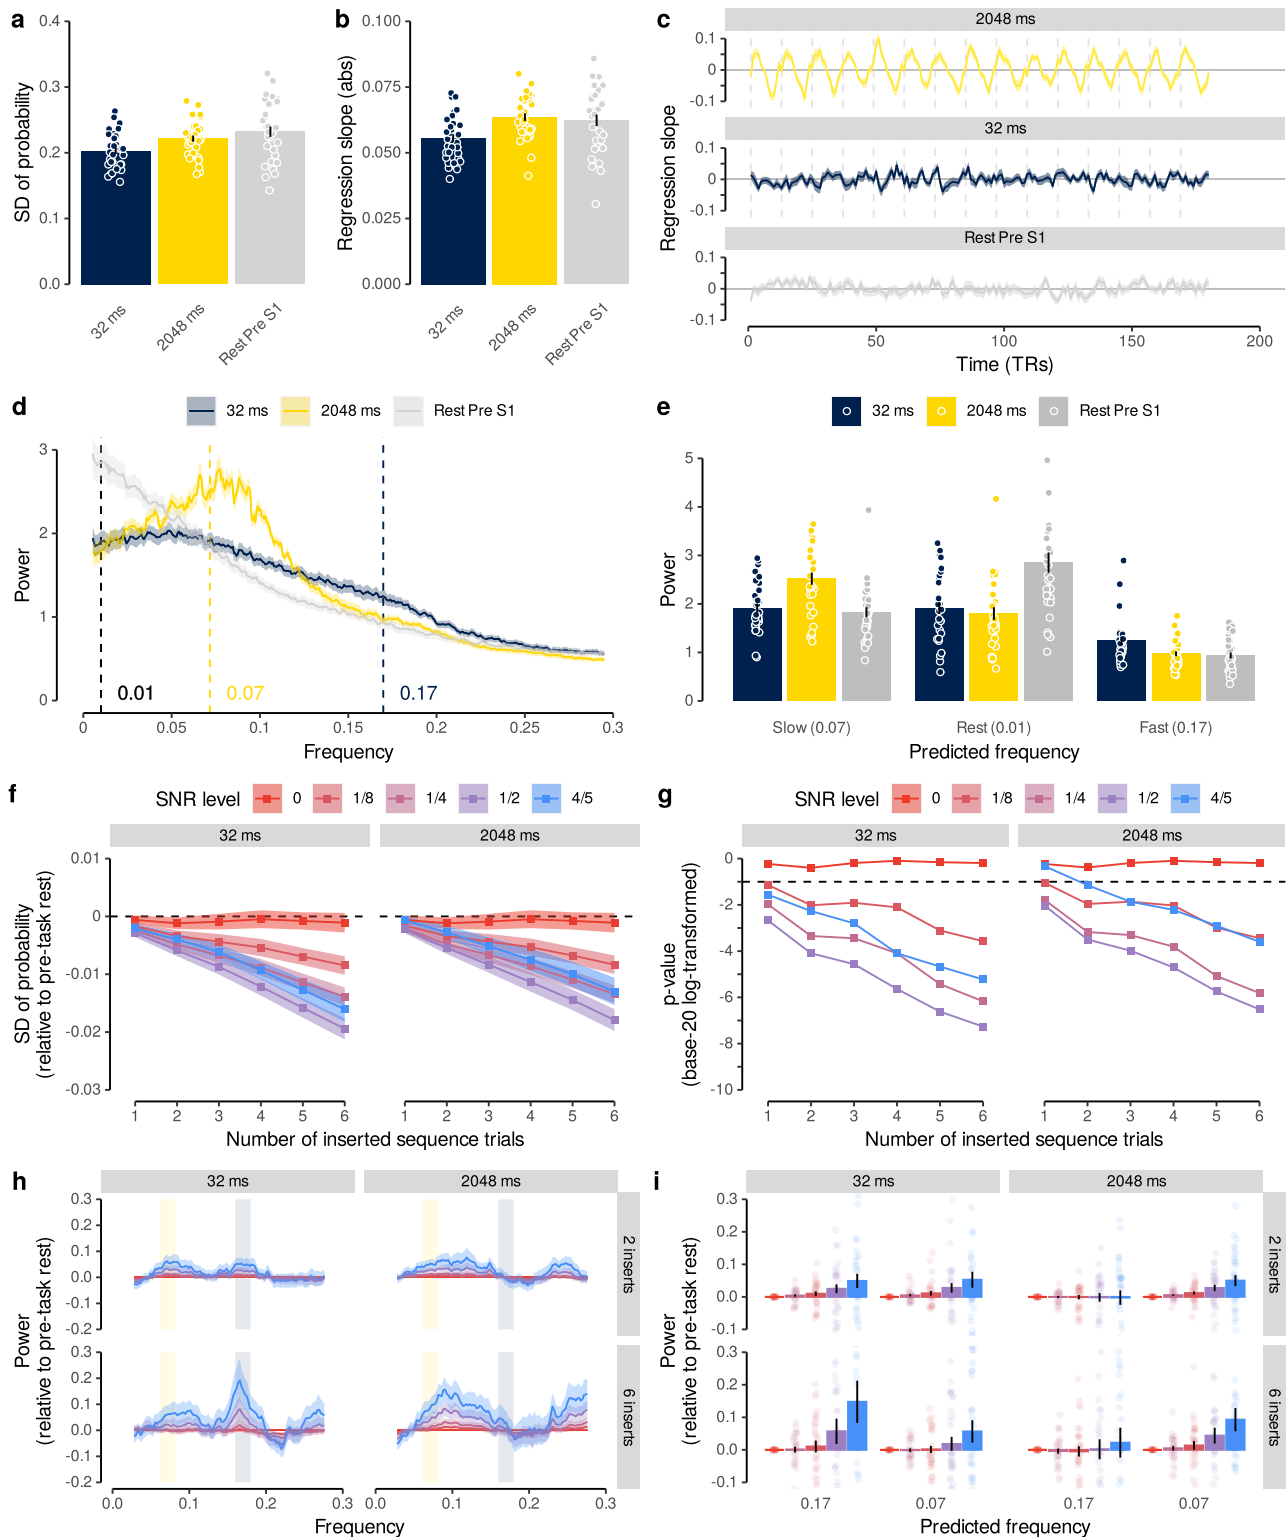
probabilities, such that category probabilities tended to be closer together and less extreme. In consequence, the average standard deviation of the probabilities per TR during rest and slow (2048 ms) sequence periods was higher ( $M = 0.23$  and  $M = 0.22$ , respectively) compared to the average standard deviation in the fast sequence condition (32 ms;  $M = 0.20$ ;  $t$ s  $\geq 4.17$ ;  $p$ s < 0.001;  $d$ s  $\geq 0.74$ ; Fig. 5a).

As before, we fitted regression coefficients through the classifier probabilities of the rest data and, for comparison, concatenated data from the 32 and 2048 ms sequence trials (Fig. 5b, c). As predicted by our modeling approach (Fig. 2e), and shown in the previous section (Fig. 3b), the time courses of regression coefficients in the sequence conditions were characterized by rhythmic fluctuations whose frequency and amplitude differed between speed conditions (Fig. 5c). To quantify the magnitude of this effect, we calculated frequency spectra of the time courses of the regression coefficients in rest and concatenated sequence data (Fig. 5d; using the Lomb-Scargle method, e.g.<sup>61</sup> to account for potential artifacts due to data concatenation, see Methods). This analysis revealed that frequency spectra of the sequence data differed from rest frequency spectra in a manner that depended on the speed condition (Fig. 5d, e). As foreshadowed by our model, power differences appeared most pronounced in the predicted frequency ranges (Fig. 5e;  $p$ s  $\leq 0.002$ ; see Eq. (5) and Methods). Specifically, when the 32 ms condition was considered, the analyses revealed an increased power around 0.17 Hz, which corresponds to the frequency predicted to occur by our model. Data from the 2048 ms condition, in contrast, exhibited an increased power around 0.07 Hz, as predicted.

Finally, we asked whether these differences would persist if (a) only few sequence events occurred during a 5-min rest period, while (b) their onset was unknown, and (c) their SNR was lower. To this end, we synthetically generated data containing a variable number of sequence events that were inserted at random times into the resting-state data acquired before any task exposure. Specifically, we inserted between 1 and 6 sequence events into the rest period by blending rest data with TRs recorded in fast (32 ms) or slow (2048 ms) sequence trials (12 TRs per trial, random selection of sequence trials and insertion of time points, without replacement). To account for possible SNR reductions, the inserted probability time courses were multiplied by a factor  $\kappa$  of  $\frac{4}{5}$ ,  $\frac{1}{2}$ ,  $\frac{1}{4}$ ,  $\frac{1}{8}$ , or 0 and added to the probability time courses of the inversely scaled  $(1-\kappa)$  resting-state data. Effectively, this led to a step-wise reduction of the inserted sequence signal from 80% to 0%, relative to the SNR obtained in the experimental conditions reported above. Thus, here we use the term SNR to describe the relative mixing proportion of (a) data from the task, which contain sequential signal, with (b) data from the pre-task resting-state session, which contain only noise. Note that this is different from the common definition of SNR in univariate fMRI as the ratio of average signal to standard deviation over time.

As expected, differences in the above-mentioned standard deviation of the probability gradually increased with both the SNR level and the number of inserted sequence events when either fast or slow sequences were inserted (Fig. 5f). In our case, this led significant differences to emerge with one insert and an SNR reduced to 12.5% in both the fast and slow conditions (Fig. 5g; comparing against zero, the expectation of no difference with a conventional false-positive rate  $\alpha$  of 5%; all  $p$ s FDR-corrected).

Importantly, the presence of sequence events was also reflected in the frequency spectrum of the regression coefficients. Inserting fast event sequences into rest led to power increases in the frequency range indicative of 32 ms events ( $\sim 0.17$  Hz, Fig. 5h, i, left panel), in line with our findings above. This effect again got



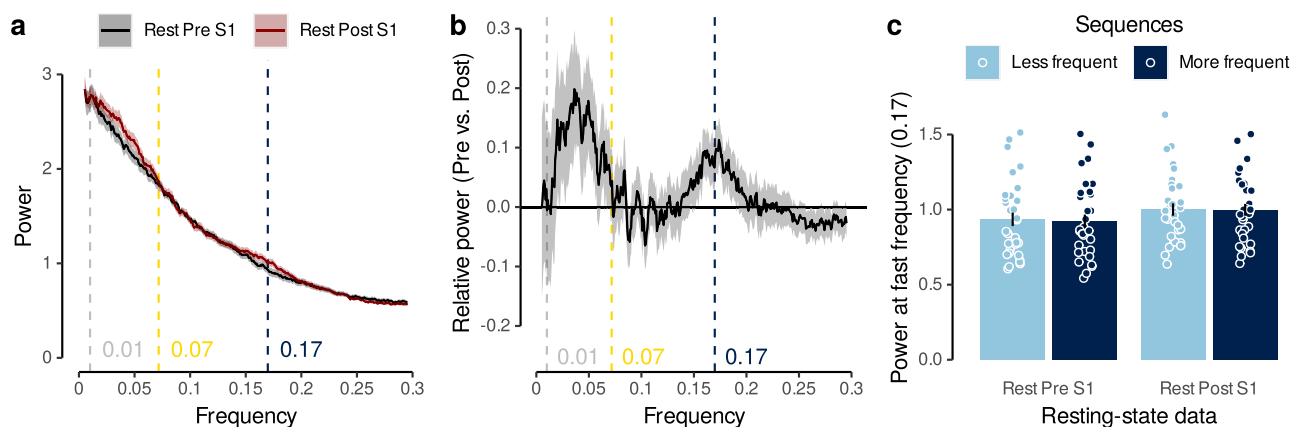
stronger with higher SNR levels and more sequence events. Inserting slow (2048 ms) sequence events into the rest period showed a markedly different frequency spectrum, with an increase around the frequency predicted for this speed (~0.07 Hz, Fig 5h, i, right panel). Comparing the power around the predicted frequency ( $\pm 0.01$  Hz) of both speed conditions indicated significant increases in power compared to sequence-free rest when six sequence events were inserted and the SNR was reduced to 80% ( $t_s \geq 2.28$ ,  $p_s \leq 0.03$ ,  $d_s \geq 0.40$ ). Hence, the

presence of spontaneously occurring sub-second sequences during rest can be detected in the frequency spectrum of our sequentiality measure, and distinguished from slower second-scale sequences that might reflect conscious thinking.

*Detecting fast reactivations in post-task resting-state data.* Finally, we asked whether our task elicited spontaneous replay of image sequences in object-selective brain areas during rest after the task. Based on the above findings, we reasoned that potentially



**Fig. 5 Detecting sparse sequence events with lower SNR.** **a** Mean standard deviation of classifier probabilities in rest and sequence data ( $n = 32$ ,  $t_s \geq 4.17$ ,  $p_s < 0.001$ ,  $d_s \geq 0.74$ , two two-sided paired  $t$ -tests comparing rest and 2048 ms conditions against 32 ms condition, FDR-corrected). **b** Mean absolute regression slopes, as in (a) ( $n = 32$ ,  $t_s \geq 4.64$ ,  $p_s < 0.001$ ,  $d_s \geq 0.82$ , two two-sided paired  $t$ -tests comparing rest and 2048 ms conditions against 32 ms condition, FDR-corrected). **c** Time courses of the regression slopes (signed values, not magnitudes) in rest and sequence data. Vertical lines indicate trial boundaries. **d** Normalized frequency spectra of regression slopes in rest and sequence data. Annotations indicate predicted frequencies based on Eq. (5). **e** Mean power of predicted frequencies in rest and sequence data, as in (a). Each dot represents data from one participant ( $n = 32$ ,  $t_s \geq 3.10$ ,  $p_s \leq 0.002$ , two-sided paired  $t$ -tests, FDR-corrected). **f** Mean standard deviation of rest data including a varying number of SNR-adjusted sequence events (fast or slow). Dashed line indicates indifference from sequence-free rest ( $n = 32$ ,  $t_s \geq 2.22$ ,  $p_s \leq 0.04$ , 30 two-sided one-sample  $t$ -tests against chance, FDR-corrected). **g** Base-20 log-transformed  $p$  values of  $t$ -tests comparing the standard deviation of probabilities in (f) with sequence-free rest. Dashed line indicates  $p = 0.05$  ( $N = 32$ ,  $t_s \geq 2.22$ ,  $p_s \leq 0.04$ , 30 two-sided one-sample  $t$ -tests against chance, FDR-corrected). **h** Frequency spectra of regression slopes in SNR-adjusted sequence-containing rest relative to sequence-free rest. Rectangles indicate predicted frequencies, as in (d). **i** Mean relative power of predicted frequencies in SNR-adjusted sequence-containing rest ( $n = 32$ ,  $t_s \geq 2.28$ ,  $p_s \leq 0.03$ , two-sided  $t$ -tests against baseline, FDR-corrected). Shaded areas/error bars represent  $\pm 1$  SEM. All statistics have been derived from data of  $n = 32$  human participants who participated in one experiment. 1 TR = 1.25 s. Source data are provided as a Source Data file.



**Fig. 6 Detecting fast task-related reactivations in post-task resting-state data.** **a** Normalized frequency spectra of regression slopes in pre- and post-task resting-state data. Annotations indicate predicted frequencies based on Eq. (5). Shaded areas represent  $\pm 1$  SEM. **b** Relative power (difference between pre- and post-task rest) of normalized frequency spectra shown in (a) ( $n = 32$ ,  $F_{1,94.99} = 6.17$ ,  $p = 0.02$ , LME model comparing pre- vs. post-task resting-state data at 0.17 Hz). **c** Mean power at predicted fast frequency (0.17 Hz) in pre- and post-task resting-state data for less and more frequent stimulus sequences ( $n = 32$ ,  $t_s \geq 4.17$ ,  $p_s < 0.001$ ,  $d_s \geq 0.74$ , two two-sided paired  $t$ -tests comparing rest and 2048 ms conditions against 32 ms condition, FDR-corrected). Each dot corresponds to averaged data from one participant. Error bars represent  $\pm 1$  SEM. All statistics have been derived from data of  $n = 32$  human participants who participated in one experiment. Source data are provided as a Source Data file.

reactivated sequences should become apparent in a frequency spectrum analysis. We therefore applied this analysis to resting-state data recorded after participants performed the task. Crucially, because the true sequence of potential replay events was not known, we repeated the analyses for all possible image orders, averaged the resulting frequency spectra, and compared the results to the same analysis performed on the pre-task rest session (see Methods). As shown in Fig. 6, the frequency spectrum analyses revealed a significant increase specifically in the power spectrum of the high frequency range (Fig. 6a,  $F_{1,94.99} = 6.17$ ,  $p = 0.02$  when testing pre- versus post-task data at the predicted frequency of 0.17 Hz, as before). Directly comparing pre- versus post-task rest revealed a large power difference at 0.17 Hz, indicative of replayed sequence speeds of 32 ms, as in our fastest sequence speed condition (Fig. 6b). In addition, we found a second peak at around 0.04 Hz, indicating long activations of individual items of several seconds. Thus, post-task rest seemed to be characterized by fast sequential reactivations as well as longer constant activations. We next asked whether specific sequences that had been experienced slightly more often by participants were more likely to be reactivated than less frequent sequences. During slow trials, all participants experienced all 120 possible sequential combinations of images. But in addition, each participant experienced only a subset of 15 image orders during the sequence trials. Hence, image orders experienced in sequence

trials were slightly more frequent and we asked if they were reactivated more strongly during the post-task resting-state session. This was not the case. A power increase in the fast frequency range when comparing pre- to post-task rest was found for both sets of sequences, i.e., the 15 image orders that occurred in sequence and slow trials, and the 105 that occurred only in slow trials (Fig. 6c,  $p_s \geq 0.13$ ). In summary, applying the frequency spectrum analyses to post-task resting-state therefore suggests that (1) task stimuli are reactivated during post-task rest, and (2) this reactivation happens fast, but (3) appears unspecific and not directly related to the sequences presented more frequently to participants during the task.

### Discussion

Here, we demonstrated that BOLD fMRI in combination with multivariate probabilistic decoding can be used to detect sub-second sequences of closely timed neural events non-invasively in humans. We combined probabilistic multivariate pattern analysis with time course modeling and investigated human brain activity recorded following the presentation of sequences of visual objects at varying speeds, as well as activity during rest. In the fastest case a sequence of five images was displayed within 628 ms (32 ms between pictures). Stimulus sequences were not masked. Even when using a TR of 1.25 s, achievable with conventional multi-

band (MB) echo-planar imaging (EPI), the image order could be detected from activity patterns in visual and ventral temporal cortex. Detection of briefly presented sequence items was also possible when their activation was affected by interfering signals from a preceding or subsequent sequence item and could be differentiated from images that were not part of the sequence. Our results withstood several robustness tests, and also indicated that detection is biased to most strongly reflect the last event of a sequence. Analyses of augmented resting data, in which neural event sequences occurred rarely, at unknown times, and with reduced signal strength, showed that our method could detect sub-second sequences even under such adverse conditions. Moreover, we showed that frequency spectrum analyses can be used to distinguish sub-second from supra-second sequences under such circumstances. Our approach therefore promises to expand the scope of BOLD fMRI to fast, sequential neural representations by extending multivariate decoding approaches into the temporal domain, in line with our previous findings<sup>54</sup>.

Importantly, we applied this method not only to experimentally controlled data, but also used it to ask whether task experience might elicit spontaneous replay of sequential stimuli in post-task resting-state data, as suggested by previous studies, for reviews, see, e.g.<sup>33,35,62,63</sup>. Indeed, our results indicate that such reactivations occur during post-task rest and can be detected using the proposed analysis. Our analyses suggest that the reactivated sequences were fast and occurred at replay-like speeds, similar to the fastest sequence trials used in our task (32 ms between activations). Evidence for fast sequential replay was accompanied by a relative increase in power in the slower frequency range (peaking at 0.04 Hz). This could reflect an increase of slower long-lasting activations, possibly reflecting conscious thinking about the task. This supports our conclusion that the frequency spectrum of the sequentiality metric is a useful approach to detect fast replay and to distinguish it from slow activations. Our analysis did not find any evidence that only those sequences were replayed that were more frequent than others or that were presented at a fast speed during the task. Rather, our results suggest that replay seemed to equally involve all stimulus orders. However, it is important to note that our task was not optimized to elicit replay of particular sequences at all. In fact, the more frequent sequences were arranged such that the same stimuli appeared equally often at the first and last position, which makes it difficult to distinguish them from other sequences.

Of note, replay during post-task rest reflected cortical reactivations in occipito-temporal brain regions. Given that we were not able to decode on-task stimulus representations in the hippocampus, it remains unclear if reactivations occurred independently from (task-related) involvement of the hippocampus or if we were simply not able to detect concurrent reactivation in the hippocampus. This possibility of hippocampus-independent cortical reactivations raises important questions regarding the functional significance of such events. One potential reason why we found no hippocampus involvement could be that the oddball detection paradigm used for slow trials to train the classifiers involved no mnemonic task component, and therefore was not suitable to activate the hippocampus. Our previous work<sup>54</sup> has already demonstrated the success of our methods in hippocampal data. Taken together, our results indicate that our method allows the uncovering of fast task-related reactivations during rest and highlight the importance of task design for detecting replay in humans using fMRI.

This contrasts with previous fMRI studies in humans (for reviews see, e.g.<sup>24,64</sup>) that measured non-sequential reactivation as increased similarity of multi-voxel patterns during experience and extended post-encoding rest compared to pre-encoding baseline<sup>49–51,53,65–69</sup> or functional connectivity of hippocampal,

cortical, and dopaminergic brain structures that support post-encoding systems-level memory consolidation<sup>66–68,70–72</sup>. In the current study we open the path toward a better understanding of the speed and sequential nature of the observed phenomena.

The fastest sequences studied in our experiments lasted 628 ms and were therefore longer than the average hippocampal replay event of about 300 ms, e.g.<sup>17</sup>. Yet, several factors support the idea that our method is still relevant for the study of replay. First, previous studies have shown that a significant proportion of replay events indeed lasts much longer than 300 ms. Davidson et al.<sup>7</sup> report sequence lengths of up to 1000 ms and the data by Kaefer et al.<sup>17</sup> indicate that about 20% of events in the hippocampus are longer than 500 ms. In addition, the median duration of replay events in medial prefrontal cortex (PFC) reported in Kaefer et al.<sup>17</sup> was 740 ms. This indicates that a significant proportion of replay events will be covered by our method. Second, our ISI was as fast as 32 ms, which corresponds to the time lag between activations reported in magnetoencephalography (MEG) studies (e.g.<sup>47</sup>) and therefore might capture the important aspect of temporal separation between activation patterns well. Third, while effect sizes showed a pronounced decrease when comparing the slower conditions (2048 ms: 3.13; 512 ms: 1.36; 128 ms: 0.75, for the backwards effect of regression slopes, Fig. 3c, effect sizes indicate Cohen's *d*), accelerating sequence speeds beyond 128 ms seemed not to be associated with a comparable decrease in effect sizes (64 ms: 0.65, 32 ms: 1.00). This indicates that the sensitivity of our methods for even faster event sequences might not be catastrophically diminished. Fourth, the sequence duration of 628 ms was to a large extent due to the stimulus duration of 100 ms. Evidence from previous work using electroencephalography (EEG) suggests that the neural response to successive visual stimuli is more strongly influenced by the ISI than the stimulus duration<sup>73,74</sup>. Hence, we speculate that our methods may also work in cases with shorter pattern activations and thereby overall shorter sequences.

Our results deepen the understanding of our previous findings<sup>54</sup> in two ways. First, we provide additional empirical evidence that our sequentiality analyses based on multivariate fMRI pattern classification are indeed sensitive to fast neural event sequences. To this end, we used an experimental setup where the order of sequential events is known—in contrast to analyses of resting-state data in Schuck and Niv<sup>54</sup> where the order and speed of event sequences can only be assumed. Second, Schuck and Niv<sup>54</sup> observed forward-ordered replay. Our present study clarifies the origins of forward and backward ordering of fMRI activation patterns. We show that probabilistic classifier evidence in earlier TRs reflects the forward order of the sequences while this pattern reverses in later TRs. Importantly, we demonstrate an asymmetry in decoding early versus late sequential events. This can therefore lead fMRI pattern sequences to appear in the reverse order relative to the underlying neural sequences. This represents a crucial insight, given the different functional roles assigned to forward and backward replay (see e.g.<sup>33</sup>). We note that future research should be careful when interpreting directionality, as the relationship between decoded and true directionality is not straightforward. One approach in this context could be to investigate the order of sequence direction itself. If items appear to be ordered first in direction A, and a few TRs later in direction B, then direction A seems to be the true one. Probabilistic classifiers might prove particularly useful for such analyses as they make it possible to characterize sequential ordering within a single measurement. The origins of this asymmetry are not entirely clear. It seems possible that they reflect the benefits of the last item not being followed by another activation that could impede its detection. A relation to the asymmetric shape of the HRF, to changing HRF variability with

time and even to inhibitory retrograde neurotransmitters (e.g.<sup>75</sup>) cannot be ruled out. Third, we have shown that the interference of activation patterns of fast sequential neural events is stronger for early events compared to late events. Importantly, early events remained detectable despite this interference, demonstrating that our method can detect the elements of a replay event with fMRI despite interference effects. The prominence of the last sequence item implies that the apparent over-occurrence of one particular item might reflect that this item was a frequent start or end point of replayed trajectories. Past research has shown that task aspects, such as goals, heavily influence which items the replayed sequences start or end with<sup>76</sup>.

In addition, our study introduces important methodological advancements that go beyond our original publication<sup>54</sup>. We show that the analyses of classifier probabilities provide major statistical improvements compared to analyses focused on the decoded category with the highest classification probability (as in Schuck and Niv<sup>54</sup>). The key advantage is that probabilistic classifiers provide a continuous metric of classification evidence and thereby allow the detection of sequential ordering within a single measurement (i.e., within a single TR). This results in significant information gain compared to the assessment of sequential ordering that considers only a single label per TR. Moreover, we leverage frequency spectrum analysis in an approach to make inferences about the speed of the sequential neural process. Although the sampling rate (i.e., the TR) of fMRI is usually less than the speed of replay events, frequency spectrum analyses can characterize the speed of fast sequential events during rest. Together, these methodological advances offer insights into previous fMRI studies investigating hippocampal replay in humans, including our own work<sup>54</sup>.

Additionally, some caveats have to be noted. Our results indicate that the sequentiality in fMRI analyses is mainly influenced by the first and last element of a fast sequence. Given that replay events are often structured by task-relevant features like the start and goal location in a spatial environment (e.g.,<sup>76</sup>), analyzing the transitions between the corresponding decoded events will offer insights into the content and functional role of fast replay events. Moreover, it is important to keep in mind that the benefits of our experimental setting came at the cost that they also introduced important differences from a replay study in various regards, including the focus on extra-hippocampal activations and sensory stimulation.

Our fMRI-based approach has advantages as well as disadvantages compared to existing EEG and MEG approaches<sup>44,46,47</sup>. In particular, it seems likely that our method has limited resolution of sequence speed. While we could distinguish between supra- and sub-second sequences, a much finer distinction might prove difficult in practice. Yet, EEG and MEG investigations suggest that the extent of temporal compression of previous experience is an important aspect of replay and other reactivation phenomena<sup>45,77–80</sup>. In addition, the differential sensitivity to activity depending on sequence position complicates interpretations of findings, and can lead to statistical aliasing of sequences with the same start and end elements but different elements in the middle. Finally, because a single sequence causes forward and backward ordering of signals, it can be difficult to determine the direction of a hypothesized sequence. One major advantage of fMRI is that it does not suffer from the low sensitivity to hippocampal activity and limited ability to anatomically localize effects that characterizes EEG and MEG. This is particularly important in the case of replay, which is hippocampus-centered but co-occurs with fast neural event sequences in other parts of the brain including primary visual cortex<sup>12</sup>, auditory cortex<sup>15</sup>, PFC<sup>13,14,16,17,81</sup>, entorhinal cortex<sup>22,82,83</sup>, and ventral striatum<sup>84</sup>. Importantly, replay events occurring in different brain areas might not be mere copies of each other, but can differ regarding their timing, content, and

relevance for cognition, e.g.<sup>16,17</sup>. Precise characterization of replay events occurring in different anatomical regions is therefore paramount. The present finding of fast and slow reactivations in visual cortex underlines the importance of knowing the anatomical origin of replay events. Because EEG and MEG cannot untangle the co-occurring events and animal research is often restricted to a single recording site, much remains to be understood about the distributed and coordinated nature of replay. One particular problem is that localizer tasks frequently used to train classifiers in MEG studies might only partially reflect hippocampal activity. In fact, our own data here show that simple visual tasks do not elicit reliable hippocampal activation patterns. Thus, EEG or MEG classifiers trained on such data risk to not reflect any hippocampal activity.

Finally, our study provides additional insights for future research. We have shown that the mere fact that detecting which elements were part of a sequence is beneficial if sequences mostly contain a local subset of all possible events. Thus, experimental setups with a larger number of possible events will be insightful. At the same time, a larger number of to-be-decoded events will likely impair baseline classification accuracy, which in turn impairs sequence detection. Researchers should thus take the trade-off between these two aspects into account. Moreover, several other factors could influence the success of future investigations: the sampling rate (the TR); the choice of brain region; and the properties of the resulting HRFs<sup>23</sup>. Whether an increased sampling rate would be beneficial for the detection of fast event sequences is difficult to predict. First, longer TRs provide better SNR as they allow more time for longitudinal magnetization. In addition, faster sampling will not affect the underlying (slow) HRF dynamics that impede the identification of temporal order of fast neural event sequences. Sampling the activation time courses at a faster rate might not reveal more information about the sequential process under investigation. Whether shorter TRs can make up for the downsides in spatial resolution and SNR therefore seems an empirical question. Moreover, the choice of brain region will impact results only if the stability of the HRF within that brain region is low, whereas between-region differences between HRF parameters might have less impact. But HRF stability is generally high<sup>30,85–87</sup>, and previous research noting this fact has therefore already indicated possibilities of disentangling temporally close events<sup>28–31,88,89</sup>. Further, increased spatial resolution might improve detection due to less partial volume averaging of non-activation-related signals. Our approach has shown how multivariate and modeling approaches can help exploit these HRF properties in order to enhance our understanding of the human brain.

## Methods

**Participants.** In all, 40 young and healthy adults were recruited from an internal participant database or through local advertisement and fully completed the experiment. No statistical methods were used to predetermine the sample size but it was chosen to be larger than similar previous neuroimaging studies, e.g.<sup>51,52,54</sup>. Four participants were excluded from further analysis because their mean behavioral performance was below the 50% chance level in either or both the sequence and repetition trials suggesting that they did not adequately process the visual stimuli used in the task. Please note that this exclusion was based on mean behavioral performance across all conditions of sequence and repetition trials. This means that participants who, for example, performed below chance in only one of the conditions of either the sequence or the repetition trials, but above chance in all other conditions, might still be included in the final sample because their mean behavioral performance across all conditions ended up to be above the level of chance performance. Thus, the final sample consisted of 36 participants (age:  $M = 24.61$  years,  $SD = 3.77$  years, range: 20–35 years, 20 female, 16 male). All participants were screened for magnetic resonance imaging (MRI) eligibility during a telephone screening prior to participation and again at the beginning of each study session according to standard MRI safety guidelines (e.g., asking for metal implants, claustrophobia, etc.). None of the participants reported to have any major physical or mental health problems. All participants were required to be right-handed, to have corrected-to-normal vision, and to speak German fluently. Furthermore, only participants with a head circumference of 58 cm or less could be

included in the study. This requirement was necessary as participants' heads had to fit the MRI head coil together with MRI-compatible headphones that were used during the experimental tasks. The ethics commission of the German Psychological Society (DGPs) approved the study protocol (reference number: NS 012018). All volunteers gave written informed consent prior to the beginning of the experiments. Every participant received 40.00 Euro and a performance-based bonus of up to 7.20 Euro upon completion of the study. None of the participants reported to have any prior experience with the stimuli or the behavioral task.

### Task

**Stimuli.** All stimuli were gray-scale images of a cat, chair, face, house, and shoe taken from Haxby et al.<sup>55</sup> with a size of  $400 \times 400$  pixels each, which have been shown to reliably elicit object-specific neural response patterns in several previous studies, e.g.<sup>55,58–60</sup>. Participants received auditory feedback to signal the accuracy of their responses. A high-pitch coin sound confirmed correct responses, whereas a low-pitch buzzer sound signaled incorrect responses. The sounds were the same for all task conditions and were presented immediately after participants entered a response or after the response time had elapsed. Auditory feedback was used to anatomically separate the expected neural activation patterns of visual stimuli and auditory feedback. While auditory feedback is more likely to engage primarily temporal brain regions, visual stimuli are more likely to activate primarily occipital brain regions. We recorded the presentation time stamps of all visual stimuli and confirmed that all experimental components were presented as expected. The task was programmed in MATLAB (version R2012b; Natick, MA, USA; The MathWorks Inc.) using the Psychophysics Toolbox extensions (Psychtoolbox; version 3.0.11)<sup>90–92</sup> and run on a Windows XP computer with a monitor refresh-rate of 16.7 ms.

**Slow trials.** The slow trials of the task were designed to elicit object-specific neural response patterns of the presented visual stimuli. The resulting patterns of neural activation were later used to train the classifiers. In order to ensure that participants maintained their attention and processed the stimuli adequately, they were asked to perform an oddball detection task (for a similar approach, see<sup>44,47</sup>). Specifically, participants were instructed to press a button each time an object was presented upside-down. Participants could answer using either the left or the right response button of an MRI-compatible button box. In contrast to similar approaches, e.g.,<sup>44,47</sup>, we intentionally did not ask participants for a response on trials with upright stimuli to avoid neural activation patterns of motor regions in our training set which could influence later classification accuracy on the test set.

Participants were rewarded with 3 cents for each oddball (i.e., stimulus presented upside-down) that was correctly identified (i.e., hit) and punished with a deduction of 3 cents for (incorrect) responses (i.e., false alarms) on non-oddball trials (i.e., when stimuli were presented upright). In case participants missed an oddball (i.e., miss), they also missed out on the reward. Auditory feedback (coin and buzzer sound for correct and incorrect responses, respectively) was presented immediately after the response (in case of hits and false alarms) or at the end of the response time limit (in case of misses) using MRI-compatible headphones (VisuaStimDigital, Resonance Technology Company Inc., Northridge, CA, USA). Correct rejections (i.e., no responses to upright stimuli) were not rewarded and were consequently not accompanied by auditory feedback. Together, participants could earn a maximum reward of 3.60 Euro in this task condition.

Across the entire experiment, all five unique images were presented in all possible sequential combinations which resulted in  $5! = 120$  sequences with each of the five unique visual objects in a different order. Thus, across the entire experiment, participants were shown  $120 \times 5 = 600$  visual objects in total for this task condition. Of all visual objects, 20% were presented upside-down (i.e., 120 oddball stimuli). All unique visual objects were shown upside-down equally often, which resulted in  $120/5 = 24$  oddballs for each individual visual object category. The order of sequences as well as the appearances of oddballs were randomly shuffled for each participant and across both study sessions.

Each trial (for the trial procedure, see Fig. 1a) started with a waiting period of 3.85 s during which a blank screen was presented. This ITI ensured a sufficient time delay between each slow trial and the preceding trial (either a sequence or a repetition trial). The five visual object stimuli of the current trial were then presented as follows: after the presentation of a short fixation dot for a constant duration of 300 ms, a stimulus was shown for a fixed duration of 500 ms followed by a variable ISI during which a blank screen was presented again. The duration of the ISI for each trial was randomly drawn from a truncated exponential distribution with a mean of 2.5 s and a lower limit of 1 s. We expected that neural activation patterns elicited by the stimuli can be well recorded during this average time period of 3 s (for a similar approach, see<sup>55</sup>). Behavioral responses were collected during a fixed time period of 1.5 s after each stimulus onset. In case participants missed an oddball target, the buzzer sound (signaling an incorrect response) was presented after the response time limit had elapsed. Only neural activation patterns related to correct trials with upright stimuli were used to train the classifiers. Slow trials were interleaved with sequence and repetition trials such that each of the 120 slow trials was followed by either one of the 75 sequence trials or 45 repetition trials (details on these trial types are given below).

**Sequence trials.** In the sequence trials of the task, participants were shown sequences of the same five unique visual objects at varying presentation speeds. In total, 15 different sequences were selected for each participant. Sequences were chosen such that each visual object appeared equally often at the first and last position of the sequence. Given five stimuli and 15 sequences, for each object category this was the case for 3 out of the 15 sequences. Furthermore, we ensured that all possible sequences were chosen equally often across all participants. Given 120 possible sequential combinations in total, the sequences were distributed across eight groups of participants. Sequences were randomly assigned to each participant following this pseudo-randomized procedure.

To investigate the influence of sequence presentation speed on the corresponding neural activation patterns, we systematically varied the ISI between consecutive stimuli in the sequence. Specifically, we chose five different speed levels of 32, 64, 128, 512, and 2048 ms, respectively (i.e., all exponents of 2 for good coverage of faster speeds). Each of the 15 sequences per participant was shown at each of the 5 different speed levels. The occurrence of the sequences was randomly shuffled for each participant and across sessions within each participant. This resulted in a total of 75 sequence trials presented to each participant across the entire experiment. To ensure that participants maintained attention to the stimuli during the sequence trials, they were instructed to identify the serial position of a previously cued target object within the shown stimulus sequence and indicate their response after a delay period without visual input.

During a sequence trial (for the trial procedure, see Fig. 1b) the target cue (the name of the visual object, e.g., *shoe*) was shown for a fixed duration of 1000 ms, followed by a blank screen for a fixed duration of 3850 ms. A blank screen was used to reduce possible interference of neural activation patterns elicited by the target cue with neural response patterns following the sequence of visual objects. A short presentation of a gray fixation dot for a constant duration of 300 ms signaled the onset of the upcoming sequence of visual objects. All objects in the sequence were presented briefly for a fixed duration of 100 ms. The ISI for each trial was determined based on the current sequence speed (see details above) and was the same for all stimuli within a sequence. The sequence of stimuli was followed by a delay period with a gray fixation dot that was terminated once a fixed duration of 16 s since the onset of the first sequence object had elapsed. This was to ensure sufficient time to acquire the aftereffects of neural responses following the sequence of objects even at a sequence speed of 2048 ms. During this waiting period, participants were listening to bird sounds in order to keep them moderately entertained without additional visual input. Subsequently, the name of the target object as well as the response mapping was presented for a fixed duration of 1.5 s (same fixed response time limit as for the slow trials, see above). In this response interval, participants had to choose the correct serial position of the target object from two response options that were presented on the left and right side of the screen. The mapping of the response options was balanced for left and right responses (i.e., the correct option appeared equally often on the left and right side; 37 times each with the mapping of the last trial being determined randomly) and shuffled randomly for every participant. The serial position of the target for each trial was randomly drawn from a Poisson distribution with  $\lambda = 1.9$  and truncated to an interval from 1 to 5. Thus, across all trials, the targets appeared more often at the later compared to earlier positions of the sequence. This was done to reduce the likelihood that participants stopped to process stimuli or diverted their attention after they identified the position of the target object. The serial position of the alternative response option was drawn from the same distribution as the serial position of the target. As for the slow trials, auditory feedback was presented immediately following a response. The coin sound indicated a reward of 3 cents for correct responses, whereas the buzzer sound signaled incorrect or missed responses (however, there was no deduction of 3 cents for incorrect responses or misses). Together, participants could earn a maximum reward of 2.25 Euro in this task condition.

**Repetition trials.** We included so-called *repetition trials* to investigate how decoding time courses would be affected by (1) the number of fast repetitions of the same neural event and (2) their interaction with the position of the switch to a subsequent stimulus category. Repetition trials included varying repetitions of two images in a sequence of nine items in total. All analyses reported in the Results section focused on the two most extreme cases, (1) the first image shown once followed by eight repetitions of the second image, and (2) eight repetitions of the first image followed by the second image shown once. Analyses of all intermediate levels of repetitions are reported in the SI. Each of the five stimulus categories was selected as the preceding stimulus for eight sequences in total. For each of these eight sequences, we systematically varied the position of the switch to the second stimulus category from serial position 2 to 9. Overall, the transition to the second stimulus happened five times at each serial position with varying stimulus material on each trial. Across the eight trials for each stimulus category, we ensured that each preceding stimulus category was followed by each of the remaining four stimulus categories equally often. Specifically, a given preceding stimulus category was followed by each of the remaining four stimulus categories two times. Also, the average serial position of the first occurrence of each of the subsequent stimuli was the same for all subsequent stimuli. That is to say, the same subsequent stimulus appeared either on position 9 and 2, 8 and 3, 7 and 4, or 6 and 5, resulting in an average first occurrence of the subsequent stimulus at position 5.5. All stimulus sequences of the repetition trials were presented with a fixed ISI of 32 ms. Note that

this is the same presentation speed as the fastest ISI of the sequence trials. Similar to the sequence trials, participants were instructed to remember the serial position at which the second stimulus within the sequence appeared for the first time. For example, if the switch to the second stimulus happened at the fifth serial position, participants had to remember this number.

Similar to the trial procedure of the sequence trials, each repetition trial (Fig. 1c) began with the presentation of the target cue (name of the visual object, e.g., *cat*), which was shown for a fixed duration of 500 ms. The target cue was followed by a blank screen that was presented for a fixed duration of 3.85 s. A briefly presented fixation dot announced the onset of the sequential visual stimuli. Subsequently, the fast sequence of visual stimuli was presented with a fixed duration for visual stimuli (100 ms each) and the ISI (32 ms on all trials). As for sequence trials, the sequence of stimuli on repetition trials was followed by a variable delay period until 16 s from sequence onset had elapsed. On repetition trials, participants had to choose the correct serial position of the first occurrence of the target stimulus from two response options. The incorrect response option was a random serial position that was at least two positions away from the correct target position. For example, if the correct option was 5, the alternative target position could either be earlier (1, 2, or 3) or later (7, 8, or 9). This was done to ensure that the task was reasonably easy to perform. Finally, we added five longer repetition trials with 16 elements per sequence. Here, the switch to the second sequential stimulus always occurred at the last serial position. Each of the five stimulus categories was the preceding stimulus once. The second stimulus of each sequence was any of the other four stimulus categories. In doing so, in the long repetition trials each stimulus category was the preceding and subsequent stimulus once. Repetition trials were randomly distributed across the entire experiment and (together with the sequence trials) interleaved with the slow trials.

**Study procedure.** The study consisted of two experimental sessions. During the first session, participants were informed in detail about the study, screened for MRI eligibility, and provided written informed consent if they agreed to participate in the study. Then they completed a short demographic questionnaire (assessing age, education, etc.) and a computerized version of the Digit-Span Test, assessing working memory capacity<sup>93</sup>. Next, they performed a 10-min practice of the main task. Subsequently, participants entered the MRI scanner. After a short localizer, we first acquired a 5-min resting-state scan for which participants were asked to stay awake and focus on a white fixation cross presented centrally on a black screen. Then, we acquired four functional task runs of about 11 min during which participants performed the main task in the MRI scanner. After the functional runs, we acquired another 5-min resting-state, 5-min fieldmaps, as well as a 4-min anatomical scan. The second study session was identical to the first session, except that participants entered the scanner immediately after another short assessment of MRI eligibility. In total, the study took about 4 h to complete (2.5 and 1.5 h for Session 1 and 2, respectively).

**MRI data acquisition.** All MRI data were acquired using a 32-channel head coil on a research-dedicated 3-Tesla Siemens Magnetom TrioTim MRI scanner (Siemens, Erlangen, Germany) located at the Max Planck Institute for Human Development in Berlin, Germany. The scanning procedure was exactly the same for both study sessions. For the functional scans, whole-brain images were acquired using a segmented k-space and steady-state T2\*-weighted multi-band (MB) echo-planar imaging (EPI) single-echo gradient sequence that is sensitive to the BOLD contrast. This measures local magnetic changes caused by changes in blood oxygenation that accompany neural activity (sequence specification: 64 slices in interleaved ascending order; anterior-to-posterior (A–P) phase-encoding direction; TR = 1250 ms; echo time (TE) = 26 ms; voxel size = 2 × 2 × 2 mm; matrix = 96 × 96; field of view (FOV) = 192 × 192 mm; flip angle (FA) = 71°; distance factor = 0%; MB acceleration factor 4). Slices were tilted for each participant by 15° forwards relative to the rostro-caudal axis to improve the quality of fMRI signal from the hippocampus (cf.<sup>94</sup>) while preserving good coverage of occipito-temporal brain regions. Each MRI session included four functional task runs. Each run was about 11 min in length, during which 530 functional volumes were acquired. For each functional run, the task began after the acquisition of the first four volumes (i.e., after 5 s) to avoid partial saturation effects and allow for scanner equilibrium. We also recorded two functional runs of resting-state fMRI data, one before and one after the task runs. Each resting-state run was about 5 min in length, during which 233 functional volumes were acquired. After the functional task runs, two short acquisitions with six volumes each were collected using the same sequence parameters as for the functional scans but with varying phase-encoding polarities, resulting in pairs of images with distortions going in opposite directions between the two acquisitions (also known as the *blip-up/blip-down* technique). From these pairs the displacement maps were estimated and used to correct for geometric distortions due to susceptibility-induced field inhomogeneities as implemented in the *fMRIPrep* preprocessing pipeline<sup>95</sup>. In addition, a whole-brain spoiled gradient recalled (GR) field map with dual echo-time images (sequence specification: 36 slices; A–P phase-encoding direction; TR = 400 ms; TE1 = 4.92 ms; TE2 = 7.38 ms; FA = 60°; matrix size = 64 × 64; FOV = 192 × 192 mm; voxel size = 3 × 3 × 3.75 mm) was obtained as a potential alternative to the method described above. However, as this field map data were not successfully recorded for four participants, we used the *blip-up/blip-down* technique for distortion correction (see details on MRI data preprocessing

below). Finally, high-resolution T1-weighted (T1w) anatomical Magnetization Prepared Rapid Gradient Echo (MPRAGE) sequences were obtained from each participant to allow registration and brain-surface reconstruction (sequence specification: 256 slices; TR = 1900 ms; TE = 2.52 ms; FA = 9°; inversion time (TI) = 900 ms; matrix size = 192 × 256; FOV = 192 × 256 mm; voxel size = 1 × 1 × 1 mm). We also measured respiration and pulse during each scanning session using pulse oximetry and a pneumatic respiration belt.

**MRI data preparation and preprocessing.** Results included in this manuscript come from preprocessing performed using *fMRIPrep* 1.2.2 (Esteban et al.<sup>95,96</sup>; RRID:SCR\_016216), which is based on *Nipype* 1.1.5 (Gorgolewski et al.<sup>97,98</sup>; RRID:SCR\_002502). Many internal operations of *fMRIPrep* use *Nilearn* 0.4.2<sup>99</sup>; RRID:SCR\_001362, mostly within the functional processing workflow. For more details of the pipeline, see <https://fmripred.readthedocs.io/en/1.2.2/workflows.html> the section corresponding to workflows in *fMRIPrep*'s documentation.

**Conversion of data to the brain imaging data structure (BIDS) standard.** The majority of the steps involved in preparing and preprocessing the MRI data employed recently developed tools and workflows aimed at enhancing standardization and reproducibility of task-based fMRI studies, for a similar preprocessing pipeline, see<sup>100</sup>. Following successful acquisition, all study data were arranged according to the BIDS specification<sup>101</sup> using the *HeuDiConv* tool (version 0.6.0. dev1; freely available from <https://github.com/nipy/heudiconv>) running inside a *Singularity* container<sup>102,103</sup> to facilitate further analysis and sharing of the data. Dicom files were converted to the NIfTI-1 format using *dcm2nii* (version 1.0.20190410 GCC6.3.0)<sup>104</sup>. In order to make identification of study participants unlikely, we eliminated facial features from all high-resolution structural images using *pydeface* (version 2.0; available from <https://github.com/poldracklab/pydeface>). The data quality of all functional and structural acquisitions was evaluated using the automated quality assessment tool *MRIQC* (for details, see<sup>105</sup>, and the [https://mriqc.readthedocs.io/en/stable/MRIQC\\_documentation](https://mriqc.readthedocs.io/en/stable/MRIQC_documentation)). The visual group-level reports of the estimated image quality metrics confirmed that the overall MRI signal quality of both anatomical and functional scans was highly consistent across participants and runs within each participant.

**Preprocessing of anatomical MRI data.** A total of two T1-weighted images were found within the input BIDS dataset, one from each study session. All of them were corrected for intensity non-uniformity (INU) using *N4BiasFieldCorrection* (Advanced Normalization Tools (ANTs) 2.2.0)<sup>106</sup>. A T1w-reference map was created after registration of two T1w images (after INU-correction) using *mri\_robust\_template* (FreeSurfer 6.0.1)<sup>107</sup>. The T1w reference was then skull-stripped using *antsBrainExtraction.sh* (ANTs 2.2.0), using OASIS as target template. Brain surfaces were reconstructed using *recon-all* (FreeSurfer 6.0.1,RRID:SCR\_001847)<sup>108</sup>, and the brain mask estimated previously was refined with a custom variation of the method to reconcile ANTs-derived and FreeSurfer-derived segmentations of the cortical gray-matter of *Mindboggle* (RRID:SCR\_002438)<sup>109</sup>. Spatial normalization to the ICBM 152 Nonlinear Asymmetrical template version 2009c<sup>110</sup> (RRID:SCR\_008796) was performed through nonlinear registration with *antsRegistration* (ANTs 2.2.0,RRID:SCR\_004757)<sup>111</sup>, using brain-extracted versions of both T1w volume and template. Brain tissue segmentation of cerebrospinal fluid (CSF), white-matter (WM), and gray-matter (GM) was performed on the brain-extracted T1w using *fast* (FSL 5.0.9,RRID:SCR\_002823)<sup>112</sup>.

**Preprocessing of functional MRI data.** For each of the BOLD runs found per participant (across all tasks and sessions), the following preprocessing was performed. First, a reference volume and its skull-stripped version were generated using a custom methodology of *fMRIPrep*. The BOLD reference was then co-registered to the T1w reference using *bbregister* (FreeSurfer) which implements boundary-based registration<sup>113</sup>. Co-registration was configured with nine degrees of freedom to account for distortions remaining in the BOLD reference. Head-motion parameters with respect to the BOLD reference (transformation matrices, and six corresponding rotation and translation parameters) are estimated before any spatiotemporal filtering using *mcfliirt* (FSL 5.0.9)<sup>114</sup>. BOLD runs were slice-time-corrected using *3dTshift* from AFNI 20160207<sup>115</sup> (RRID:SCR\_005927). The BOLD time-series (including slice-timing correction) were resampled onto their original, native space by applying a single, composite transform to correct for head-motion and susceptibility distortions. These resampled BOLD time-series will be referred to as preprocessed BOLD in original space, or just preprocessed BOLD. The BOLD time-series were resampled to MNI152NLin2009cAsym standard space, generating a preprocessed BOLD run in MNI152NLin2009cAsym space. First, a reference volume and its skull-stripped version were generated using a custom methodology of *fMRIPrep*. Several confounding time-series were calculated based on the preprocessed BOLD: frame-wise displacement (FD), DVARS, and three region-wise global signals. FD and DVARS are calculated for each functional run, both using their implementations in *Nipype* (following the definitions by Power et al.)<sup>116</sup>. The three global signals are extracted within the CSF, the WM, and the whole-brain masks. Additionally, a set of physiological regressors were extracted to allow for component-based noise correction (*CompCor*)<sup>117</sup>. Principal components are estimated after high-pass filtering the preprocessed BOLD time-series (using a

discrete cosine filter with 128 s cut-off) for the two *CompCor* variants: temporal (*tCompCor*) and anatomical (*aCompCor*). Six *tCompCor* components are then calculated from the top 5% variable voxels within a mask covering the subcortical regions. This subcortical mask is obtained by heavily eroding the brain mask, which ensures it does not include cortical GM regions. For *aCompCor*, six components are calculated within the intersection of the aforementioned mask and the union of CSF and WM masks calculated in T1w space, after their projection to the native space of each functional run (using the inverse BOLD-to-T1w transformation). The head-motion estimates calculated in the correction step were also placed within the corresponding confounds file. The BOLD time-series were resampled to surfaces on the following spaces: *fsnative*, *fsaverage*. All resamplings can be performed with a single interpolation step by composing all the pertinent transformations (i.e., head-motion transform matrices, susceptibility distortion correction when available, and co-registrations to anatomical and template spaces). Gridded (volumetric) resamplings were performed using `antsApplyTransforms` (ANTs), configured with Lanczos interpolation to minimize the smoothing effects of other kernels<sup>118</sup>. Non-gridded (surface) resamplings were performed using `mri_vol2surf` (FreeSurfer). Following preprocessing using `fMRIPrep`, the fMRI data were spatially smoothed using a Gaussian mask with a standard deviation (full-width at half-maximum (FWHM) parameter) set to 4 mm using an example `Nipype` smoothing workflow (see the [Nipype documentation](#) for details) based on the SUSAN algorithm as implemented in the FMRIB Software Library (FSL)<sup>119</sup>.

### Multivariate fMRI pattern analysis

**Leave-one-run-out cross-validation procedure.** All fMRI pattern classification analyses were conducted using open-source packages from the Python (Python Software Foundation, Python Language Reference, version 3.7) modules `Nilearn` (version 0.5.0)<sup>99</sup> and `scikit-learn` (version 0.20.3)<sup>120</sup>. fMRI pattern classification was performed using a leave-one-run-out cross-validation procedure for which data from seven task runs were used for training and data from the left-out run (i.e., the eighth run) were used for testing. This procedure was repeated eight times so that each task run served as the testing set once. We trained an ensemble of five independent classifiers, one for each of the five stimulus classes (cat, chair, face, house, and shoe). For each class-specific classifier, labels of all other classes in the data were relabeled to a common *other* category. In order to ensure that the classifier estimates were not biased by relative differences in class frequency in the training set, the weights associated with each class were adjusted inversely proportional to the class frequencies in each training fold. Given that there were five classes to decode, the frequencies used to adjust the classifiers' weights were 1/5 for the class of interest, and 4/5 for the *other* class, comprising any other classes. Adjustments to minor imbalances caused by the exclusion of erroneous trials were performed in the same way. Training was performed on data from all trials of the seven runs in the respective cross-validation fold using only the trials of the slow task where the visual object stimuli were presented upright and participants did not respond correctly (i.e., correct rejection trials). In each iteration of the classification procedure, the classifiers trained on seven out of eight runs were then applied separately to the data from the left-out run. Specifically, the classifiers were applied to (1) data from the slow trials of the left-out run, selecting volumes capturing the expected activation peaks to determine classification accuracy, (2) data from the slow trials of the left-out run, selecting all volumes from stimulus onset to the end of the trial (seven volumes in total per trial) to identify temporal dynamics of classifier predictions on a single trial basis, (3) data from the sequence trials of the left-out run, selecting all volumes from sequence onset to the end of the delay period (13 volumes in total per trial), and (4) data from the repetition trials of the left-out run, also selecting all volumes from sequence onset to the end of the delay period (13 volumes in total per trial). When the classifiers were applied to sequence and repetition trials, data from both accurate and inaccurate trials were used to allow for an equal number of test trials across participants and maximize statistical power within the current study design. As shown in Fig. 1e, f, behavioral performance on sequence and repetition trials was high and significantly above chance.

We used separate multinomial logistic regression classifiers with identical parameter settings. All classifiers were regularized using L2 regularization. The  $C$  parameter of the cost function was fixed at the default value of 1.0 for all participants. The classifiers employed the `lbfgs` algorithm to solve the multi-class optimization problem and were allowed to take a maximum of 4000 iterations to converge. Pattern classification was performed within each participant separately, never across participants. For each stimulus in the training set, we added 4 s to the stimulus onset and chose the volume closest to that time point (i.e., rounded to the nearest volume) to center the classifier training on the expected peaks of the BOLD response (for a similar approach, see, e.g.<sup>49</sup>). At a TR of 1.25 s, this corresponded to the fourth MRI volume which thus compromised a time window of 3.75–5 s after each stimulus onset. We detrended the fMRI data separately for each run across all task conditions to remove low-frequency signal intensity drifts in the data due to signal noise from the MRI scanner. For each classifier and run, the features were standardized ( $z$ -scored) by removing the mean and scaling to unit variance separately for each test set.

For fMRI pattern classification analysis performed on resting-state data, we created a new mask for each participant through additive combination of the eight masks used for cross-validation (see above). This mask was then applied to all task and resting-state fMRI runs which were then separately detrended and

standardized ( $z$ -scored). The classifiers were trained on the peak activation patterns from all slow trials combined.

**Feature selection.** Feature selection is commonly used in multi-voxel pattern analysis (MVPA) to determine the voxels constituting the activation patterns used for classification in order to improve the predictive performance of the classifier<sup>121,122</sup>. Here, we combined a functional ROI approach based on thresholded  $t$ -maps with anatomical masks to select image-responsive voxels within a predefined anatomical brain region.

We ran eight standard first-level general linear models (GLMs) for each participant, one for each of the eight cross-validation folds using `SPM12` (version 12.7219; <https://www.fil.ion.ucl.ac.uk/spm/software/spm12/>) running inside a Singularity container built using `neurodocker` (version 0.7.0; <https://github.com/ReproNim/neurodocker>) implemented in a custom analysis workflow using `Nipype` (version 1.4.0)<sup>97</sup>. In each cross-validation fold, we fitted a first-level GLM to the data in the training set (e.g., data from run 1 to 7) and modeled the stimulus onset of all trials of the slow task when a stimulus was presented upright and was correctly rejected (i.e., participants did not respond correctly). These trial events were modeled as boxcar functions with the length of the modeling event corresponding to the duration of the stimulus on the screen (500 ms for all events). If present in the training data, we also included trials with hits (correct response to upside-down stimuli), misses (missed response to upside-down stimuli), and false alarms (incorrect response to upright stimuli) as regressors of no interest, thereby explicitly modeling variance attributed to these trial types (cf.<sup>123</sup>). Finally, we included the following nuisance regressors estimated during preprocessing with `fMRIPrep`: the frame-wise displacement for each volume as a quantification of the estimated bulk-head-motion, the six rigid-body motion-correction parameters estimated during realignment (three translation and rotation parameters, respectively), and six noise components calculated according to the anatomical variant of *CompCorr* (for details, see<sup>95</sup>, and the <https://fmripred.readthedocs.io/en/stable/fMRIPrep> documentation). All regressors were convolved with a canonical HRF and did not include model derivatives for time and dispersion. Serial correlations in the fMRI time-series were accounted for using an autoregressive AR (1) model. This procedure resulted in fold-specific maps of  $t$ -values that were used to select voxels from the left-out run of the cross-validation procedure. Note that this approach avoids circularity (or so-called *double-dipping*) as the selective analysis (here, fitting of the GLMs to the training set) is based on data that are fully independent from the data that voxels are later selected from (here, testing set from the left-out run; cf.<sup>124</sup>).

The resulting brain maps of voxel-specific  $t$ -values resulting from the estimation of the described  $t$ -contrast were then combined with an anatomical mask of occipito-temporal brain regions. All participant-specific anatomical masks were created based on automated anatomical labeling of brain-surface reconstructions from the individual T1w-reference image created with `FreeSurfer's recon-all`<sup>108</sup> as part of the `fMRIPrep` workflow<sup>95</sup>, in order to account for individual variability in macroscopic anatomy and to allow reliable labeling<sup>125,126</sup>. For the anatomical masks of occipito-temporal regions we selected the corresponding labels of the cuneus, lateral occipital sulcus, pericalcarine gyrus, superior parietal lobule, lingual gyrus, inferior parietal lobule, fusiform gyrus, inferior temporal gyrus, parahippocampal gyrus, and the middle temporal gyrus (cf.<sup>95</sup>). Only gray-matter voxels were included in the generation of the masks as BOLD signal from non-gray-matter voxels cannot be generally interpreted as neural activity<sup>122</sup>. However, note that due to the whole-brain smoothing performed during preprocessing, voxel activation from brain regions outside the anatomical mask but within the sphere of the smoothing kernel might have entered the anatomical mask (thus, in principle, also including signal from surrounding non-gray-matter voxels).

Finally, we combined the  $t$ -maps derived in each cross-validation fold with the anatomical masks. All voxels with  $t$ -values above or below a threshold of  $t = 3$  (i.e., voxels with the most negative and most positive  $t$ -values) inside the anatomical mask were then selected for the left-out run of the classification analysis and set to 1 to create the final binarized masks ( $M = 11,162$  voxels on average,  $SD = 2,083$ ).

**Classification accuracy and multivariate decoding time courses.** In order to assess the classifiers' ability to differentiate between the neural activation patterns of individual visual objects, we compared the predicted visual object of each example in the test set to the visual object that was actually shown to the participant on the corresponding trial. We obtained an average classification accuracy score for each participant by calculating the mean proportion of correct classifier predictions across all correctly answered, upright slow trials (Fig. 2a). The mean accuracy scores of all participants were then compared to the chance baseline of  $100\%/5 = 20\%$  using a one-sided one-sample  $t$ -test, testing the a priori hypothesis that classification accuracy would be higher than the chance baseline. The effect size (Cohen's  $d$ ) was calculated as the difference between the mean of accuracy scores and the chance baseline, divided by the standard deviation of the data<sup>127</sup>. Furthermore, we assessed the classifiers' ability to accurately detect the presence of visual objects on a single trial basis. For this analysis, we applied the trained classifiers to seven volumes from the volume closest to the stimulus onset and examined the time courses of the probabilistic classification evidence in response to the visual stimuli on a single trial basis (Fig. 2b). In order to test if the time series of classifier probabilities reflected the expected increase of classifier probability for the

stimulus shown on a given trial, we compared the time series of classifier probabilities related to the classified class with the mean time courses of all other classes using a two-sided paired *t*-test at every time point (i.e., at every TR). Here, we used the Bonferroni-correction method<sup>128</sup> across time points and stimulus classes to adjust for multiple comparisons of 35 observations (7 TRs and 5 stimulus classes). In the main text, we only report the results for the peak in classification probability of the true class, corresponding to the fourth TR after stimulus onset. The effect size (Cohen's *d*) was calculated as the difference between the means of the probabilities of the current versus all other stimuli, divided by the standard deviation of the difference<sup>127</sup>.

**Response and difference function modeling.** As reported above, analyzing probabilistic classifier evidence on single slow trials revealed multivariate decoding time courses that can be characterized by a slow response function that resembles single-voxel hemodynamics. For simplicity, we modeled this response function as a sine wave that was flattened after one cycle, scaled by an amplitude, and adjusted to baseline. The model was specified as follows:

$$h(t) = \frac{A}{2} \sin(2\pi ft - 2\pi fd - 0.5\pi) + b + \frac{A}{2} \quad (1)$$

whereby *t* is time, *A* is the response amplitude (the peak deviation of the function from baseline), *f* is the angular frequency (unit: 1/ TR, i.e., 0.8 Hz), *d* is the onset delay (in TRs), and *b* is the baseline (in %). The restriction to one cycle was achieved by converting the sine wave in accordance with the following piece-wise function:

$$H(t) = \begin{cases} h(t) & \text{if } d \leq t \leq (d + \frac{1}{f}) \\ b & \text{otherwise} \end{cases} \quad (2)$$

We fitted the four model parameters (*A*, *f*, *d*, and *b*) to the mean probabilistic classifier evidence of each stimulus class at every TR separately for each participant. For convenience, we count time *t* in TRs. To approximate the time course of the difference between two response functions, we utilized the trigonometric identity for the subtraction of two sine functions, e.g.<sup>129</sup>:

$$\cos(z_1) - \cos(z_2) = -2 \sin\left(\frac{z_1 + z_2}{2}\right) \sin\left(\frac{z_1 - z_2}{2}\right) \quad (3)$$

Considering the case of two sine waves with identical frequency but differing by a temporal shift  $\delta$  one obtains

$$\begin{aligned} \cos(2\pi ft) - \cos(2\pi ft - 2\pi f\delta) &= -2 \sin\left(\frac{4\pi ft - 2\pi f\delta}{2}\right) \sin\left(\frac{2\pi f\delta}{2}\right) \\ &= -2 \sin\left(2\pi f\frac{\delta}{2}\right) \sin\left(2\pi ft - 2\pi f\frac{\delta}{2}\right) \end{aligned} \quad (4)$$

which corresponds to a flipped sine function with an amplitude scaled by  $2 \sin(2\pi f\frac{\delta}{2})$ , a shift of  $\frac{\delta}{2}$  and an identical frequency *f*.

To apply this equation to our scenario, two adjustments have to be made since the single-cycle nature of our response function is not accounted for in Eq. (3). First, one should note that properties of the amplitude term in Eq. (4) only hold as long as shifts of no greater than half a wavelength are considered (the wavelength  $\lambda$  is the inverse of the frequency *f*). The term  $\sin(2\pi f\frac{\delta}{2})$  can be written as  $\sin(2\pi\frac{\delta}{2\lambda})$ , which illustrates that the term monotonically increases until  $\delta > \frac{\lambda}{2}$ . Second, the frequency term has to be adapted as follows: The flattening of the sine waves to the left implies that the difference becomes positive at 0 rather than  $\frac{\delta}{2}$ , thus undoing the phase shift and stretching the wave by  $\frac{1}{2}\delta$  TRs. The flattening on the right also leads to a lengthening of the wave by an additional  $\frac{1}{2}\delta$  TRs, since the difference becomes 0 at  $2\pi f + 2\pi f\delta$ , instead of only  $2\pi f + 2\pi f\frac{\delta}{2}$ . Thus, the total wavelength has to be adjusted by a factor of  $\delta$  TRs, and no phase shift relative to the first response is expected. The difference function therefore has frequency

$$f_\delta = (f^{-1} + \delta)^{-1} = \frac{f}{1 + f\delta} \quad (5)$$

instead of *f*, and Eq. (4) becomes  $-2A \sin(2\pi f_\delta t) \sin(2\pi\frac{f}{1+f\delta}t)$ . We can now apply Eq. (3) to the fitted response function as follows:

$$\begin{aligned} h_\delta(t) &= \left(\frac{1}{2}\hat{A} \cos(2\pi\hat{f}t - 2\pi\hat{f}\hat{d} - 0.5\pi) + \hat{b} + \frac{1}{2}\hat{A}\right) \\ &\quad - \left(\frac{1}{2}\hat{A} \cos(2\pi\hat{f}t - 2\pi\hat{f}\hat{d} - 2\pi\hat{f}\delta - 0.5\pi) + \hat{b} + \frac{1}{2}\hat{A}\right) \\ &= -\hat{A} \sin\left(2\pi\hat{f}\frac{\delta}{2}\right) \sin\left(2\pi\frac{\hat{f}}{1+\hat{f}\delta}t - 2\pi\frac{\hat{f}}{1+\hat{f}\delta}\hat{d} - \pi\right) \\ &= \hat{A} \sin\left(2\pi\hat{f}\frac{\delta}{2}\right) \sin(2\pi\hat{f}_\delta t - 2\pi\hat{f}_\delta\hat{d}) \end{aligned} \quad (6)$$

whereby  $\hat{f}$ ,  $\hat{d}$ ,  $\hat{b}$ , and  $\hat{A}$  indicate fitted parameters.

We determined the relevant TRs in the forward and backward periods for sequence trials by calculating  $\delta$  depending on the sequence speed (the ISI). The resulting values for  $\delta$  and corresponding forward and backward periods are shown

**Table 1 Relevant time periods depending on sequence speed.**

Speed (in ms)	$\delta$ (in TRs)	Forward period	Backward period
32	0.42	TRs 2-4	TRs 5-7
64	0.52	TRs 2-4	TRs 5-7
128	0.73	TRs 2-4	TRs 5-8
512	1.96	TRs 2-5	TRs 6-9
2048	6.87	TRs 2-7	TRs 8-13

Forward periods were calculated as  $[0.56; 0.5 * \lambda_s + d = 0.5 * (5.24 + \delta) + 0.56]$ . Backward period were calculated as  $[0.5 * \lambda_s + d = 0.5 * (5.24 + \delta) + 0.56; \lambda_s + d = 5.24 + \delta + 0.56]$ .  $\delta$  reflects the interval between the onsets of the first and last of five sequence items that is dependent on the sequence speed (the ISI) and the stimulus duration (here, 100 ms). For example, for an ISI of 32 ms,  $\delta$  (in TRs) is calculated as  $(0.032 * 4 + 0.1 * 4) / 1.25 = 0.42$  TRs. *d* reflects the fitted onset delay (here, 0.56 TRs). All values were then rounded to the closest TRs resulting in the speed-adjusted time periods (two rightmost columns).

in Table 1. Model fitting was performed using `NLOptr`, an R interface to the `NLOptr` library for nonlinear optimization<sup>130</sup> employing the COBYLA (Constrained Optimization BY Linear Approximation) algorithm (version 1.2.2.1)<sup>131,132</sup>. The resulting parameters were then averaged across participants, yielding the mean parameters reported in the main text. To assess if the model fitted the data reasonably, we inspected the fits of the sine-wave response function for each stimulus class and participant using individual parameters (Supplementary Fig. 4).

**Detecting sequentiality in fMRI patterns on sequence trials.** In order to analyze the neural activation patterns following the presentation of sequential visual stimuli for evidence of sequentiality, we first determined the true serial position of each decoded event for each trial. Specifically, applying the trained classifiers to each volume of the sequence trials yielded a series of predicted event labels and corresponding classification probabilities that were assigned their sequential position within the true sequence that was shown to participants on the corresponding trial.

The main question we asked for this analysis was to what extent we can infer the serial order of image sequences from relative activation differences in fMRI pattern strength within single measurements (a single TR). To this end, we applied the trained classifiers to a series of 13 volumes following sequence onset (spanning a total time window of about 16 s) on sequence trials and analyzed the time courses of the corresponding classifier probabilities related to the five image categories (Fig. 3a). Classification probabilities were normalized by dividing the probabilities by their trial-wise sum for each image class. As detailed in the task description, the time window was selected such that the neural responses to the image sequences could be fully captured without interference from upcoming trials. We examined relative differences in decoding probabilities between serial events at every time point (i.e., at every TR) and quantified the degree of sequential ordering in two different analyses.

First, we conducted a linear regression between the serial position of the five images and their classification probabilities at every TR in the relevant forward and backward period (adjusted by sequence speed) and extracted the slope of the linear regression as an index of linear association. The slopes were then averaged at every TR separately for each participant and sequence speed across data from all 15 sequence trials (Fig. 3b). Here, if later events have a higher classification probability compared to earlier events, the slope coefficient will be negative. In contrast, if earlier events have a higher classification probability compared to later events, the slope coefficient will be positive. Note that, for convenience, we flipped the sign of the mean regression slopes so that positive values indicate forward ordering and negative values indicate backward ordering. To determine if we can find evidence for significant sequential ordering of classification probabilities in the forward and backward periods, we conducted a series of ten separate two-tailed one-sample *t*-tests comparing the mean regression slope coefficients of each speed condition against zero (the expectation of no order information). All *p* values were adjusted for ten comparisons by controlling the FDR (Fig. 3c;<sup>133</sup>). As an estimate of the effect size, we calculated Cohen's *d* as the difference between the sample mean and the null value in units of the sample standard deviation<sup>127</sup>. As reported in the main text, we conducted the same analysis using rank correlation coefficients (Kendall's  $\tau$ ) and the mean step size between probability-ordered events within TRs as alternative indices of linear association (for details, see SI). In order to directly compare the predicted time courses of regression slopes based on our modeling approach with the observed time courses, we computed the Pearson's correlation coefficient between the two time series, both on data averaged across participants and within each participant (Fig. 2d, e). The mean within-participant correlation coefficients were tested against zero (the expectation of no correlation) using a separate two-sided one-sample *t*-test for each speed condition. All *p* values were adjusted for five comparisons by controlling the FDR<sup>133</sup>.

We hypothesized that sequential order information of fast neural events will translate into order structure in the fMRI signal and successively decoded events in turn. Therefore, we analyzed the fMRI data from sequence trials for evidence of sequentiality across consecutive measurements. The analyses were restricted to the

expected forward and backward periods which were adjusted depending on the sequence speed. For each TR, we obtained the image with the most likely fMRI signal pattern based on the classification probabilities. First, we asked if we are more likely to decode earlier serial events earlier and later serial events later in the decoding time window of 13 TRs. To this end, we averaged the serial position of the most likely event at every TR, separately for each trial and participant, resulting in a time course of average serial event position across the decoding time window (Fig. 3d). We then compared the average serial event position against the mean serial position (position 3) as a baseline across participants at every time point in the forward and backward period using a series of two-sided one-sample *t*-tests, adjusted for 38 multiple comparisons (across all five speed conditions and TRs in the forward and backward period) by controlling the FDR<sup>133</sup>. These results are reported in the SI. Next, in order to assess if the average serial position differed between the forward and backward period for the five different speed conditions, we conducted a linear mixed effects (LME) model and entered the speed condition (with five levels) and trial period (forward versus backward) as fixed effects including by-participant random intercepts and slopes. Finally, we conducted a series of two-sided one-sample *t*-tests to assess whether the mean serial position in the forward and backward periods differed from the expected mean serial position (baseline of 3) for every speed condition (all *p* values adjusted for 10 comparisons using FDR-correction<sup>133</sup>).

Second, we analyzed how this progression through the involved sequence elements affected transitions between consecutively decoded serial events. As before, we extracted the most likely pattern for each TR (i.e., the pattern with the highest classification probability), and calculated the step sizes between consecutively decoded serial events, as in Schuck and Niv<sup>54</sup>. For example, decoding Event 2 → Event 4 in consecutive TRs would correspond to a step size of +2, while a Event 3 → Event 2 transition would reflect a step size of −1, etc. We then calculated the mean step-size of the first (early) and second (late) halves of the forward and backward periods, respectively, which were adjusted for sequence speed. Specifically, the transitions were defined as follows: at speeds of 32, 64, and 128 ms these transitions included the 2 → 3 (early forward), 3 → 4 (late forward), 5 → 6 (early backward), and 6 → 7 (late backward); at speeds of 512 ms these transitions included 2 → 3 (early forward), 4 → 5 (late forward), 6 → 7 (early backward), and 8 → 9 (late backward); at 2048 ms these transitions included 2 → 3 → 4 (early forward), 5 → 6 → 7 (late backward) 8 → 9 → 10 (early backward), and 11 → 12 → 13 (late backward). Finally, we compared the mean step size in the early and late half of the forward versus backward period for every speed condition using ten separate two-sided one-sample *t*-tests. All *p* values were adjusted for multiple comparisons by controlling the FDR (cf.<sup>133</sup>).

**Analysis of repetition trials for sensitivity of within-sequence items.** Applying the classifiers trained on slow trials to data from repetition trials yielded a classification probability estimate for each stimulus class given the data at every time point (i.e., at every TR; Fig. 4a and Supplementary Fig. 9). As described in the main text, we then analyzed the classification probabilities to answer which fMRI patterns were activated during a fast sequence under conditions of extreme forward or backward interference. Specifically, sequences with forward interference entailed a brief presentation of a single image that was followed by eight repetitions of a second image; whereas backward interference was characterized by a condition where eight image repetitions were followed by a single briefly presented item. As predicted by the sine-based response functions, the relevant time period included TRs 2–7. All analyses reported in the Results section were conducted using data from these selected TRs as described. Results based on data from all TRs are reported in the SI.

First, we calculated the mean probability of each event type (first, second, and non-sequence events) across all selected TRs and trials in the relevant time period separately for each repetition condition across participants. In order to examine whether the event type (first, second, and non-sequence events) had an influence on the mean probability estimates on repetition trials, we conducted a LME model<sup>134</sup> and entered the event type (with three factor levels: first, second, and non-sequence events) as a fixed effect and included by-participant random intercepts and slopes (Fig. 4b). Post hoc comparisons between the means of the three factor levels were conducted using Tukey's honest significant difference (HSD) test<sup>135</sup>.

Second, in order to jointly examine the influence of event duration (number of repetitions) and event type (first, second, and non-sequence events), we conducted a LME model<sup>134</sup> with fixed effects of event type (with three factor levels: first, second, and non-sequence events) and repetition condition (number of individual event repetitions with two factor levels: (1) *forward interference* trials, where one briefly presented event is followed by eight repetitions of a second event, and (2) *backward interference* trials, where eight repetitions of a first event are followed by one briefly presented second event), also adding an interaction term for the two effects. Again, the model included both by-participant random intercepts and slopes (Fig. 4c). Post hoc multiple comparisons among interacting factor levels were performed separately for each repetition condition by conditioning on each level of this factor (i.e., forward interference versus backward interference trials), using Tukey's HSD test.

Third, we asked if we are more likely to find transitions between decoded events that were part of the sequence (the two within-sequence items) compared to items that were not part of the sequence (non-sequence items). To this end, we classified each transition as follows: forward (from Event 1 to Event 2), backward (from

Event 2 to Event 1), repetitions of each sequence item, outwards (from sequence items to any non-sequence item), inwards (from non-sequence items to sequence items), outside (among non-sequence items), and repetitions among non-sequence events (the full transition matrix is shown in Fig. 4e). We then compared the average proportion of forward transitions within the sequence (i.e., decoding Event 1 → Event 2) with the average proportions of (1) transitions from sequence items to items that were not part of the sequence (outwards transitions), and (2) transitions between events not part of the sequence (outside transitions) using paired two-sample *t*-tests with *p* values adjusted for four comparisons using Bonferroni correction (Fig. 4d).

**Analysis of sparse sequence events with lower SNR.** We only used resting-state data from the first study session before participants had any experience with the task (except a short training session outside the scanner). These resting-state data could not be successfully recorded in four participants. Therefore, the analyses were restricted to *N* = 32 of 36 participants. Participants were instructed to rest as calmly as possible with eyes opened while focusing on a white fixation cross that was presented centrally on the screen. For decoding on resting-state data, we used the union of all eight masks created for the functional task runs during the cross-validation procedure. Logistic regression classifiers were trained on masked data from slow trials of all eight functional runs and applied to all TRs of the resting-state data, similar to our sequence trial analysis. We assigned pseudo serial positions to each class randomly for every participant, assuming one fixed event ordering. We first characterized and compared the behavior of sequence detection metrics on resting-state and concatenated sequence trial data. For sequence trials, we only considered data from TRs within the expected forward and backward periods (TRs 2–13) and focused on the fastest (32 ms) and slowest (2048 ms) speed condition. Accordingly, we restricted the resting-state data to the first 180 TRs to match it to the length of concatenated sequence trial data (15 concatenated trials of 12 TRs each). For both fast and slow sequence trials and rest data, we then calculated the standard deviation of the probabilities (Fig. 5a) as well as the slope of a linear regression between serial position and their classification probabilities (Fig. 5b, c) at every TR. We then compared both the standard deviation of probabilities and the mean regression slopes over the entire rest period with the mean regression slopes in fast (32 ms) sequence trials using two-sided paired *t*-tests (Fig. 5a, b; *p*s adjusted for four comparisons using Bonferroni correction). The effect sizes (Cohen's *d*) were calculated as the difference between the means of the resting and sequence data, divided by the standard deviation of the differences<sup>127</sup>. Given the rhythmic fluctuations of the regression slope dynamics (Fig. 2e), we calculated the frequency spectra across the resting-state and concatenated sequence trial data using the Lomb-Scargle method (using the `lsp` function from the R package `lomb`, e.g.<sup>61</sup> that is suitable for unevenly sampled data, and therefore accounts for potential artifacts due to data concatenation (Fig. 5d). The resulting frequency spectra were smoothed with a running average filter with width 0.005. Next, we extracted the mean power of the frequencies for fast and slow event sequences as predicted by Eq. (5) in both resting and sequence data. For example, for a 32 ms sequence with  $\delta = 0.032 * 4 + 0.1 * 5 = 0.628$ , one obtains the predicted frequency as  $f_{\delta} = \frac{f}{1+f*0.628} = 0.17$ , whereby *f* equals the fitted single trial frequency  $f = 1/5.24$ . The mean power at the predicted frequencies were then compared between resting as well as fast and slow sequence data using two-sided paired *t*-tests with *p* values adjusted for multiple comparisons using FDR-correction<sup>133</sup>.

We then inserted 1–6 sequence events into the pre-task resting-state period by blending TRs during resting state with TRs recorded during fast (32 ms) or slow (2048 ms) sequence trials. Specifically, we randomly selected six sequence trials for each speed condition, without replacement. Only TRs from the relevant time period (see above; 12 TRs for both speed conditions, respectively) were blended into the resting-state data. To investigate the effects of a reduced SNR, we systematically multiplied the probabilities of the inserted sequence TRs by a factor  $\kappa$  of  $\frac{4}{5}$ ,  $\frac{1}{2}$ ,  $\frac{1}{4}$ ,  $\frac{1}{8}$ , or 0, step-wise reducing the signal from 80% to 0% and added these scaled probabilities to the probability time courses of the resting-state data. The resting-state data used for blending were independently sampled from non-overlapping random locations within the resting-state data of the same participant. This ensured that even in the 0 SNR condition, potential artifacts due to data concatenation were present and would therefore not impact our comparisons between SNR levels. For each combination of the number of inserts and SNR levels, we then compared the mean standard deviation of the probabilities during sequence-inserted rest with sequence-free rest using a series of two-sided paired *t*-tests. *p* values were adjusted accordingly for 30 comparisons using FDR-correction<sup>133</sup> and log-transformed (base 20) to make them easier to visualize (here, a log-transformed *p* value of 1 corresponds to  $p < 0.05$ ).

Finally, we calculated the frequency spectra of sequence-inserted rest data as before, separately for data with fast and slow sequence inserts. To achieve comparable resolution obtained in the above analyses, we over-sampled the frequency space by a factor of 2. Smoothing was then applied again as before. We then calculated the relative power of each frequency compared to sequence-free rest and averaged the relative frequency spectra across participants (Fig. 5h). As before, we extracted the mean power within the predicted fast and slow frequency range ( $\pm 0.01$  Hz, given the smoothing) and compared them between fast and slow sequence-inserted rest and for different numbers of inserts and SNR levels. We



then compared the relative power for each sequence-inserted rest dataset, number of inserts, and SNR level against zero (no difference from sequence-free rest) using a series of two-sided one-sample *t*-tests (*p* values uncorrected).

**Analysis of task-related reactivations in post-task resting-state data.** We investigated whether the frequency spectrum analyses described above could be used to detect task-related reactivations of stimulus sequences in post-task resting-state data in Session 1 (i.e., after participants performed four runs of the task). As the pre-task resting-state acquisition, post-task resting-state data consisted of a 5-min fMRI run during which participants rested calmly with eyes open but without any additional task. We calculated the frequency spectra (using the Lomb-Scargle method) across the pre- and post-task resting-state data as described above (see Fig. 6a). To this end, we calculated the slopes and frequency spectra in the two resting-state runs considering all permutations of possible sequential orderings of classification probabilities at every TR, rather than assuming a random ordering (as for the sequence-inserted rest analyses described above), then averaging across all data from all permutations. We then compared pre- and post-task rest directly by calculating the relative power of the frequency spectra as the difference between pre- and post-task rest (Fig. 6b). Finally, we assessed if the power difference in the fast frequency range (0.17 Hz), indicative of fast sequential neural events, between pre- and post-task rest was specific to the sequential combinations of stimuli that participants experienced during the task. To this end, we split the data depending on whether they were created based on sequences the participants experienced more or less frequently during the task. As described above, the 15 sequences that were selected for the sequence trials for each participant were considered more frequent compared to all other sequential permutations that participants experienced during the slow trials. Lastly, to examine if the increases in power in the fast frequency range were specific to the more frequent sequences, we conducted a LME model with the resting-state run (pre- vs. post-task) and the sequence frequency (less vs. more frequent) as the main fixed effects of interest, and by-participant random intercepts and slopes (Fig. 6c). Post hoc multiple comparisons among the interacting factors were performed using Tukey's HSD test.

**Statistical analysis.** Main statistical analyses were conducted using LME models employing the `lmer` function of the `lme4` package (version 1.1.21<sup>134</sup>) in R (version 3.6.1<sup>136</sup>). If not stated otherwise, all models were fit with participants considered as a random effect on both the intercept and slopes of the fixed effects, in accordance with results from Barr et al.<sup>137</sup> who recommend to fit the most complex model consistent with the experimental design<sup>137</sup>. If applicable, explanatory variables were standardized to a mean of 0 and a standard deviation of 1 before they entered the models. If necessary, we removed by-participant slopes from the random effects structure to achieve a non-singular fit of the model<sup>137</sup>. Models were fitted using the BOBYQA (Bound Optimization BY Quadratic Approximation) optimizer<sup>138,139</sup> with a maximum of 500,000 function evaluations and no calculation of gradient and Hessian of nonlinear optimization solution. The likelihoods of the fitted models were assessed using Type III analysis of variance (ANOVA) with Satterthwaite's method. A single-step multiple comparison procedure between the means of the relevant factor levels was conducted using Tukey's HSD test<sup>135</sup>, as implemented in the `emmeans` package in R (version 1.3.4<sup>136,140</sup>). In all other analyses, we used one-sample *t*-tests if group data were compared to, e.g., a baseline, or paired *t*-tests if two samples from the same population were compared. If applicable, correction for multiple hypothesis testing was performed using the FDR-correction method<sup>133</sup>. If not stated otherwise, *t*-tests were two-sided and the  $\alpha$  level set to 0.05.

**Analysis of behavioral data.** The main goal of the current study was to investigate the statistical properties of BOLD activation patterns following the presentation of fast visual object sequences. Therefore, attentive processing of all visual stimuli was a prerequisite to ensure that we would be able to decode neural representations of the stimuli from occipito-temporal fMRI data. If behavioral performance was low, we could expect that participants did not attend well to the stimuli. We thus calculated the mean behavioral accuracy on sequence and repetition trials and excluded all participants that had a mean behavioral accuracy below the 50% chance level (Supplementary Fig. 1a). Mean behavioral accuracy scores of the remaining participants in the final sample are reported in the main text (Fig. 1e, f) and the SI (Supplementary Fig. 1). In order to assess how well participants detected upside-down stimuli on slow trials, we conducted a one-sided one-sample *t*-test against the 50% chance level, testing the a priori hypothesis that mean behavioral accuracy would be higher than chance. Cohen's *d* quantified the effect size and was calculated as the difference between the mean of the data and the chance level, divided by the standard deviation of the data<sup>127</sup>. As low performance in this task condition could be indicated by both false alarms (incorrect response to upright stimuli) and misses (missed response to upside-down stimuli), we also checked whether the frequency of false alarms and misses differed (Supplementary Fig. 1b). Furthermore, we assessed if behavioral accuracy on slow trials used for classifier training was stable across task runs (Supplementary Fig. 1c). In order to examine the effect of sequence speed on behavioral accuracy in sequence trials, we conducted a LME model including the sequence speed condition as the main fixed effect of interest, and by-participant random intercepts and slopes (Fig. 1e). We then examined whether performance was above chance for all five speed conditions

and conducted five separate one-sided one-sample *t*-tests testing the a priori hypothesis that mean behavioral accuracy would be higher than a 50% chance level. All *p* values were adjusted for multiple comparisons using FDR-correction<sup>133</sup>. The effect of serial position of the cued target image on behavioral accuracy is reported in the SI (Supplementary Fig. 1d). For repetition trials with forward and backward interference we conducted separate one-sided one-sample *t*-test for each repetition condition to test the a priori hypothesis that behavioral accuracy would be higher than the 50% chance level (Fig. 1f). Results for all repetition conditions are reported in the SI (Supplementary Fig. 1e). The effect sizes (Cohen's *d*) were calculated as for slow trials.

**Reporting summary.** Further information on research design is available in the Nature Research Reporting Summary linked to this article.

## Data availability

We publicly share all data used in this study. Data and code management was realized using DataLad [version 0.13.0<sup>142</sup>, for details, see <https://www.datalad.org/>]. An overview of all the resources is publicly available on our project website: <https://wittkuhn.mpib.berlin/highspeed/>. All individual datasets can be found at <https://gin.g-node.org/lnnrtwttkhn>. Please note that each dataset is associated with a unique URL and Digital Object Identifier (DOI). We share all MRI and behavioral data adhering to the BIDS standard (cf.<sup>101</sup>) (<https://github.com/lnnrtwttkhn/highspeed-bids>; <https://gin.g-node.org/lnnrtwttkhn/highspeed-bids>; <https://doi.org/10.12751/g-node.4ivuv8>), all MRI quality metrics and reports based on MRIQC (cf.<sup>105</sup>) (<https://github.com/lnnrtwttkhn/highspeed-mriqc>; <https://gin.g-node.org/lnnrtwttkhn/highspeed-mriqc>; <https://doi.org/10.12751/g-node.0vmyuh>), all preprocessed MRI data using fMRIprep (cf.<sup>96,143</sup>) (<https://github.com/lnnrtwttkhn/highspeed-fmriprep>; <https://gin.g-node.org/lnnrtwttkhn/highspeed-fmriprep>; <https://doi.org/10.12751/g-node.0fi06t>), all binarized anatomical masks used for feature selection (<https://github.com/lnnrtwttkhn/highspeed-masks>; <https://doi.org/10.12751/g-node.omirok>), all first-level GLM results used for feature selection (<https://github.com/lnnrtwttkhn/highspeed-glm>; <https://gin.g-node.org/lnnrtwttkhn/highspeed-glm>; <https://doi.org/10.12751/g-node.d21zpv>), all results of the multivariate decoding approach (<https://github.com/lnnrtwttkhn/highspeed-decoding>; <https://gin.g-node.org/lnnrtwttkhn/highspeed-decoding>; <https://doi.org/10.12751/g-node.9zfi1r>), and the unprocessed data of the behavioral task acquired during MRI acquisition (<https://github.com/lnnrtwttkhn/highspeed-data-behavior>; <https://gin.g-node.org/lnnrtwttkhn/highspeed-data-behavior>; <https://doi.org/10.12751/g-node.p7dabb>). Bird sounds used as stimuli can be downloaded from <https://audiojungle.net/item/british-bird-song-dawn-chorus/98074>. The visual stimulus material is freely available from <http://data.pymvpa.org/datasets/haxby2001/>. The original authors of<sup>55</sup> hold the copyright of this dataset and made it available under the terms of the Creative Commons Attribution-Share Alike 3.0 license (see <http://creativecommons.org/licenses/by-sa/3.0/>) for details). The images selected for the task were not modified. Source Data to reproduce the main parts of all figures are provided with this paper.

## Code availability

We share all code used in this study. An overview of all the resources is publicly available on our project website: <https://wittkuhn.mpib.berlin/highspeed/>. All code for the main statistical analyses can be found at <https://github.com/lnnrtwttkhn/highspeed-analysis>; <https://gin.g-node.org/lnnrtwttkhn/highspeed-analysis>; <https://doi.org/10.12751/g-node.eqqdqt>). All code to run the behavioral task can be found at (<https://github.com/lnnrtwttkhn/highspeed-task>; <https://doi.org/10.5281/zenodo.4305888>). Please note that we share all data listed in the Data availability section in modularized units alongside the code that created the data, usually in a dedicated `code` directory in each dataset, instead of separate data and code repositories. This approach allows to better establish the provenance of data (i.e., a better understanding which code and input data produced which output data), loosely following the DataLad YODA principles (for details, see the chapter "YODA: Best practices for data analyses in a dataset" in the DataLad handbook (version 0.13<sup>144</sup>), available at <https://handbook.datalad.org/>).

Received: 25 March 2020; Accepted: 16 February 2021;

Published online: 19 March 2021

## References

- Wilson, M. & McNaughton, B. L. Reactivation of hippocampal ensemble memories during sleep. *Science* **265**, 676–679 (1994).
- Skaggs, W. E. & McNaughton, B. L. Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* **271**, 1870–1873 (1996).
- Louie, K. & Wilson, M. A. Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* **29**, 145–156 (2001).

4. Lee, A. K. & Wilson, M. A. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* **36**, 1183–1194 (2002).
5. Diba, K. & Buzsáki, G. Forward and reverse hippocampal place-cell sequences during ripples. *Nat. Neurosci.* **10**, 1241–1242 (2007).
6. Foster, D. J. & Wilson, M. A. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**, 680–683 (2006).
7. Davidson, T. J., Kloosterman, F. & Wilson, M. A. Hippocampal replay of extended experience. *Neuron* **63**, 497–507 (2009).
8. Karlsson, M. P. & Frank, L. M. Awake replay of remote experiences in the hippocampus. *Nat. Neurosci.* **12**, 913–918 (2009).
9. Gupta, A. S., van der Meer, M. A., Touretzky, D. S. & Redish, A. D. Hippocampal replay is not a simple function of experience. *Neuron* **65**, 695–705 (2010).
10. Johnson, A. & Redish, A. D. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.* **27**, 12176–12189 (2007).
11. Kay, K. et al. Constant sub-second/ cycling between representations of possible futures in the hippocampus. *Cell* **180**, 552–567.e25 (2020).
12. Ji, D. & Wilson, M. A. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat. Neurosci.* **10**, 100–107 (2006).
13. Euston, D. R., Tatsuno, M. & McNaughton, B. L. Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science* **318**, 1147–1150 (2007).
14. Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S. I. & Battaglia, F. P. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nat. Neurosci.* **12**, 919–926 (2009).
15. Rothschild, G., Eban, E. & Frank, L. M. A cortical–hippocampal–cortical loop of information processing during memory consolidation. *Nat. Neurosci.* **20**, 251–259 (2016).
16. Shin, J. D., Tang, W. & Jadhav, S. P. Dynamics of awake hippocampal–prefrontal replay for spatial learning and memory-guided decision making. *Neuron* **104**, 1110–1125 (2019).
17. Kaefer, K., Nardin, M., Blahna, K. & Csicsvari, J. Replay of behavioral sequences in the medial prefrontal cortex during rule switching. *Neuron* **106**, 154–165.e6 (2020).
18. Rich, E. L. & Wallis, J. D. Decoding subjective decisions from orbitofrontal cortex. *Nat. Neurosci.* **19**, 973–980 (2016).
19. Egleman, S. L. & Dragoi, V. Image sequence reactivation in awake V4 networks. *Proc. Natl Acad. Sci. USA* **109**, 19450–19455 (2012).
20. Ekman, M., Kok, P. & de Lange, F. P. Time-compressed preplay of anticipated events in human primary visual cortex. *Nat. Commun.* **8**, 1–9 (2017).
21. Xu, S., Jiang, W., Poo, M.-m. & Dan, Y. Activity recall in a visual cortical ensemble. *Nat. Neurosci.* **15**, 449–455 (2012).
22. Ólafsdóttir, H. F., Carpenter, F. & Barry, C. Coordinated grid and place cell replay during rest. *Nat. Neurosci.* **19**, 792–794 (2016).
23. Ghuman, A. S. & Martin, A. Dynamic neural representations: an inferential challenge for fMRI. *Trends Cogn. Sci.* **23**, 534–536 (2019).
24. Tambini, A. & Davachi, L. Awake reactivation of prior experiences consolidates memories and biases cognition. *Trends Cogn. Sci.* **23**, 876–890 (2019).
25. Heeger, D. J. & Ress, D. What does fMRI tell us about neuronal activity? *Nat. Rev. Neurosci.* **3**, 142–151 (2002).
26. Ogawa, S., Lee, T. M., Kay, A. R. & Tank, D. W. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc. Natl Acad. Sci. USA* **87**, 9868–9872 (1990).
27. Kwong, K. K. et al. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc. Natl Acad. Sci. USA* **89**, 5675–5679 (1992).
28. Miezin, F., Maccotta, L., Ollinger, J., Petersen, S. & Buckner, R. Characterizing the hemodynamic response: effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. *NeuroImage* **11**, 735–759 (2000).
29. Kim, S.-G., Richter, W. & Uğurbil, K. Limitations of temporal resolution in functional MRI. *Magn. Reson. Med.* **37**, 631–636 (1997).
30. Aguirre, G., Zarahn, E. & D’Esposito, M. The variability of human, BOLD hemodynamic responses. *NeuroImage* **8**, 360–369 (1998).
31. Menon, R. S., Luknowsky, D. C. & Gati, J. S. Mental chronometry using latency-resolved functional MRI. *Proc. Natl Acad. Sci. USA* **95**, 10902–10907 (1998).
32. Mísaki, M., Luh, W.-M. & Bandettini, P. A. Accurate decoding of sub-TR timing differences in stimulations of sub-voxel regions from multi-voxel response patterns. *NeuroImage* **66**, 623–633 (2013).
33. Foster, D. J. Replay comes of age. *Annu. Rev. Neurosci.* **40**, 581–602 (2017).
34. Ólafsdóttir, H. F., Bush, D. & Barry, C. The role of hippocampal replay in memory and planning. *Curr. Biol.* **28**, R37–R50 (2018).
35. Pezzulo, G., Donnarumma, F., Maisto, D. & Stoianov, I. Planning at decision time and in the background during spatial navigation. *Curr. Opin. Behav. Sci.* **29**, 69–76 (2019).
36. Staresina, B. P. & Wimber, M. A neural chronometry of memory recall. *Trends Cogn. Sci.* **23**, 1071–1085 (2019).
37. Redish, A. D. Beyond replay: introduction to the special issue on hippocampal replay. *Hippocampus* **30**, 3–5 (2020).
38. Axmacher, N., Elger, C. E. & Fell, J. Ripples in the medial temporal lobe are relevant for human memory consolidation. *Brain* **131**, 1806–1817 (2008).
39. Logothetis, N. K. et al. Hippocampal–cortical interaction during periods of subcortical silence. *Nature* **491**, 547–553 (2012).
40. Staresina, B. P. et al. Hierarchical nesting of slow oscillations, spindles and ripples in the human hippocampus during sleep. *Nat. Neurosci.* **18**, 1679–1686 (2015).
41. Zhang, H., Fell, J. & Axmacher, N. Electrophysiological mechanisms of human memory consolidation. *Nat. Commun.* **9**, 4103 (2018).
42. Vaz, A. P., Inati, S. K., Brunel, N. & Zaghoul, K. A. Coupled ripple oscillations between the medial temporal lobe and neocortex retrieve human memory. *Science* **363**, 975–978 (2019).
43. Jafarpour, A., Fuentemilla, L., Horner, A. J., Penny, W. & Düzel, E. Replay of very early encoding representations during recollection. *J. Neurosci.* **34**, 242–248 (2013).
44. Kurth-Nelson, Z., Economides, M., Dolan, R. J. & Dayan, P. Fast sequences of non-spatial state representations in humans. *Neuron* **91**, 194–204 (2016).
45. Michelmann, S., Staresina, B. P., Bowman, H. & Hanslmayr, S. Speed of time-compressed forward replay flexibly changes in human episodic memory. *Nat. Hum. Behav.* **3**, 143–154 (2018).
46. Huang, Q., Jia, J., Han, Q. & Luo, H. Fast-backward replay of sequentially memorized items in humans. *eLife* **7**, e35164 (2018).
47. Liu, Y., Dolan, R. J., Kurth-Nelson, Z. & Behrens, T. E. Human replay spontaneously reorganizes experience. *Cell* **178**, 640–652 (2019).
48. Wimmer, G. E., Liu, Y., Vehar, N., Behrens, T. E. & Dolan, R. J. Episodic memory retrieval is associated with rapid replay of episode content. *Nat. Neurosci.* **23**, 1025–1033 (2020).
49. Deuker, L. et al. Memory consolidation by replay of stimulus-specific neural activity. *J. Neurosci.* **33**, 19373–19383 (2013).
50. Staresina, B. P., Alink, A., Kriegeskorte, N. & Henson, R. N. Awake reactivation predicts memory in humans. *Proc. Natl Acad. Sci. USA* **110**, 21159–21164 (2013).
51. Tambini, A. & Davachi, L. Persistence of hippocampal multivoxel patterns into postencoding rest is related to memory. *Proc. Natl Acad. Sci. USA* **110**, 19591–19596 (2013).
52. Momennejad, I., Otto, A. R., Daw, N. D. & Norman, K. A. Offline replay supports planning in human reinforcement learning. *eLife* **7**, e32548 (2018).
53. Schapiro, A. C., McDevitt, E. A., Rogers, T. T., Mednick, S. & Norman, K. A. Human hippocampal replay during rest prioritizes weakly learned information and predicts memory performance. *Nat. Commun.* **9**, 3920 (2018).
54. Schuck, N. W. & Niv, Y. Sequential replay of nonspatial task states in the human hippocampus. *Science* **364**, eaaw5181 (2019).
55. Haxby, J. V. et al. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* **293**, 2425–2430 (2001).
56. Dale, A. M. Optimal experimental design for event-related fMRI. *Hum. Brain Mapp.* **8**, 109–114 (1999).
57. Cai, M. B., Schuck, N. W., Pillow, J. W. & Niv, Y. Representational structure or task structure? Bias in neural representational similarity analysis and a Bayesian method for reducing bias. *PLoS Comput. Biol.* **15**, e1006299 (2019).
58. Spiridon, M. & Kanwisher, N. How distributed is visual category information in human occipito-temporal cortex? An fMRI study. *Neuron* **35**, 1157–1165 (2002).
59. Hanson, S. J., Matsuka, T. & Haxby, J. V. Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: is there a “face” area? *NeuroImage* **23**, 156–166 (2004).
60. O’Toole, A. J., Jiang, F., Abdi, H. & Haxby, J. V. Partially distributed representations of objects and faces in ventral temporal cortex. *J. Cogn. Neurosci.* **17**, 580–590 (2005).
61. VanderPlas, J. T. Understanding the Lomb–Scargle periodogram. *Astrophys. J. Suppl. Ser.* **236**, 16 (2018).
62. Carr, M. F., Jadhav, S. P. & Frank, L. M. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nat. Neurosci.* **14**, 147–153 (2011).
63. Schmidt, B., Wikenheiser, A. M. & Redish, A. D. Goal-directed sequences in the hippocampus. In *Goal-Directed Decision Making: Computations and Neural Circuits* 1st edn. (eds. Morris, R., Bornstein, A. & Shenhav, A.) Ch. 6, 125–151 (Elsevier, 2018).
64. Zhang, H., Deuker, L. & Axmacher, N. Replay in humans - first evidence and open questions. In *Cognitive Neuroscience of Memory Consolidation* (eds. Axmacher, N. & Rasch, B.) 251–263 (Springer, 2017).
65. Schlichting, M. L. & Preston, A. R. Memory reactivation during rest supports upcoming learning of related content. *Proc. Natl Acad. Sci. USA* **111**, 15845–15850 (2014).

66. Gruber, M. J., Ritchey, M., Wang, S.-F., Doss, M. K. & Ranganath, C. Post-learning hippocampal dynamics promote preferential retention of rewarding events. *Neuron* **89**, 1110–1120 (2016).
67. Hermans, E. J. et al. Persistence of amygdala–hippocampal connectivity and multi-voxel correlation structures during awake rest after fear learning predicts long-term expression of fear. *Cereb. Cortex* **27**, 3028–3041 (2017).
68. de Voogd, L. D., Fernández, G. & Hermans, E. J. Awake reactivation of emotional memory traces through hippocampal–neocortical interactions. *NeuroImage* **134**, 563–572 (2016).
69. Sadeh, T., Chen, J., Goshen-Gottstein, Y. & Moscovitch, M. Overlap between hippocampal pre-encoding and encoding patterns supports episodic memory. *Hippocampus* **29**, 836–847 (2019).
70. Tambini, A., Ketz, N. & Davachi, L. Enhanced brain correlations during rest are related to memory for recent experiences. *Neuron* **65**, 280–290 (2010).
71. Tompary, A., Duncan, K. & Davachi, L. Consolidation of associative and item memory is related to post-encoding functional connectivity between the ventral tegmental area and different medial temporal lobe subregions during an unrelated task. *J. Neurosci.* **35**, 7326–7331 (2015).
72. Murty, V. P., Tompary, A., Adcock, R. A. & Davachi, L. Selectivity in postencoding connectivity with high-level visual cortex is associated with reward-motivated memory. *J. Neurosci.* **37**, 537–545 (2017).
73. Retter, T. L., Jiang, F., Webster, M. A. & Rossion, B. Dissociable effects of inter-stimulus interval and presentation duration on rapid face categorization. *Vision Res.* **145**, 11–20 (2018).
74. Robinson, A. K., Grootswagers, T. & Carlson, T. A. The influence of image masking on object representations during rapid serial visual presentation. *NeuroImage* **197**, 224–231 (2019).
75. Wilson, R. I. & Nicoll, R. A. Endogenous cannabinoids mediate retrograde signalling at hippocampal synapses. *Nature* **410**, 588–592 (2001).
76. Pfeiffer, B. E. & Foster, D. J. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**, 74–79 (2013).
77. Yaffe, R. B., Shaikhouni, A., Arai, J., Inati, S. K. & Zaghoul, K. A. Cued memory retrieval exhibits reinstatement of high gamma power on a faster timescale in the left temporal lobe and prefrontal cortex. *J. Neurosci.* **37**, 4472–4480 (2017).
78. Wimber, M., Maaß, A., Staudigl, T., Richardson-Klavehn, A. & Hanslmayr, S. Rapid memory reactivation revealed by oscillatory entrainment. *Curr. Biol.* **22**, 1482–1486 (2012).
79. Zhang, H. et al. Gamma power reductions accompany stimulus-specific representations of dynamic events. *Curr. Biol.* **25**, 635–640 (2015).
80. Staudigl, T., Vollmar, C., Noachtar, S. & Hanslmayr, S. Temporal-pattern similarity analysis reveals the beneficial and detrimental effects of context reinstatement on human memory. *J. Neurosci.* **35**, 5373–5384 (2015).
81. Jadhav, S. P., Rothschild, G., Roumis, D. K. & Frank, L. M. Coordinated excitation and inhibition of prefrontal ensembles during awake hippocampal sharp-wave ripple events. *Neuron* **90**, 113–127 (2016).
82. O’Neill, J., Boccarda, C., Stella, F., Schoenenberger, P. & Csicsvari, J. Superficial layers of the medial entorhinal cortex replay independently of the hippocampus. *Science* **355**, 184–188 (2017).
83. Trettel, S. G., Trimper, J. B., Hwaun, E., Fiete, I. R. & Colgin, L. L. Grid cell co-activity patterns during sleep reflect spatial overlap of grid fields during active behaviors. *Nat. Neurosci.* **22**, 609–617 (2019).
84. Lansink, C. S., Goltstein, P. M., Lankelma, J. V., McNaughton, B. L. & Pennartz, C. M. A. Hippocampus leads ventral striatum in replay of place-reward information. *PLoS Biol.* **7**, e1000173 (2009).
85. Leontiev, O. & Buxton, R. B. Reproducibility of BOLD, perfusion, and CMRO<sub>2</sub> measurements with calibrated-BOLD fMRI. *NeuroImage* **35**, 175–184 (2007).
86. de Zwart, J. A. et al. Hemodynamic nonlinearities affect BOLD fMRI response timing and amplitude. *NeuroImage* **47**, 1649–1658 (2009).
87. Lin, F.-H. et al. Relative latency and temporal variability of hemodynamic responses at the human primary visual cortex. *NeuroImage* **164**, 194–201 (2018).
88. Menon, R. S. & Kim, S.-G. Spatial and temporal limits in cognitive neuroimaging with fMRI. *Trends Cogn. Sci.* **3**, 207–216 (1999).
89. Lin, F.-H. et al. fMRI hemodynamics accurately reflects neuronal timing in the human brain measured by MEG. *NeuroImage* **78**, 372–384 (2013).
90. Brainard, D. H. The psychophysics toolbox. *Spat. Vis.* **10**, 433–436 (1997).
91. Kleiner, M. et al. What’s new in Psychtoolbox-3? A free cross-platform toolkit for psychophysics with Matlab and GNU/Octave. *Cogn. Comput. Psychophys.* **36**, 1–89 (2007).
92. Pelli, D. G. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* **10**, 437–442 (1997).
93. Petermann, F. & Wechsler, D. *Wechsler Adult Intelligence Scale* 4th edn. (Pearson, 2008).
94. Weiskopf, N., Hutton, C., Josephs, O. & Deichmann, R. Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: a whole-brain analysis at 3 T and 1.5 T. *NeuroImage* **33**, 493–504 (2006).
95. Esteban, O. et al. fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nat. Methods* **16**, 111–116 (2018).
96. Esteban, O. et al. fMRIPrep 1.2.2. (fMRIPrep, 2019).
97. Gorgolewski, K. J. et al. Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in Python. *Front. Neuroinform.* **5**, 13 (2011).
98. Gorgolewski, K. J. et al. Nipype (Nipype, 2019).
99. Abraham, A. et al. Machine learning for neuroimaging with scikit-learn. *Front. Neuroinform.* **8**, 14 (2014).
100. Esteban, O. et al. Analysis of task-based functional MRI data preprocessed with fMRIPrep. *Nat. Protoc.* **15**, 2186–2202 (2020).
101. Gorgolewski, K. J. et al. The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Sci. Data* **3**, 160044 (2016).
102. Kurtzer, G. M., Sochat, V. & Bauer, M. W. Singularity: scientific containers for mobility of compute. *PLoS ONE* **12**, e0177459 (2017).
103. Sochat, V. V., Prybol, C. J. & Kurtzer, G. M. Enhancing reproducibility in scientific computing: metrics and registry for Singularity containers. *PLoS ONE* **12**, e0188511 (2017).
104. Li, X., Morgan, P. S., Ashburner, J., Smith, J. & Rorden, C. The first step for neuroimaging data analysis: dicom to nifti conversion. *J. Neurosci. Methods* **264**, 47–56 (2016).
105. Esteban, O. et al. MRIQC: advancing the automatic prediction of image quality in MRI from unseen sites. *PLoS ONE* **12**, e0184661 (2017).
106. Tustison, N. J. et al. N4ITK: improved n3 bias correction. *IEEE Trans. Med. Imaging* **29**, 1310–1320 (2010).
107. Reuter, M., Rosas, H. D. & Fischl, B. Highly accurate inverse consistent registration: a robust approach. *NeuroImage* **53**, 1181–1196 (2010).
108. Dale, A. M., Fischl, B. & Sereno, M. I. Cortical surface-based analysis. *NeuroImage* **9**, 179–194 (1999).
109. Klein, A. et al. Mindboggling morphometry of human brains. *PLoS Comput. Biol.* **13**, e1005350 (2017).
110. Fonov, V., Evans, A., McKinstry, R., Almlí, C. & Collins, D. Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage* **47**, S102 (2009).
111. Avants, B., Epstein, C., Grossman, M. & Gee, J. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Anal.* **12**, 26–41 (2008).
112. Zhang, Y., Brady, M. & Smith, S. Segmentation of brain MR images through a hidden markov random field model and the expectation-maximization algorithm. *IEEE Trans. Med. Imaging* **20**, 45–57 (2001).
113. Greve, D. N. & Fischl, B. Accurate and robust brain image alignment using boundary-based registration. *NeuroImage* **48**, 63–72 (2009).
114. Jenkinson, M., Bannister, P., Brady, M. & Smith, S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage* **17**, 825–841 (2002).
115. Cox, R. W. & Hyde, J. S. Software tools for analysis and visualization of fMRI data. *NMR Biomed.* **10**, 171–178 (1997).
116. Power, J. D. et al. Methods to detect, characterize, and remove motion artifact in resting state fMRI. *NeuroImage* **84**, 320–341 (2014).
117. Behzadi, Y., Restom, K., Liu, J. & Liu, T. T. A component based noise correction method (compcor) for bold and perfusion based fMRI. *NeuroImage* **37**, 90–101 (2007).
118. Lanczos, C. Evaluation of noisy data. *J. Soc. Indust. Appl. Math.* **B 1**, 76–85 (1964).
119. Smith, S. M. & Brady, J. M. SUSAN - a new approach to low level image processing. *Int. J. Comput. Vis.* **23**, 45–78 (1997).
120. Pedregosa, F. et al. Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
121. Guyon, I. & Elisseeff, A. An introduction to variable and feature selection. *J. Mach. Learn. Res.* **3**, 1157–1182 (2003).
122. Kunz, L., Deuker, L., Zhang, H. & Axmacher, N. Tracking human engrams using multivariate analysis techniques. (Manahan-Vaughan, D., editor), *Handbook of in Vivo Neural Plasticity Techniques In Handbook of Behavioral Neuroscience*, Vol. 28, 481–508 (Elsevier, 2019).
123. Mumford, J. A., Turner, B. O., Ashby, F. G. & Poldrack, R. A. Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses. *NeuroImage* **59**, 2636–2643 (2012).
124. Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S. F. & Baker, C. I. Circular analysis in systems neuroscience: the dangers of double dipping. *Nat. Neurosci.* **12**, 535–540 (2009).
125. Fischl, B. et al. Automatically parcellating the human cerebral cortex. *Cereb. Cortex* **14**, 11–22 (2004).
126. Poldrack, R. A. Region of interest analysis for fMRI. *Soc. Cogn. Affect. Neurosci.* **2**, 67–70 (2007).
127. Cohen, J. *Statistical Power Analysis For The Behavioral Sciences* (Lawrence Erlbaum Associates, 1988).

128. Bonferroni, C. E. Teoria statistica delle classi e calcolo delle probabilità. *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze* **8**, 3–62 (1936).
129. Abramowitz, M. & Stegun, I. A. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables* 10th edn. (Dover Publications, Inc., 1964).
130. Johnson, S. G. The NLOpt Nonlinear-Optimization Package (NLOpt, 2019).
131. Powell, M. J. D. A direct search optimization method that models the objective and constraint functions by linear interpolation. In *Advances in Optimization and Numerical Analysis* (eds. Gomez, S. & Hennart, J.-P.) 51–67 (Springer, 1994).
132. Powell, M. J. D. Direct search algorithms for optimization calculations. *Acta Numerica* **7**, 287–336 (1998).
133. Benjamini, Y. & Hochberg, Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Statist. Soc. B* **57**, 289–300 (1995).
134. Bates, D., Mächler, M., Bolker, B. & Walker, S. Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* **67**, 1–48 (2015).
135. Tukey, J. W. Comparing individual means in the analysis of variance. *Biometrics* **5**, 99–114 (1949).
136. R Core Team. R: a language and environment for statistical computing. <https://www.R-project.org/> (2019).
137. Barr, D. J., Levy, R., Scheepers, C. & Tily, H. J. Random effects structure for confirmatory hypothesis testing: keep it maximal. *J. Mem. Lang.* **68**, 255–278 (2013).
138. Powell, M. J. D. Developments of NEWUOA for unconstrained minimization without derivatives. *IMA J. Num. Anal.* **28**, 649–664 (2007).
139. Powell, M. J. D. The BOBYQA algorithm for bound constrained optimization without derivatives. 26–46, Cambridge NA Report NA2009/06, (Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, 2009).
140. Lenth, R. emmeans: Estimated marginal means, aka least-squares means. <https://CRAN.R-project.org/package=emmeans>. R package version 1.3.4. (2019).
141. Allen, M., Poggiali, D., Whitaker, K., Marshall, T. R. & Kievit, R. A. Raincloud plots: a multi-platform tool for robust data visualization. *Wellcome Open Res.* **4**, 63 (2019).
142. Halchenko, Y. O. et al. datalad/datalad 0.11.5. <https://zenodo.org/record/3233911> (2019).
143. Esteban, O. et al. Analysis of task-based functional MRI data preprocessed with fMRIPrep. *Nat. Protoc.* **15**, 2186–2202 (2020).
144. Wagner, A. S. et al. *The DataLad Handbook*. <https://zenodo.org/record/3905791> (Zenodo, 2020).
145. Brand, A., Allen, L., Altman, M., Hlava, M. & Scott, J. Beyond authorship: attribution, contribution, collaboration, and credit. *Learned Publishing* **28**, 151–155 (2015).

## Acknowledgements

This work was funded by an Independent Max Planck Research Group grant awarded to N.W.S by the Max Planck Society (M.TN.A.BILD0004) and a Starting Grant awarded to N.W.S by the European Union (ERC-2019-StG REPLAY-852669). We also acknowledge financial support by the Max Planck Institute for Human Development. We thank Eran Eldar, Sam Hall-McMaster, and Ondrej Zika for helpful comments on a previous version of this manuscript, Gregor Caregnato for help with participant recruitment and data collection, Sonali Beckmann and Nadine Taube for assistance with MRI data acquisition,

Anika Löwe for assistance with data collection, Lion Schulz for help with behavioral data analysis, Michael Krause for support with cluster computing, members of the Max Planck Research Group NeuroCode for helpful feedback throughout the project, and all participants for their participation. L.W. is a pre-doctoral fellow of the International Max Planck Research School on Computational Methods in Psychiatry and Ageing Research (IMPRS COMP2PSYCH). The participating institutions are the Max Planck Institute for Human Development, Berlin, Germany, and University College London, London, UK. For more information, see <https://www.mps-ucl-centre.mpg.de/en/comp2psych>.

## Author contributions

The following list of author contributions is based on the CRediT taxonomy<sup>145</sup>. For details on each type of author contribution, please see Brand et al.<sup>145</sup>. Conceptualization: L.W., N.W.S.; Data curation: L.W.; Formal analysis: L.W., N.W.S.; Funding acquisition: N.W.S.; Investigation: L.W.; Methodology: L.W., N.W.S.; Project administration: L.W., N.W.S.; Resources: N.W.S.; Software: L.W., N.W.S.; Supervision: N.W.S.; Validation: L.W., N.W.S.; Visualization: L.W., N.W.S.; Writing - original draft: L.W., N.W.S.; Writing - review & editing: L.W., N.W.S.

## Funding

Open Access funding enabled and organized by Projekt DEAL.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41467-021-21970-2>.

**Correspondence** and requests for materials should be addressed to L.W. or N.W.S.

**Peer review information** *Nature Communications* thanks the anonymous reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

**Reprints and permission information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021

Dynamics of fMRI patterns reflect sub-second activation  
sequences and reveal replay in human visual cortex  
- Supplementary Information -

Lennart Wittkuhn<sup>1,2\*</sup> and Nicolas W. Schuck<sup>1,2\*</sup>

<sup>1</sup>Max Planck Research Group NeuroCode, Max Planck Institute for Human Development, Berlin, Germany

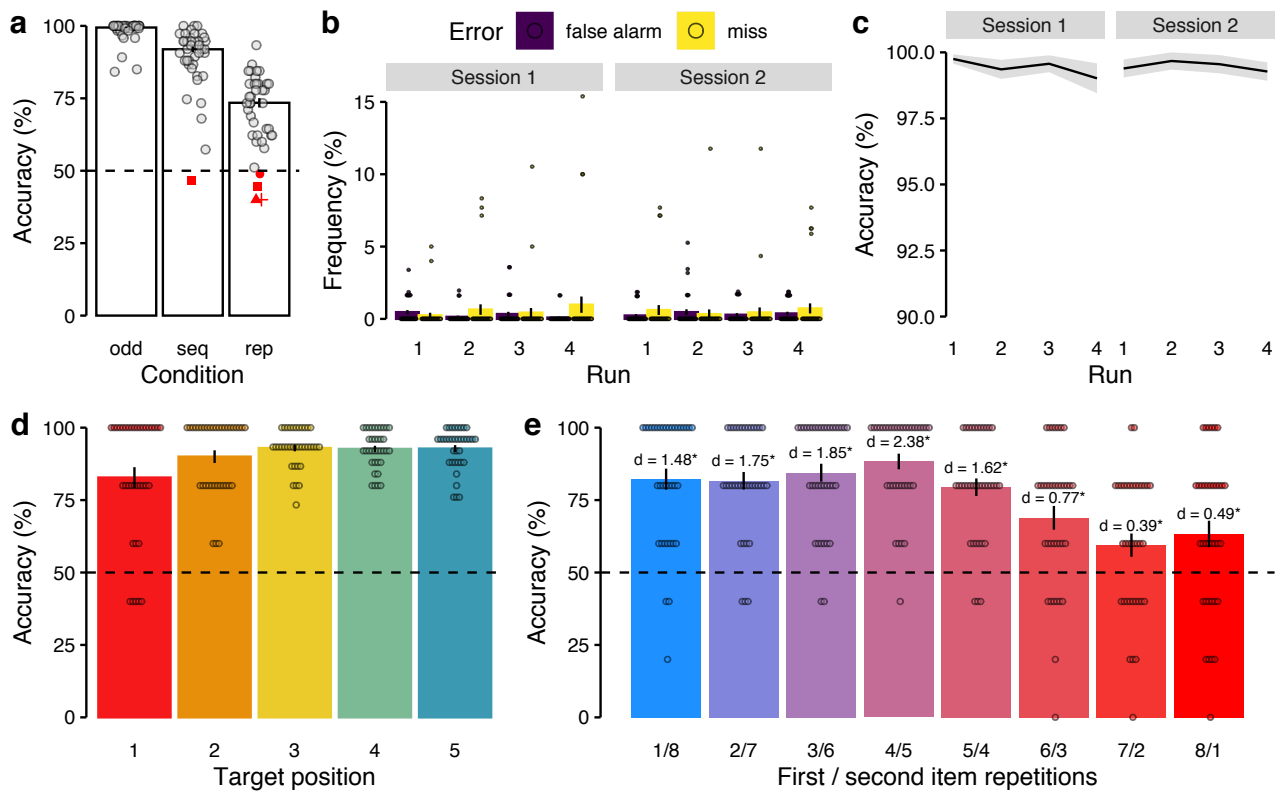
<sup>2</sup>Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Berlin, Germany

\*Corresponding authors, email [wittkuhn@mpib-berlin.mpg.de](mailto:wittkuhn@mpib-berlin.mpg.de) and [schuck@mpib-berlin.mpg.de](mailto:schuck@mpib-berlin.mpg.de)

## Supplementary Notes

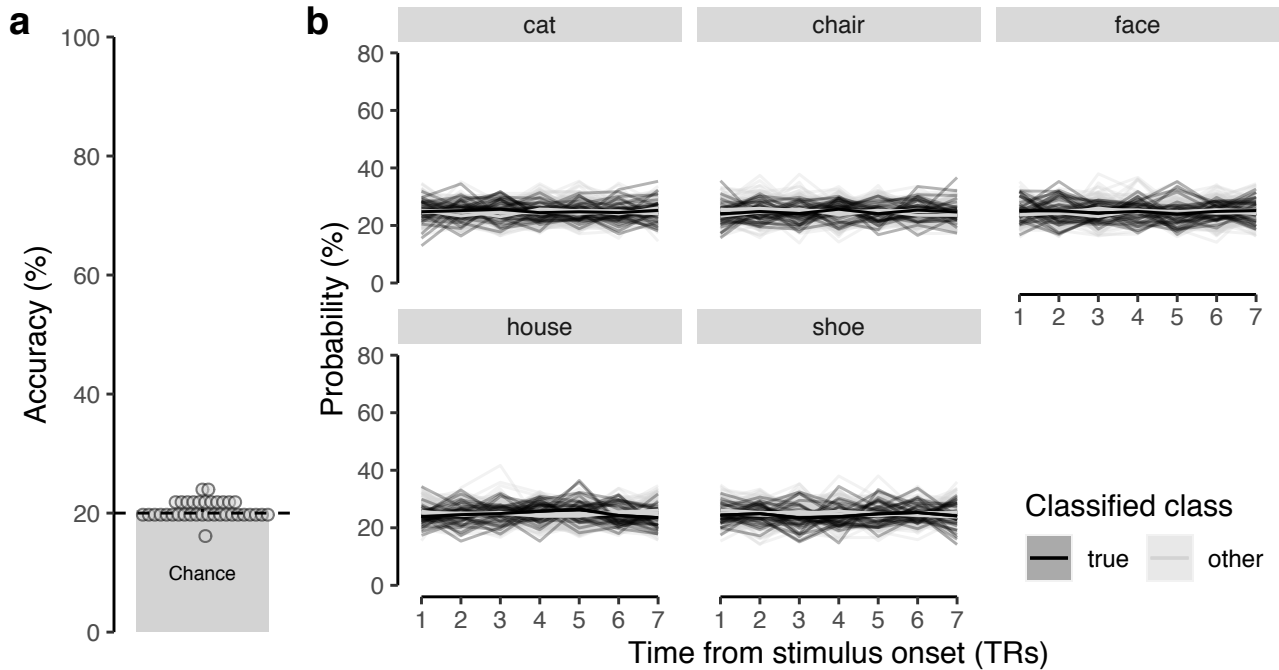
**Additional behavioral results** Attentive processing of the visual stimuli was a prerequisite to study the evoked activation patterns in visual and ventral temporal cortex. We therefore excluded all participants that performed below chance on either or both the repetition and sequence trials of the task. To this end, we removed all participants with a mean behavioral accuracy below the 50% chance level from all further analyses (Supplementary Fig. S1a). We also compared the relative proportion of misses and false alarms for each of the eight fMRI task runs in the experiment. To this end, we conducted a LME model with trial type (miss, false alarm), session (first, second) and session run (run 1–4) as fixed effects and included by-participant random intercepts and slopes. As shown in Supplementary Fig. S1b, misses ( $M = 0.55\%$ ) consistently occurred more frequently than false alarms ( $M = 0.30\%$ ),  $F_{1,501.00} = 4.12$ ,  $p = .043$ , which was consistent across task runs (no effects of session or run,  $ps \leq .703$ ). Our classification was performed using a leave-one-run-out approach. In order to examine whether the accuracy of behavioral performance on slow trials was stable across all task runs of the study, we conducted a LME model that included the eight task runs as the fixed effect of interest as well as random intercepts and slopes for each participant. The results showed no effect of task run indicating that the accuracy of behavioral performance was relatively stable across task runs,  $F_{1,92.72} = 0.13$ ,  $p = .72$  (Supplementary Fig. S1c). We examined whether mean behavioral accuracy on sequence trials was influenced by either the sequence speed or the serial position of the cued target image. A LME model including the sequence speed as a fixed effect and by-participant random intercepts and slopes indicated slightly lower but clear above-chance performance if the sequences were displayed at faster speeds,  $F_{1,35} = 4.27$ ,  $p = .046$  (Fig. 1f). A separate LME model including the target position as a fixed effect and by-participant random intercepts and slopes indicated lower but above-chance performance if the target image appeared at earlier serial positions,  $F_{1,42.02} = 9.92$ ,  $p = .003$  (Supplementary Fig. S1d). We focused the analysis of repetition trials on the forward and backward interference condition in the main text, but also examined performance for all intermediate repetition conditions and conducted a LME model with repetition condition as a fixed effect and by-participant random intercepts and slopes. Mean behavioral performance decreased with the number of second item repetitions,  $F_{1,39} = 57.43$ ,  $p < .001$  (Supplementary Fig. S1e). A series of eight one-sided one-sample t-tests indicated that for all repetition conditions mean behavioral accuracy was above the 50% chance level ( $t_{35} \geq 2.35$ ,  $ps \leq .012$ , FDR-corrected;  $ds \geq 0.39$ ).

**Decoding in the hippocampus is at chance level** We also conducted a separate leave-one-run-out classification analysis to decode the five stimulus categories from activation patterns in the hippocampus. To this end, the same decoding approach was used but activity patterns were extracted from an anatomical ROI centered on the hippocampus. The ROI was based on the same automated anatomical labeling of brain surface reconstructions from the individual T1w reference images that were used to create the anatomical masks of occipito-temporal brain regions. No GLM-based feature selection was performed on activity patterns from the hippocampus. Using the hippocampal masks in the leave-one-run-out cross-validation approach revealed that the average classification accuracy ( $M = 20.52\%$ ,  $SD = 1.49\%$ , range = 17–24%) did not differ from the chance baseline of 20%,  $t_{(35)} = 2.10$ , 95% CI [20.02, 21.03],  $p = .05$ ,  $d = 0.35$  (two-sided one-sample t-test, no multiple comparisons; see Supplementary Fig. S2). The implications of this finding are further discussed in the main manuscript.



**Supplementary Figure S1: Additional behavioral results.** (a) Mean behavioral performance (in %; y-axis) for the three trial conditions (x-axis). Dots / symbols represent mean data of a single participant with below-chance performance colored in red. Note, that the errorbars were derived from data of  $N = 36$  human participants after participants with below-chance performance were excluded and indicate mean values  $\pm 1$  SEM ( $N = 36$ ,  $t_s \geq 14.50$ ,  $p_s \leq .001$ ,  $d$ 's  $\geq 2.42$ , one-sided one-sample t-test per condition, no correction for multiple comparisons) (b) Mean frequency of incorrect slow trials (in %; y-axis) across the four task runs (x-axis) of each study session (panels), separately for false alarms (violet bars) and misses (yellow bars;  $N = 36$ ,  $F_{1,501.00} = 4.12$ ,  $p = .043$  for main effect of error type, LME model). (c) Mean behavioral accuracy on slow trials (in %; y-axis) across the four task runs (x-axis) of each study session (panels;  $N = 36$ ,  $F_{1,92.71} = 0.13$ ,  $p = .72$  for main effect of task run, LME model). (d) Mean behavioral accuracy on sequence trials (in %; y-axis) as a function of serial position of the target stimulus (x-axis;  $N = 36$ ,  $F_{1,42.02} = 9.92$ ,  $p = .003$  for main effect of target position, LME model). (e) Mean behavioral accuracy on repetition trials (in %; y-axis) for all repetition conditions (x-axis) compared to the 50% chance-level ( $N = 36$ ,  $t_s \geq 2.35$ ,  $p_s \leq .012$ ,  $d$ 's  $\geq 0.39$ , eight one-sided one-sample t-tests, FDR-corrected). Asterisks indicate  $p < .05$ , FDR-corrected. Effect sizes are indicated by Cohen's  $d$ . Horizontal dashed lines (in a, d, e) indicate 50% chance level. Errorbars (in a, b, d, e) and shaded areas (in c) represent  $\pm 1$  SEM. All statistics have been derived from data of  $N = 36$  human participants. Source data are provided as a Source Data file.

**Spatial correlations between classifier patterns** According to previous fMRI studies that investigated the neural representations of visual objects, corresponding multi-voxel patterns are often found to be widely distributed and largely overlapping within occipito-temporal brain regions [e.g., 55, 58–60]. To investigate the spatial distribution and overlap between the voxel activation patterns of the five visual stimuli used in our study, we visually examined their mean average activation patterns (an example from one participant is shown in Supplementary Fig. S3). Overall, the five stimuli appeared to activate a mix of overlapping and non-overlapping sets of voxels. In order to quantify this impression, we calculated the mean spatial correlation between voxel-activations for each participant. These analyses indicated that classifier patterns were slightly negatively correlated, with Pearson's correlations ranging from  $r = .02$  for the correlation between cat and face to  $r = -.44$  for the



**Supplementary Figure S2: Classification accuracy in the hippocampal mask.** (a) Cross-validated classification accuracy in decoding the five unique visual objects in hippocampal data during task performance (in %;  $N = 36$ ,  $t_{(35)} = 2.10$ , 95% CI [20.02, 21.03],  $p = 0.05$ ,  $d = 0.35$ , one two-sided one-sample t-test, no multiple comparisons). Chance level is 20% (dashed line). Each dot corresponds to averaged data from one participant. The errorbar represents  $\pm 1$  SEM. (b) Time courses (in TRs from stimulus onset; x-axis) of probabilistic classification evidence (in %; y-axis) for all five stimulus classes. No probability increases for any stimulus presented (black lines) on a given trial (gray panels) were found. Each line represents one participant. All statistics have been derived from data of  $N = 36$  human participants. Source data are provided as a Source Data file.

correlation between cat and house (see Supplementary Table S1 below).

**Supplementary Table S1:** Average correlation between average spatial patterns associated with each image category. Source data are provided as a Source Data file.

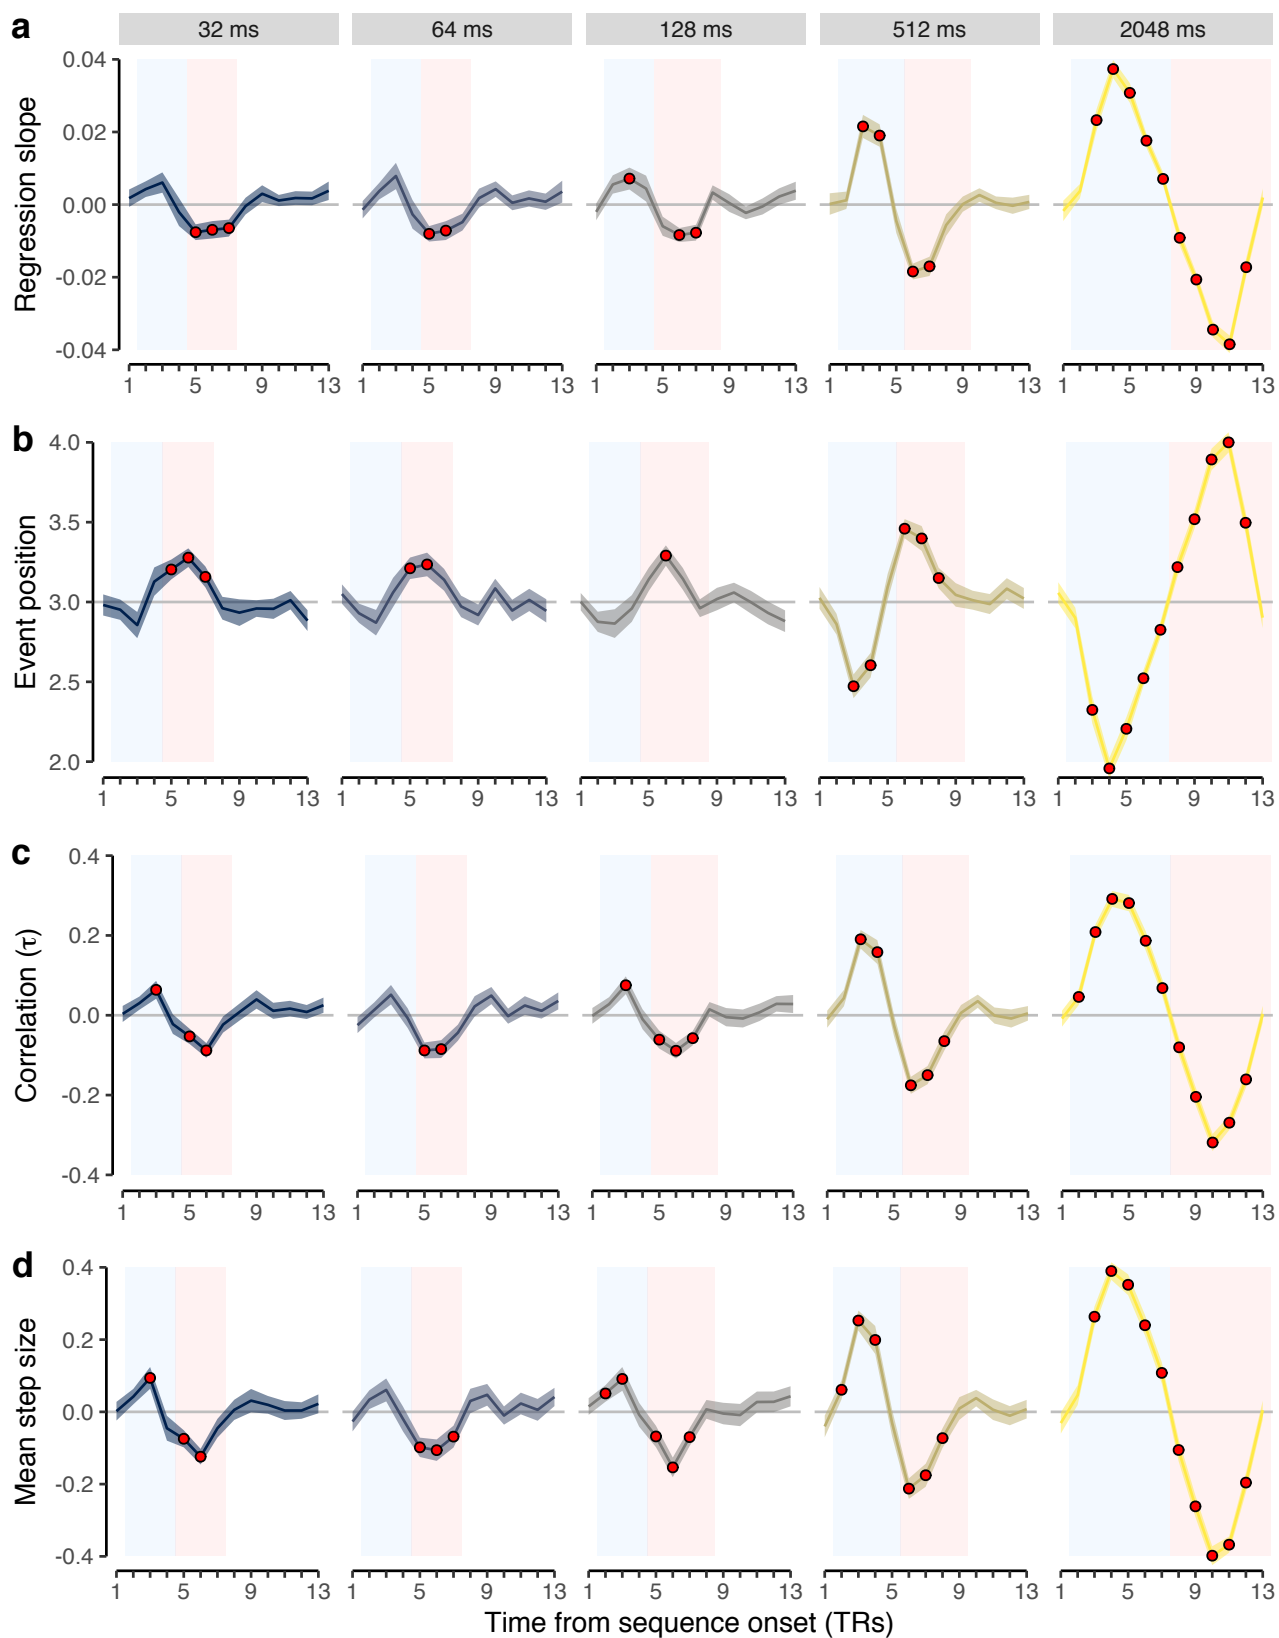
	Cat	Chair	Face	House	Shoe
Shoe	-.33	-.16	.29	-.15	1
House	-.44	-.20	-.37	1	
Face	.02	-.31	1		
Chair	-.23	1			
Cat	1				

**Additional information on single event and event sequence modeling** As reported in the main text, we described multivariate decoding time courses on slow trials by a sine wave response function that was fitted to the decoding time courses of all participants separately. Evaluating a single sine wave response function for three randomly selected example participants based on the individually fitted parameters indicated that the response functions capture the individual participant data well (Supplementary Fig. S4a). Based on the mean parameters across all participants we derived the mean response functions for each stimulus class which looked qualitatively similar (Supplementary Fig. S4b).



**Additional results for sequence trials** As reported in the main text, we investigated whether sequence order was evident in the relative pattern activation strength within a single measurement (i.e., within a single TR) and quantified sequential ordering by the slope of a linear regression between serial events and their classification probabilities. In addition, we repeated the same analysis using two different indices of linear association which produced qualitatively similar results. First, using ranked correlation coefficients (Kendall’s  $\tau$ ) between the serial event position and their classification probabilities as the index of linear association, we also found significant forward ordering in the forward period at sequence speeds of 128, 512 and 2048 ms ( $ts \geq 2.13$ ;  $ps \leq .05$ , FDR-corrected;  $ds \geq 0.36$ ) and significant backward ordering in the backward period for all speed conditions ( $ts \geq 4.24$ ;  $ps < .001$ , FDR-corrected;  $ds \geq 0.71$ ; Supplementary Fig. S5a–b). Second, we ordered the probabilities at every TR and calculated the mean step size (i.e., difference) between the probability-ordered event positions. Again, this analysis revealed qualitatively similar results, as we found significant forward ordering in the forward period at sequence speeds of 128, 512 and 2048 ms ( $ts \geq 2.25$ ;  $ps \leq .04$ , FDR-corrected;  $ds \geq 0.37$ ) and significant backward ordering in the backward period for all speed conditions ( $ts \geq 4.73$ ;  $ps < .001$ , FDR-corrected;  $ds \geq 0.79$ ; Supplementary Fig. S5c–d).

Next, we analyzed the time courses of linear associations in more detail. Specifically, for each index of linear association, we tested for sequentiality at every time point (i.e., at every TR) and conducted a series of two-sided one-sample t-tests comparing the sample mean at every time point against zero (the expectation of no order information). All  $p$  values were adjusted for multiple comparisons by controlling the FDR across all time-points within the forward and backward period and speed conditions (38 comparisons in total). This analysis produced consistent results for each index of linear association that was tested. For the mean regression slopes, this analysis revealed significant forward sequentiality at earlier time points for all speed conditions (TR 3 at 128 ms,  $t_{35} = 2.37$ ,  $p = .04$ ,  $d = 0.40$ ; TRs 3–4 at 512 ms,  $ts = 6.16$ ,  $ps < .001$ ,  $ds \geq 1.03$ ; TRs 3–7 at 2048 ms,  $ts = 7.78$ ,  $ps < .001$ ,  $ds \geq 1.03$ ; all  $ps$  FDR-corrected for 38 comparisons) except the 32 and 64 ms speed condition ( $ps \geq .08$ ). Furthermore, we found significant backward sequentiality at later time points for all speed conditions (TRs 5–7 at 32 ms,  $ts = 2.77$ ,  $ps \leq .02$ ,  $ds \geq 0.46$ ; TRs 5–6 at 64 ms,  $ps \leq .02$ ,  $ds \geq 0.46$ ; TRs 6–7 at 128 ms,  $ts = 3.53$ ,  $ps \leq .003$ ,  $ds \geq 0.59$ ; TRs 6–7 at 512 ms,  $ts = 6.41$ ,  $ps < .001$ ,  $ds \geq 1.07$ ; TRs 8–12 at 2048 ms,  $ts = 4.21$ ,  $ps < .001$ ,  $ds \geq 0.70$ ; all  $ps$  FDR-corrected for 38 comparisons; S6a). As can be seen in Supplementary Figs. S6b–d these results were qualitatively similar for all other indices of linear association tested (rank correlation coefficients and mean step size between probability-ordered event positions).



**Supplementary Figure S6: Classification time courses on sequence trials.** Time courses (in TRs from sequence onset; x-axis) of **(a)** mean linear regression coefficients (slope), **(b)** mean decoded serial event position with maximum probability for each sequence presentation speed (in ms; panels / colors), **(c)** mean correlation coefficients (Kendall's  $\tau$ ), and **(d)** mean step size between probability-ordered within-TR events. Shaded areas represent  $\pm 1$  SEM. All statistics have been derived from data of  $n = 36$  human participants. The blue and red rectangles indicate forward and backward period, respectively. Red dots indicate significant differences from baseline (horizontal gray line at zero; all  $ps \leq .05$ , FDR-corrected for 38 comparisons; two-sided one-sample t-tests). 1 TR = 1.25 s. Source data are provided as a Source Data file.

As reported in the main text, we verified that the sequentiality effects observed on sequence trials (Fig. 3b) are not only driven by the event with the maximum probability but that sequentiality is also present if the event with the maximum probability is removed. Examining the mean slope coefficients within the expected forward and backward period (adjusted by considering only four sequence events) after removing the event with the maximum probability showed that we could still find evidence for sequential ordering (Supplementary Fig. S7a). Significant forward ordering in the forward period was still evident at sequence speeds of 512 and 2048 ms ( $ts \geq 3.31$ ;  $ps \leq .004$ , ten two-sided one-sample t-tests, FDR-corrected;  $ds \geq 0.55$ ) and significant backward ordering in the backward period for all speed conditions ( $ts \geq 3.74$ ;  $ps \leq .002$ , ten two-sided one-sample t-tests, FDR-corrected;  $ds \geq 0.62$ ; Supplementary Fig. S7b) except the 32 and 128 ms speed conditions ( $p \geq .20$ ). The main analysis reported in the Results section highlighted an apparent asymmetry in detecting forward and backward sequentiality. To determine the extent to which this asymmetry was driven by the first or last item in the sequence, we conducted two additional control analyses by either removing the first or last sequence item from the analysis. Removing the first sequence item did not change the observed sequentiality effects qualitatively (Supplementary Fig. S7c) as we still found significant forward ordering in the forward period at sequence speeds of 512 and 2048 ms ( $ts \geq 5.72$ ;  $ps < .001$ , FDR-corrected;  $ds \geq 0.95$ ) and significant backward ordering in the backward period for all speed conditions ( $ts \geq 2.65$ ;  $ps \leq .02$ , ten two-sided one-sample t-tests, FDR-corrected;  $ds \geq 0.44$ ; Supplementary Fig. S7d). Removing the last sequence item, in contrast, made any significant sequentiality disappear for speed conditions of 128 ms or faster ( $p \geq .27$ ), while forward and backward sequentiality were still evident at sequence speeds of 512 ms and 2048 ms ( $ts \geq 3.55$ ;  $ps \leq .05$ , ten two-sided one-sample t-tests, FDR-corrected;  $ds \geq 0.59$ ; Supplementary Figs. S7e–f).

**Additional analyses of repetition trials** We conducted two additional analyses for the data on repetition trials. First, we analyzed the effect of event duration (number of repetitions) on event probability in more detail by calculating the average event probability for each event type (first, second, and averaged non-sequence) as a function of event duration (number of repetitions). Importantly, while we focused only on the two repetition conditions with the highest degree of interference before, we now also included the data from all intermediate repetition trial types. As before, we averaged the probabilities for each serial event type but this time as a function of how often each item type was repeated in any given trial. Then, in order to test how likely we were in decoding each serial event type (first, second, non-sequence), when each item was only shown briefly once, we conducted three independent pairwise two-sample t-tests comparing the mean probabilities of all three event types with one another (correcting for multiple comparisons using Bonferroni correction). The results reported in the main text focused on the two repetition conditions with the strongest expected effects of forward and backward interference. Additionally, we characterized the effect of event duration (number of repetitions) in more detail by analyzing the average probability of event types (first, second, non-sequence) as a function of event duration also for all intermediate repetition conditions. The results revealed a main effect of event type (first, second, non-sequence),  $F_{2,278.97} = 23.99$ ,  $p < .001$  and event duration (number of repetitions),  $F_{1,58.73} = 183.10$ ,  $p < .001$  as well as an interaction between event type and event duration,  $F_{2,753.00} = 52.53$ ,  $p < .001$  (see Supplementary Fig. S8). In order to further characterize the origin of this interaction, we also conceived a reduced model that did not include the data from non-sequence events. The results of this reduced model again showed a main effect

of event type (first, second),  $F_{1,350.19} = 12.11$ ,  $p < .001$  and event duration (number of repetitions),  $F_{1,125.87} = 187.86$ ,  $p < .001$  but no interaction between event type and event duration,  $F_{1,501.90} = 0.10$ ,  $p = .75$ . If only shown briefly, the second event had a mean probability ( $M = 17.07\%$ ,  $SD = 5.42\%$ ) that was higher than for the first event ( $M = 13.50\%$ ,  $SD = 6.04\%$ ),  $t_{(35)} = 2.45$ ,  $p = .02$  and the averaged non-sequence items ( $M = 7.75\%$ ,  $SD = 2.93\%$ ),  $t_{(35)} = 8.98$ ,  $p < .001$  while the average probability of the first event was also higher compared to the out-of-sequence items,  $t_{(35)} = 5.53$ ,  $p < .001$  (all  $ps$  were adjusted for six multiple comparisons, using the Bonferroni correction). If the event duration was prolonged (eight consecutive repetitions) the second event had a mean probability ( $M = 31.63\%$ ,  $SD = 6.94\%$ ) that was significantly different from the first event ( $M = 26.87\%$ ,  $SD = 8.34\%$ ),  $t_{(39)} = 2.59$ ,  $p = .01$  and the averaged non-sequence items ( $M = 7.69\%$ ,  $SD = 2.69\%$ ),  $t_{(35)} = 18.96$ ,  $p < .001$  while the average probability of the first event was also higher compared to the non-sequence items,  $t_{(35)} = 12.52$ ,  $p < .001$  (all  $ps$  were adjusted for six multiple comparisons, using the Bonferroni correction). These effects were attenuated but qualitatively similar when data from all TRs were considered.

We asked whether we would be more likely to decode items that were part of the sequence actually shown to participants (within-sequence items) as compared to items not part of the sequence (out-of-sequence items). To this end, we assessed if the serial events 1 and 2 were more likely to be decoded in the repetition trials than other events. As before, we identified the item with the highest classifier probability at every TR of each trial and then calculated the relative frequency of each item in the decoded sequence of events. These frequencies were then averaged separately for each repetition condition across all trials and participants. Next, using paired t-tests, we performed two statistical tests: First, we tested how well we were able to decode a single briefly presented item in a 32 ms sequence compared to items that were not presented, when the item is followed by a statistical representation that could mask its activation pattern (short  $\rightarrow$  long trials). Second, we tested how well we were able to decode a single briefly presented item (first serial event) in a 32 ms sequence compared to items that were not part of the sequence, when the item (last serial event) is followed by a random statistical signal, for example, during an ITI (long  $\rightarrow$  short trials).

Analyzing the average proportion of decoded serial events across all TRs for the *backward interference* and *forward interference* conditions separately revealed a main effect of serial event type (first, second, averaged out-of-sequence),  $F_{2,234} = 40.70$ ,  $p = 6.80 \times 10^{-16}$ . No main effect of repetition condition (short  $\rightarrow$  long versus long  $\rightarrow$  short) was found,  $F_{1,234} = 0.08$ ,  $p = .78$ , but an interaction between serial event position and repetition condition,  $F_{2,234} = 23.92$ ,  $p = 3.54 \times 10^{-10}$  (see Fig. 4e). Post-hoc comparisons indicated that in the short  $\rightarrow$  long condition the longer second event had a higher frequency ( $M = 29.0\%$ ) compared to the out-of-sequence ( $M = 17.4\%$ ) as well as the short, first event ( $M = 18.9\%$ ,  $ps < .0001$ ). The short first event did not differ from the out-of-sequence events ( $p = .47$ , Tukey-correction for three comparisons). In the long  $\rightarrow$  short condition, in contrast, there was no difference between the long first ( $M = 24.6\%$ ) and short second event ( $M = 22.3\%$ ,  $p = .17$ , Tukey-correction for three comparisons) but significant differences between both within-sequence items and the averaged out-of-sequence ( $M = 17.7\%$ ) items (both  $ps < .001$ , Tukey-correction for three comparisons).

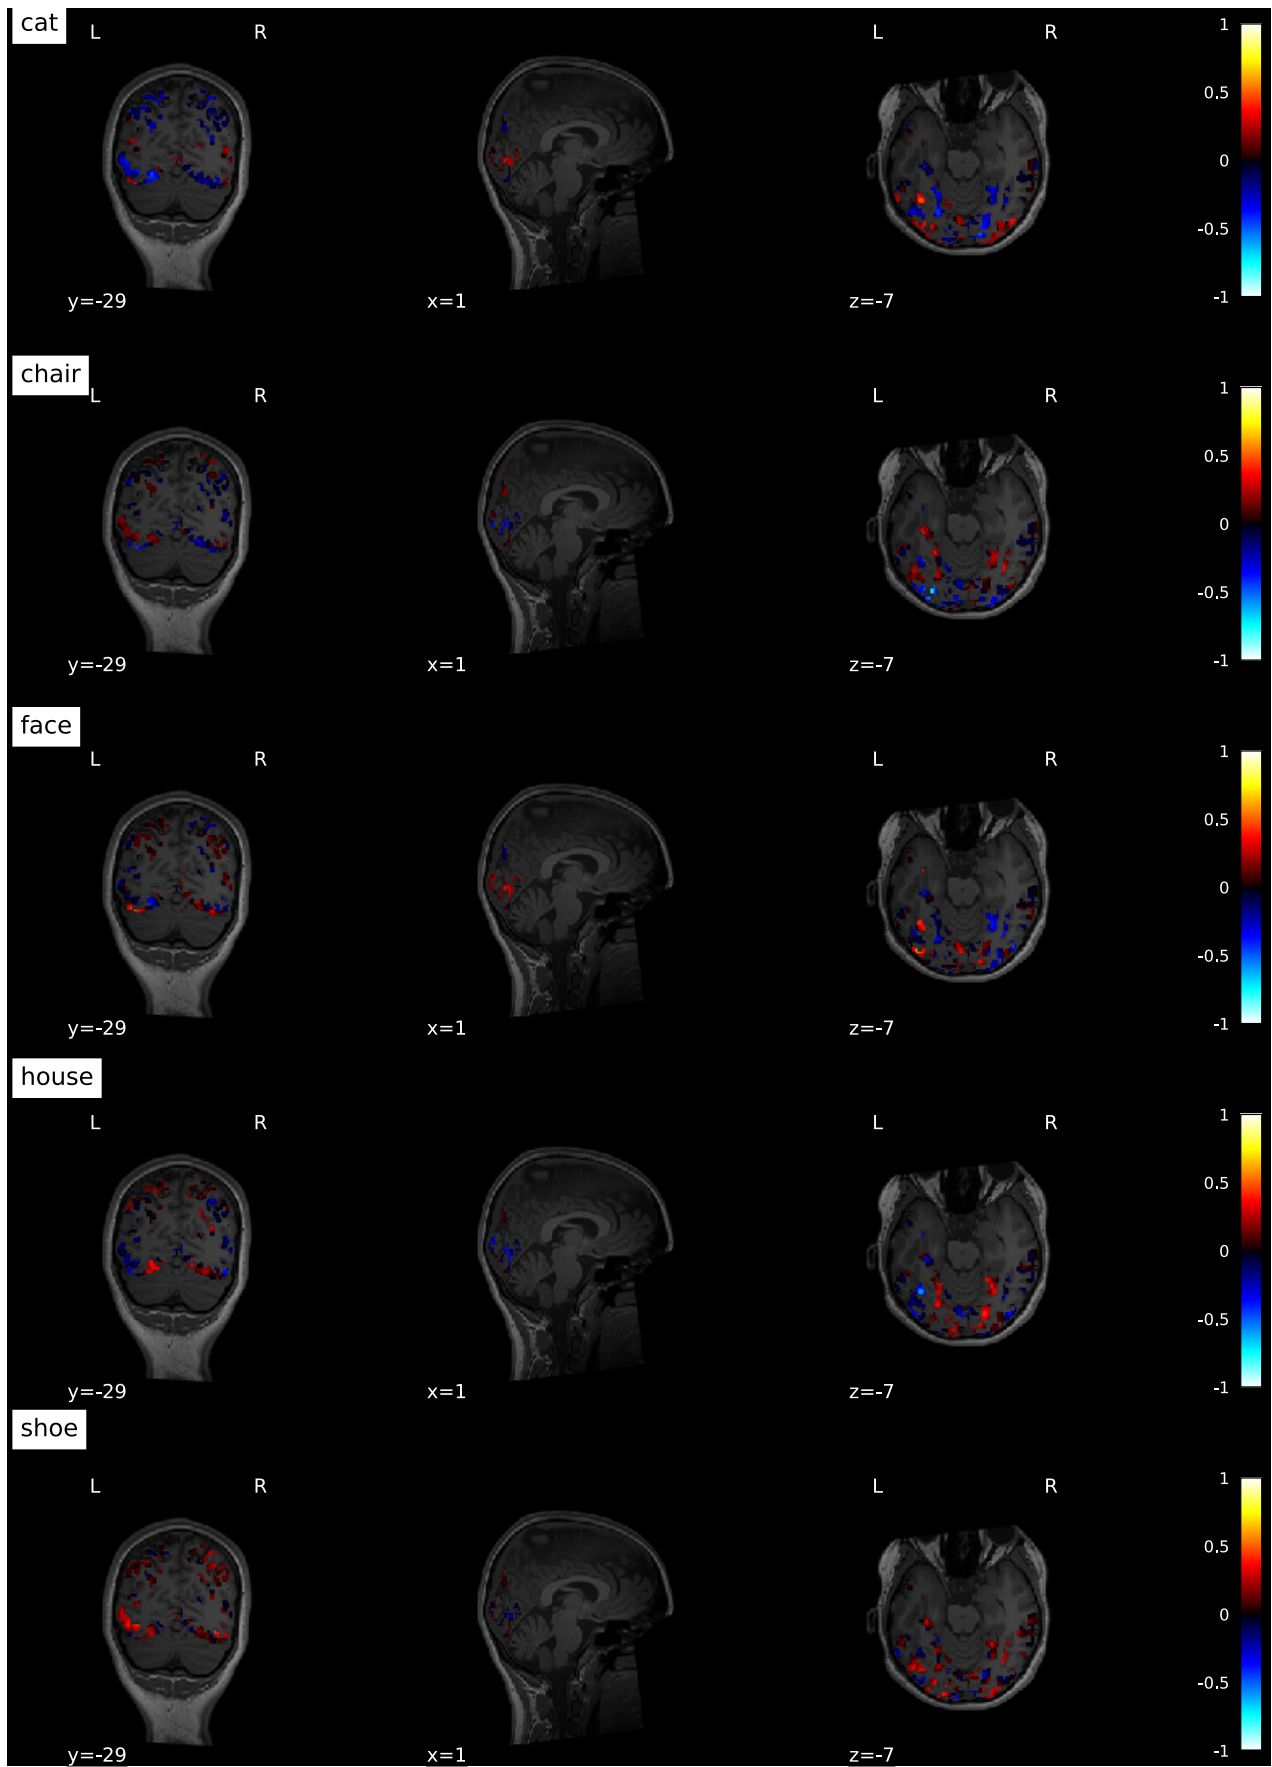
Analyzing the mean probability for the three event types (first, second, and out-of-sequence events) on repetition trials as a function of the absolute event occurrence per trial using data from all 13 TRs revealed a main effect of event type (first, second, out-of-sequence),  $F_{2,915} = 14.31$ ,  $p < .001$  and event

duration (number of repetitions),  $F_{1,915} = 68.97, p < .001$  as well as an interaction between event type and event duration,  $F_{2,915} = 17.90, p < .001$  (see Fig. 4d). In order to further characterize the origin of this interaction, we also conceived a reduced model that did not include the data from out-of-sequence events. The results of this reduced model again showed a main effect of event type (first, second),  $F_{1,597} = 10.92, p = .001$  and event duration (number of repetitions),  $F_{1,597} = 78.92, p < .001$  but no interaction between event type and event duration,  $F_{1,597} = 0.18, p = 0.68$ . If only shown briefly, the second event had a mean probability ( $M = 14.41\%$ ,  $SD = 4.53\%$ ) that was higher than for the first event ( $M = 12.02\%$ ,  $SD = 4.78\%$ ),  $t_{(39)} = 2.46, p = .03$  and the averaged out-of-sequence items ( $M = 10.28\%$ ,  $SD = 2.88\%$ ),  $t_{(39)} = 5.80, p < .001$  while the average probability of the first event was also higher compared to the out-of-sequence items,  $t_{(39)} = 2.52, p = .03$  (all  $p$  values were adjusted for six multiple comparisons, using the FDR correction). If the event duration was prolonged (eight consecutive repetitions) the second event had a mean probability ( $M = 19.37\%$ ,  $SD = 6.44\%$ ) that was not significantly different from the first event ( $M = 16.54\%$ ,  $SD = 4.75\%$ ),  $t_{(39)} = 2.27, p = .06$  but from the averaged out-of-sequence items ( $M = 9.75\%$ ,  $SD = 3.05\%$ ),  $t_{(39)} = 9.36, p < .001$  while the average probability of the first event was also higher compared to the out-of-sequence items,  $t_{(39)} = 7.99, p < .001$  (all  $p$  values were adjusted for six multiple comparisons, using the FDR correction).


We also analyzed the trial-wise proportion of transition types between consecutively decoded events using data from all 13 TRs following stimulus onset. This analysis revealed that in the short  $\rightarrow$  long condition the mean trial-wise proportion of forward transitions ( $M = 6.50$ ) was higher than the mean proportion of outward transitions ( $M = 2.48$ ),  $t_{(39)} = 4.82, p < .001$  and also differed from the mean trial-wise proportion of outside transitions ( $M = 1.28$ ),  $t_{(39)} = 6.14, p < .001$  (all  $p$  values were corrected for four comparisons using Bonferroni correction; see Fig. 4f). Similarly, in the long  $\rightarrow$  short condition, the mean trial-wise proportion of forward transitions ( $M = 6.80$ ) was higher than the mean proportion of outward transitions ( $M = 2.58$ ),  $t_{(39)} = 6.11, p < .001$  and also differ compared to the mean trial-wise proportion of outside transitions ( $M = 1.18$ ),  $t_{(39)} = 7.71, p < .001$  (all  $p$  values were corrected for four comparisons using Bonferroni correction).

**Repeating analyses of repetition trials using data from all TRs** As reported in the main text, we focused the analyses of repetition trials on data from a relevant period of six TRs (from the second to the seventh TR) and the two trial conditions with maximum forward and backward interference, respectively. Here, we report results of the same analyses repeated using data from all TRs. The estimated probabilities of each stimulus class given the data for all repetition conditions are shown in Supplementary Fig. S9. Analyzing the mean probabilities of the different event types (first, second, out-of-sequence) using data from all TRs (see Supplementary Fig. S10a) revealed qualitatively similar results. Event type still influenced the average decoding probability,  $F_{2,54.79} = 41.67, p < .001$  (see Supplementary Fig. S10b). Post-hoc comparisons indicated that sequence items had a higher mean probability than out-of-sequence (9.94%) items (both  $ps < .001$ , Tukey-correction for three comparisons), while the second (16.73%) and first (14.42%) within-sequence event type also differed ( $p = .045$ , Tukey-correction for three comparisons). Repeating the analysis for the forward and backward interference conditions using data from all TRs again revealed smaller but qualitatively similar effects, with a main effect of event type (first, second, out-of-sequence),  $F_{2,44.29} = 55.22, p < .001$ , an interaction between event type and duration,  $F_{2,140.000} = 40.38, p < .001$ , and no main

effect of duration (number of repetitions),  $F_{1,116.04} = 0.15$ ,  $p = .70$  (see Supplementary Fig. S10c). Post-hoc comparisons indicated that in the forward interference condition the longer second event had a higher probability (19.32%) compared to both the out-of-sequence ( $M = 10.11\%$ ) and the short, first event ( $M = 11.94\%$ ,  $ps < .001$ , Tukey-correction for three comparisons). As reported in the main text, when using data from all TRs, the short first event did not differ from the out-of-sequence events ( $p = .09$ , Tukey-correction for three comparisons). In the backward interference condition, in contrast, there was a significant difference between the long first (16.91%) and short second event (14.15%,  $p = .04$ , Tukey-correction for three comparisons) as well as significant differences between both within-sequence items and the averaged out-of-sequence (9.77%) items ( $ps < .001$ , Tukey-correction for three comparisons). We also repeated the analysis investigating trial-wise proportions of transitions between consecutively decoded events using data from all TRs. Based on the full transition matrix (see Supplementary Fig. S10e), this analysis revealed qualitatively similar effects (Supplementary Fig. S10d): Forward transitions (3.84%) between the two sequence items were as frequent as outward transitions (2.74%,  $t_{(35)} = 2.61$ ,  $p = .05$ , Bonferroni-corrected for four comparisons) but more frequent than outside transitions (2.32%,  $t_{(35)} = 2.61$ ,  $p = .02$ , Bonferroni-corrected for four comparisons) in the forward interference condition. The same was true for the backward interference condition (forward transitions: 4.54%; outwards transitions: 2.97%; outside transitions: 2.34%, all  $ts \geq 3.56$ , all  $ps < .001$ ; Bonferroni-corrected for four comparisons).

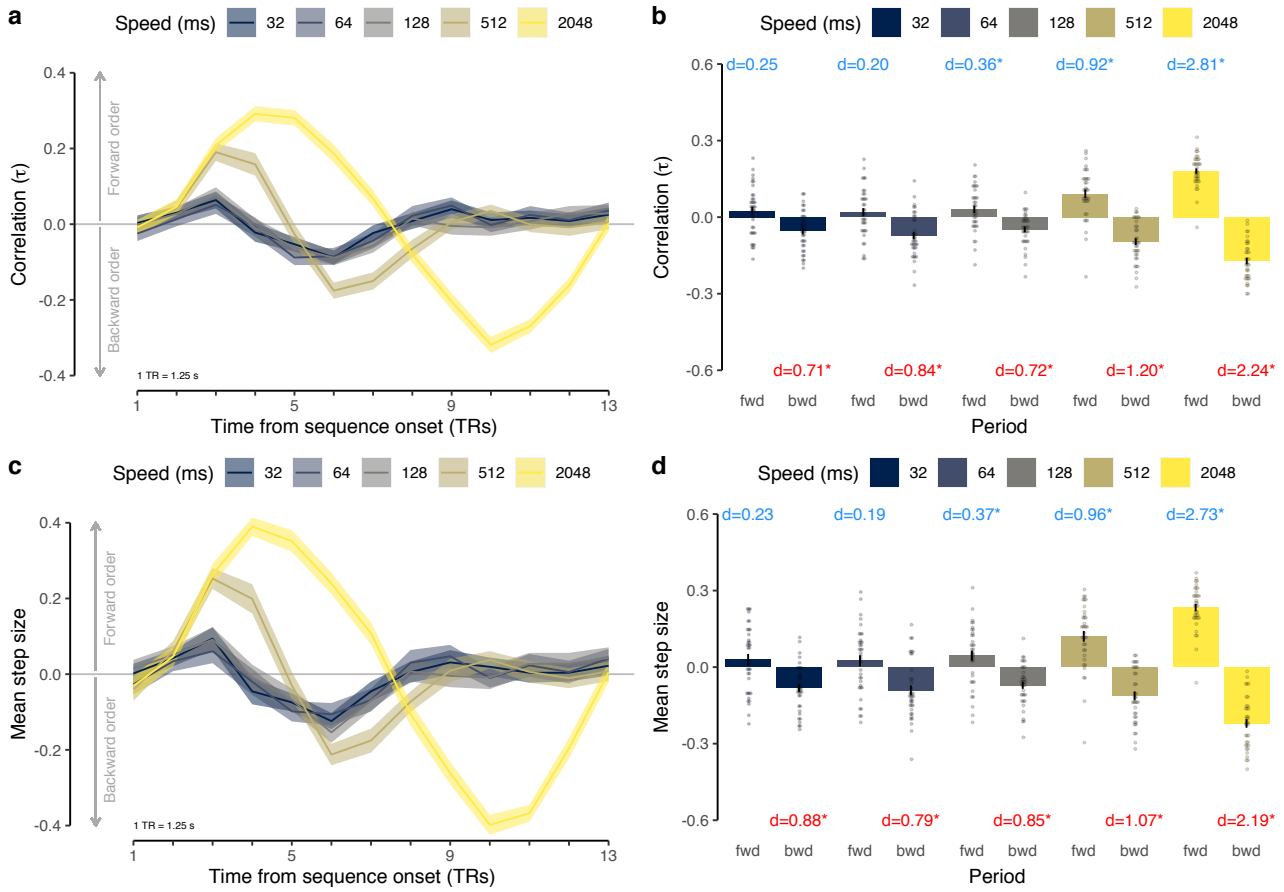


**Supplementary Figure S3: Spatial distribution of mean voxel activations in one example participant for five stimuli** Averaged patterns of voxel activations used for multivariate pattern analysis (colors indicate z-scored voxel activations) for the five decoded stimuli (horizontal panels) in one example participant (sub-01) shown against the participant's individual defaced structural scan.

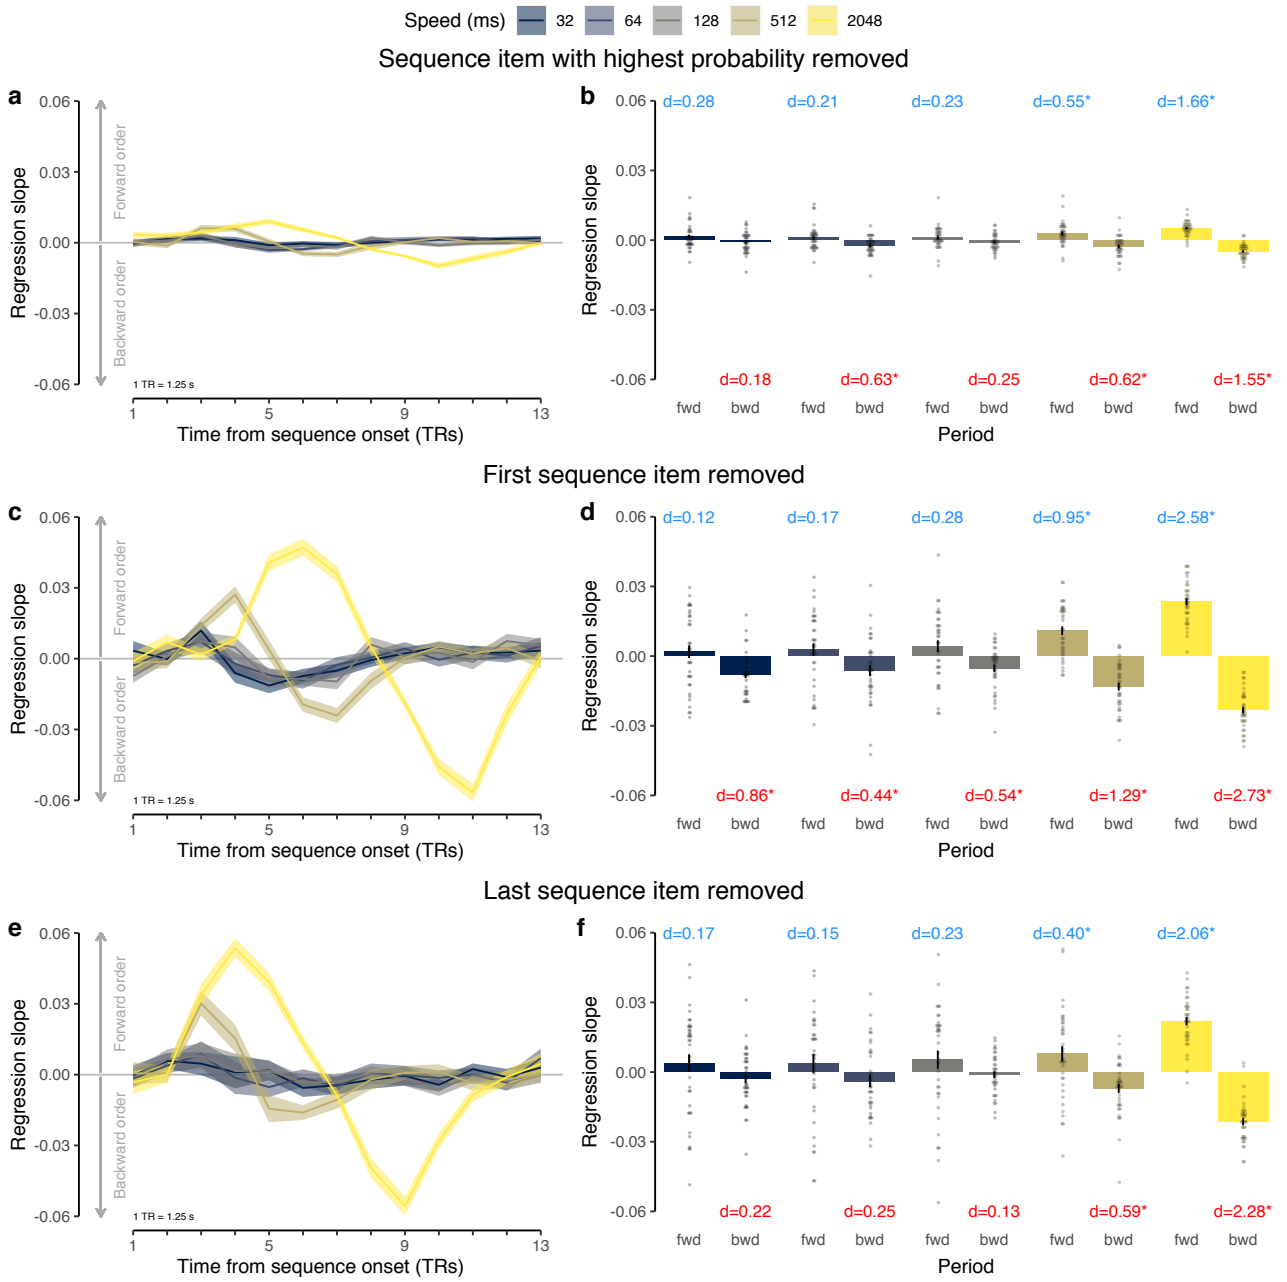


**Supplementary Figure S4: Individual fits of sine wave response function to probabilistic classifier evidence.** (a) Time courses (in TRs from stimulus onset; x-axis) of probabilistic classifier evidence (in %; y-axis) generated by the sine wave response function with fitted parameters (black dotted line) or the true data (gray line and dots) separately for the five stimulus classes (vertical panels) and three randomly chosen example participants (horizontal panels). (b) Time courses (in TRs from stimulus onset; x-axis) of mean probabilistic classifier evidence (in %; y-axis) averaged separately for each participant (gray semi-transparent lines) and stimulus class (vertical panels) or predicted by the sine wave response model based on fitted parameters averaged across all participants (black line). 1 TR = 1.25 s. Source data are provided as a Source Data file.

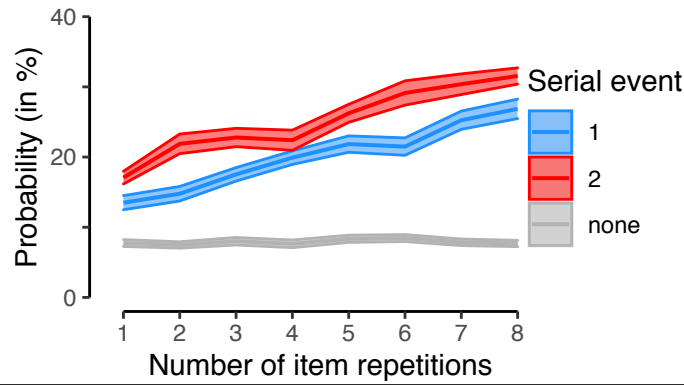




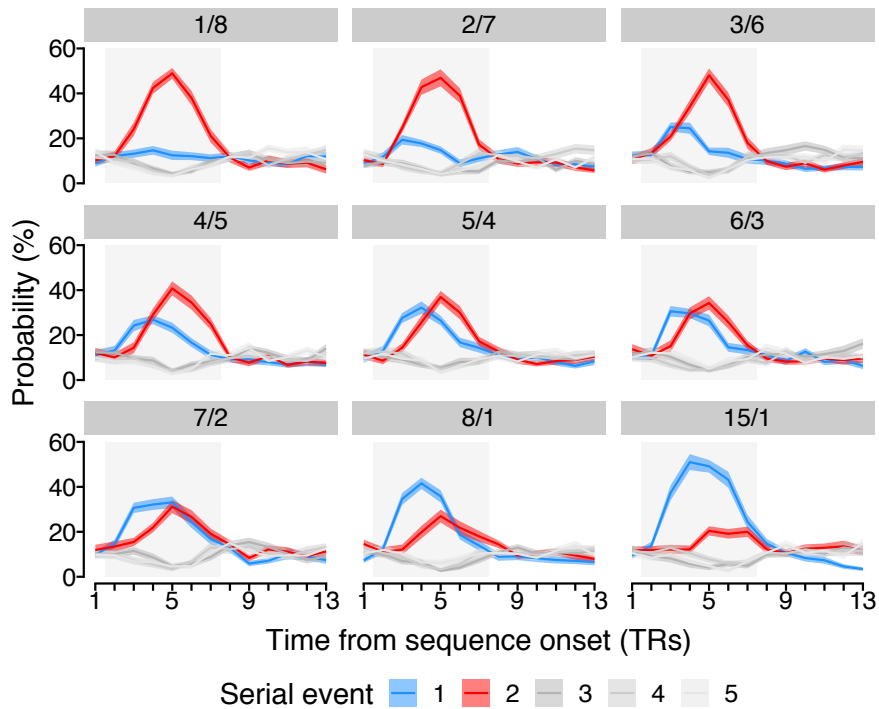
**Supplementary Figure S5:** (a) Time courses (in TRs from sequence onset; x-axis) of mean ranked correlation coefficients between serial event position and classification probabilities (Kendall's  $\tau$ ; y-axis) for each speed condition (in ms; colors) on sequence trials. (b) Mean ranked correlation coefficients (Kendall's  $\tau$ ; y-axis) as a function of time period (forward versus backward; x-axis) and sequence speed (in ms; colors;  $N = 36$ ,  $ts \geq 2.13$ ,  $ps \leq .05$ ,  $ds \geq 0.36$  (significant tests only), ten two-sided one-sample t-tests against zero, FDR-corrected). (c) Time courses (in TRs from sequence onset; x-axis) of the mean step size between probability-ordered within-TR events (y-axis) for each speed condition (in ms; colors) on sequence trials. (d) Mean within-TR step-size (y-axis) as a function of time period (forward versus backward; x-axis) and sequence presentation speed (in ms; colors;  $N = 36$ ,  $ts \geq 2.25$ ,  $ps \leq .04$ ,  $ds \geq 0.37$  (significant tests only), ten two-sided one-sample t-tests against zero, FDR-corrected). Each dot in (b) and (d) represents averaged data of one participant. Shaded areas in (a), (c) and errorbars in (b), (d) represent mean values  $\pm 1$  SEM. All statistics have been derived from data of  $n = 36$  human participants. 1 TR = 1.25 s. Asterisks indicate significant differences from baseline. Source data are provided as a Source Data file.



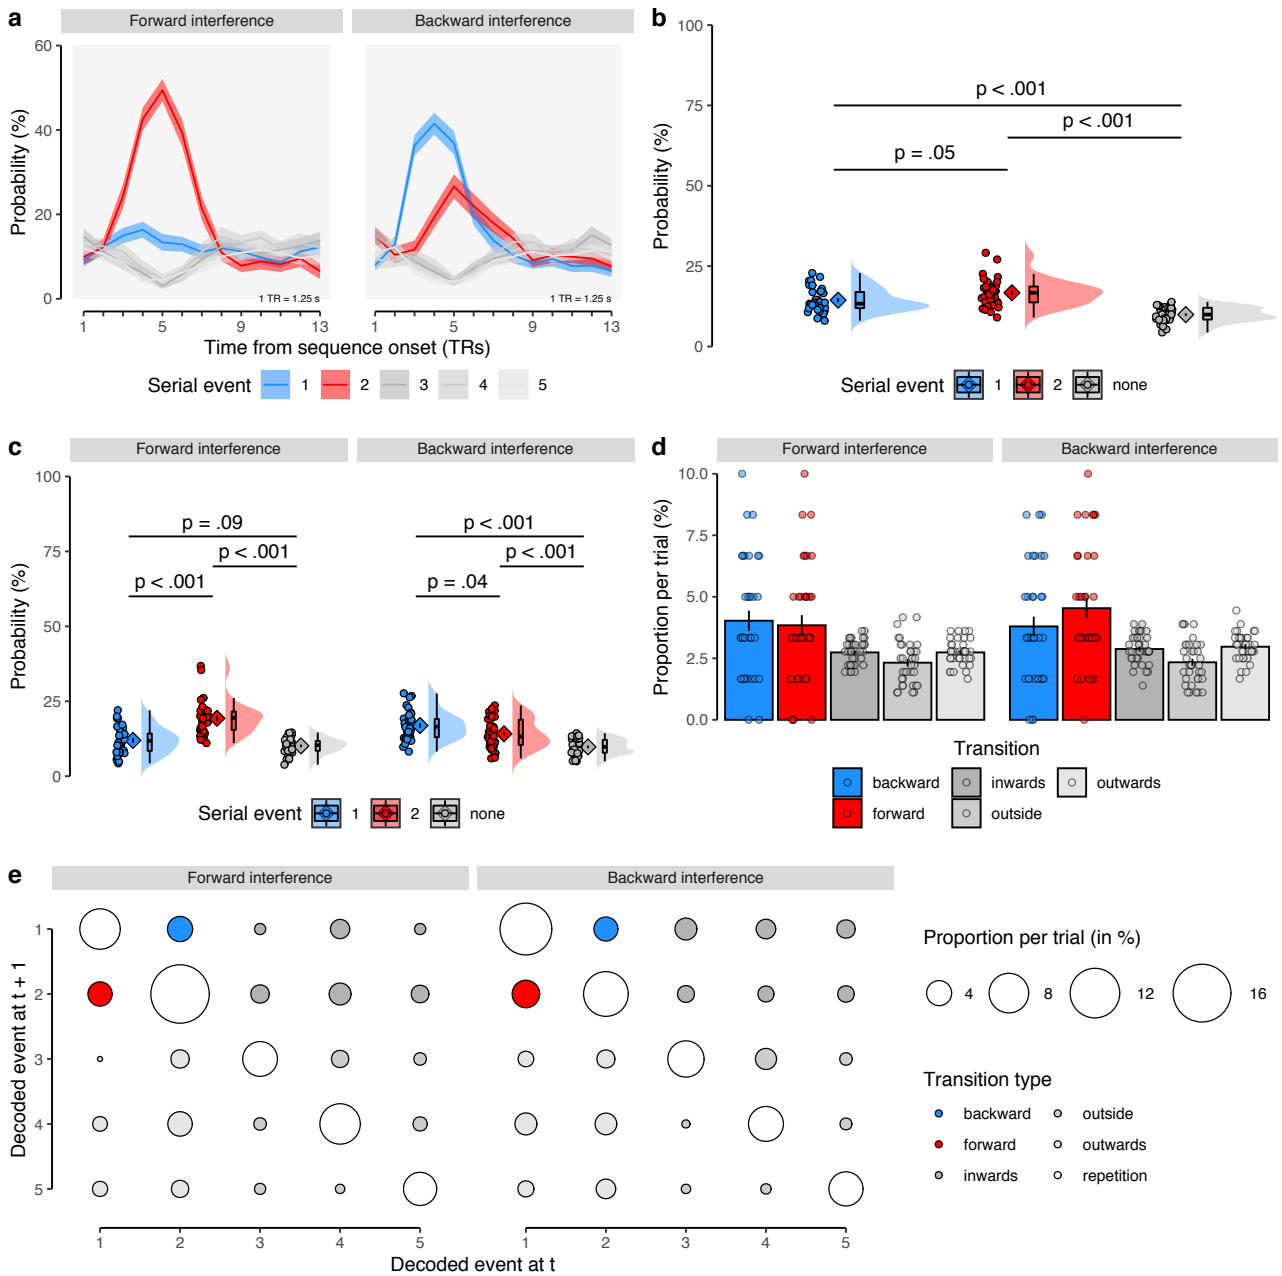
**Supplementary Figure S7: Effects of sequence item removal on sequentiality metrics.** (a, c, e) Time courses (in TRs from sequence onset; x-axis) of mean slope coefficients of a linear regression between serial event position and classifier probability (y-axis) for each speed condition (in ms; colors) on sequence trials after removal of (a) the sequence item with the highest classification probability, (c) the first sequence item, (e) the last sequence item. (b, d, f) Mean slope coefficients (y-axis) as a function of time period (forward versus backward; x-axis) and sequence speed (in ms; colors) after removal of (b) the sequence item with the highest classification probability, (d) the first sequence item, (f) the last sequence item ( $N = 36$ , ten two-sided one-sample t-tests against zero for each panel, FDR-corrected). Each dot represents averaged data of one participant. Shaded areas in (a, c, e) and errorbars in (b, d, f) represent  $\pm 1$  SEM. All statistics have been derived from data of  $n = 36$  human participants. 1 TR = 1.25 s. Source data are provided as a Source Data file.



**Supplementary Figure S8: Effects of event duration (element repetition)** Mean probability (in %; y-axis) as a function of the number of item repetitions (i.e., total event duration), separately for event types (first, second, and out-of-sequence events; colors) based on data of all TRs. Shaded areas represent  $\pm 1$  SEM. All statistics have been derived from data of  $n = 36$  human participants. Source data are provided as a Source Data file.



**Supplementary Figure S9: Time courses of probabilistic classifier evidence for all repetition conditions.** Time courses (in TR from sequence onset; x-axis) of probabilistic classifier evidence (in %; y-axis) on repetition trials grouped by event type (colors), separately for each repetition condition (gray panels). Each panel indicates the number of repetitions per sequence event (e.g., the top-left panel indicates 1 versus 8 repeats of the first versus second event). Time-courses of classifier evidence for the first and second event are shown in blue and red, respectively, while all other stimuli that were not part of the sequence are shown in three shades of gray. Shaded areas represent  $\pm 1$  SEM. All statistics have been derived from data of  $n = 36$  human participants. 1 TR = 1.25 s. Source data are provided as a Source Data file.



**Supplementary Figure S10: Ordering of two-item pairs on repetition trials across all TRs.** (a) Time courses (in TRs from sequence onset; x-axis) of probabilistic classifier evidence (in %) in repetition trials, color-coded by event type (first/second/non-sequence, see legend). Data shown separately for forward (left) and backward (right) interference conditions. Gray background indicates relevant time period across all TRs. Shaded areas represent  $\pm 1$  SEM. All statistics have been derived from data of  $n = 36$  human participants. 1 TR = 1.25 s. (b) Mean probability of event types averaged across all TRs in the relevant time period, as in (a). Each dot represents one participant, the probability density of the data is shown as rain cloud plots [cf. 61]. Boxplots indicate the median and interquartile range (IQR). The lower and upper hinges correspond to the first and third quartiles (the 25<sup>th</sup> and 75<sup>th</sup> percentiles). The upper whisker extends from the hinge to the largest value no further than  $1.5 \times$  IQR from the hinge (where IQR is the inter-quartile range, or distance between the first and third quartiles). The lower whisker extends from the hinge to the smallest value at most  $1.5 \times$  IQR of the hinge. The barplots show the sample mean and errorbars indicate  $\pm 1$  SEM.  $N = 36$ ,  $t_s \geq 2.49$ ,  $p_s \leq .045$ , LME model with post-hoc Tukey's HSD tests. (c) Average probability of event types, separately for conditions as in (a), plots as in (b). ( $N = 36$ ,  $t_s \geq 2.11$ ,  $p_s < .09$ ), LME model with post-hoc Tukey's HSD tests). (d) Mean trial-wise proportion of each transition type, separately for forward/backward conditions, as in (a) ( $N = 36$ ,  $t_s \geq 2.61$ ,  $p_s \leq .05$ , four two-sided paired t-tests, Bonferroni-corrected). (e) Transition matrix of decoded images indicating mean proportions per trial, separately for the forward and backward condition (left/right). Transition types highlighted in colors (see legend). All statistics have been derived from data of  $n = 36$  human participants who participated in one experiment. Source data are provided as a Source Data file.

## E Paper III

**Wittkuhn, L.**, Krippner, L. M., & Schuck, N. W. (2022). Statistical learning of successor representations is related to on-task replay. *bioRxiv*. doi: [10.1101/2022.02.02.478787](https://doi.org/10.1101/2022.02.02.478787)

This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license. To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc-nd/4.0/>.

# Statistical learning of successor representations is related to on-task replay

Lennart Wittkuhn<sup>1,2,\*</sup>, Lena M. Krippner<sup>1,3</sup> & Nicolas W. Schuck<sup>1,2,\*</sup>

<sup>1</sup>Max Planck Research Group NeuroCode, Max Planck Institute for Human Development, Berlin, Germany

<sup>2</sup>Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Berlin, Germany  
Lentzeallee 94, D-14195 Berlin, Germany

<sup>3</sup>Harding Center for Risk Literacy, University of Potsdam, Faculty of Health Sciences, Potsdam, Germany  
Virchowstraße 2-4, D-14482 Potsdam, Germany

\*Correspondence to

[wittkuhn@mpib-berlin.mpg.de](mailto:wittkuhn@mpib-berlin.mpg.de) (ORCID: 0000-0003-2966-6888)

[schuck@mpib-berlin.mpg.de](mailto:schuck@mpib-berlin.mpg.de) (ORCID: 0000-0002-0150-8776)

## Abstract

Humans automatically infer higher-order relationships between events in the environment from their statistical co-occurrence, often without conscious awareness. Neural replay of task representations, which has been described as sampling from a learned transition structure of the environment, is a candidate mechanism by which the brain could use or even learn such relational information in the service of adaptive behavior. Human participants viewed sequences of images that followed probabilistic transitions determined by ring-like graph structures. Behavioral modeling revealed that participants acquired multi-step transition knowledge through gradual updating of an internal successor representation (SR) model, although half of participants did not indicate any knowledge about the sequential task structure. To investigate neural replay, we analyzed dynamics of multivariate functional magnetic resonance imaging (fMRI) patterns during short pauses from the ongoing statistical learning task. Evidence for sequential replay consistent with the probabilistic task structure was found in occipito-temporal and sensorimotor cortices during short on-task intervals. These findings indicate that implicit learning of higher-order relationships establishes an internal SR-based map of the task, and is accompanied by cortical on-task replay.

## 27 Introduction

28 The representation of structural knowledge in the brain in form of a so-called *cognitive map* has been  
29 a topic of great interest. A common assumption is that a cognitive map provides the basis for flexible  
30 learning, inference, and generalization (Tolman, 1948; Wilson et al., 2014; Schuck et al., 2016; Behrens  
31 et al., 2018), and yet is based on individual experiences that provide structural information only  
32 indirectly (Schapiro et al., 2013; Garvert et al., 2017). The brain must therefore extract statistical  
33 regularities from continuous experiences, and then use these regularities as the starting point for the  
34 formation of abstract, map-like knowledge. A mechanism through which abstract knowledge could  
35 be used to generate flexible behavior is on-task replay (e.g., Sutton, 1991; Kurth-Nelson et al., 2016),  
36 the rapid reactivation of trajectories simulated from an internal cognitive map. In this paper, we  
37 investigated whether on-task replay of cognitive map-like knowledge occurs in the human brain while  
38 participants learn statistical regularities.

39 The extraction of statistical regularities from experience is known as *statistical learning* (Schapiro  
40 and Turk-Browne, 2015; Garvert et al., 2017; Sherman et al., 2020). Statistical learning is automatic  
41 and incidental, as it occurs without any instructions or premeditated intention to learn, and often leads  
42 to implicit knowledge that is not consciously accessible (Reber, 1989; Seger, 1994; Turk-Browne et al.,  
43 2005). This contrasts with research on cognitive maps and planning that often relies on instruction-  
44 based task knowledge (e.g., Schuck et al., 2016; Constantinescu et al., 2016; Kurth-Nelson et al.,  
45 2016). In a statistical learning setting, relationships between events are typically described by pairwise  
46 transition probabilities (i.e., the probability that  $A$  is followed by  $B$ ) to which humans show great  
47 sensitivity from an early age on (Saffran et al., 1996). Intriguingly, many experiments have shown that  
48 humans extract higher-order relational structures among individual events that go beyond pairwise  
49 transition probabilities (for reviews, see e.g., Karuza et al., 2016; Lynn and Bassett, 2020). This  
50 includes knowledge about ordinal and hierarchical information that structures individual subsequences  
51 (Schuck et al., 2012a,b; Solway et al., 2014; Balaguer et al., 2016), graph topological aspects such as  
52 bottlenecks and community structure (Schapiro et al., 2013; Karuza et al., 2017; Kahn et al., 2018),  
53 and macro-scale aspects of graph structures (Lynn et al., 2020a,b).

54 A main benefit of abstracted knowledge in the context of transition structures is that it allows to  
55 plan multi-step sequences (Miller and Venditto, 2021; Hunt et al., 2021). Specifically, while experienced  
56 transition structure can be used to learn about the probability that a given event will be followed by a  
57 specific other event, it can also be used to compute long-term visitation probabilities, i.e., which events  
58 can be expected over a given future horizon. This idea is formalized in the successor representation  
59 (SR) (Dayan, 1993), a predictive map that reflects the (discounted) expected visitations of future events  
60 (Garvert et al., 2017; Bellmund et al., 2020; Brunec and Momennejad, 2021; Russek et al., 2021), and  
61 can be learned from the experience of individual transitions. Critically, the predictive horizon of the  
62 SR depends on a discount parameter  $\gamma$  which determines how far into the future upcoming states are  
63 considered (Momennejad and Howard, 2018; Momennejad, 2020). One goal of our study was therefore  
64 to investigate whether statistical learning leads to knowledge of expected future visitations over a  
65 predictive horizon, as required for mental planning.

66 The second main interest of our study was to understand whether abstract knowledge derived from  
67 statistical learning would be reflected in on-task replay. Replay is characterized by the fast sequential  
68 reactivation of neural representations that reflect previously experienced transition structure (see e.g.,  
69 Wikenheiser and Redish, 2015a; Schuck and Niv, 2019; Wittkuhn et al., 2021; Yu et al., 2021). Replay  
70 occurs in hippocampal but also cortical brain areas (Ji and Wilson, 2006; Wittkuhn and Schuck, 2021)

71 and has been observed during short pauses from the ongoing task in rodents (Johnson and Redish,  
72 2007; Carr et al., 2011) as well as humans (Kurth-Nelson et al., 2016; Tambini and Davachi, 2019).  
73 Sequential reactivation observed during brief pauses is often referred to as *online* or *on-task* replay,  
74 and likely reflects planning of upcoming choices (Kurth-Nelson et al., 2016; Eldar et al., 2020).

75 Previous studies have shown that expectations about upcoming visual stimuli elicit neural signals  
76 that are very similar to those during actual perception (Kok et al., 2012, 2014; Hindy et al., 2016;  
77 Kok and Turk-Browne, 2018) and anticipatory activation sequences have been found in visual cortex  
78 following perceptual sequence learning (Xu et al., 2012; Eagleman and Dragoi, 2012; Gavornik and  
79 Bear, 2014; Ekman et al., 2017). It remains unknown, however, whether on-task replay mirrors  
80 predictive knowledge that is stored in SR-based cognitive maps. In addition, while most research has  
81 focused on hippocampal reactivation, the above evidence suggests that statistical knowledge is also  
82 reflected in sensory and motor brain areas.

83 In the present study, we therefore examined whether on-task neural replay in visual and motor  
84 cortex reflects anticipation of sequentially structured stimuli in an automatic and incidental statisti-  
85 cal learning context. This may elucidate if (non-hippocampal) neural replay during on-task pauses  
86 contributes to learning of probabilistic cognitive maps. To this end, participants performed an in-  
87 cidental statistical learning paradigm (cf. Schapiro et al., 2012; Lynn et al., 2020a) in which visual  
88 presentation order and motor responses followed statistical regularities that were determined by a  
89 ring-like graph structure. The nature of the graph structure allowed us to dissociate knowledge about  
90 individual transition probabilities from an SR-based cognitive map that entails long-term visitation  
91 probabilities. Moreover, the transition probabilities among the task stimuli changed halfway through  
92 the experiment without prior announcement, which allowed us to understand the dynamical updating  
93 of task knowledge and replay within the same participants.

## 94 Results

95 Thirty-nine human participants took part in an fMRI experiment over two sessions. Participants  
96 were first informed that the experiment involves six images of animals (cf. Snodgrass and Vanderwart,  
97 1980; Rossion and Pourtois, 2004) and six response buttons mapped onto their index, middle, and  
98 ring fingers of both hands. Participants then began the first session of magnetic resonance imaging  
99 (MRI), during which they learned the stimulus-response (S-R) mappings between images and response  
100 buttons through feedback (*recall trials*, Fig. 1a, 8 runs with 60 trials each, 480 trials in total). In recall  
101 trials, animal images were shown without any particular sequential order, i.e., all pairwise sequential  
102 orderings of the images were presented equally often per run. Participants had to press the correct  
103 button in response to briefly presented images (500 milliseconds (ms)) during a response window (800  
104 ms; jittered stimulus-response interval (SRI) of 2500 ms on average). If the response was incorrect,  
105 a feedback about the correct button was provided (500 ms; no feedback on correct trials). The trial  
106 ended with a jittered inter-trial interval (ITI) of 2500 ms on average.

107 The second session started with one additional run of recall trials that was followed by five runs  
108 of *graph trials* (Fig. 1b, 240 trials per run, 1200 trials in total). As before, participants had to press  
109 the correct button in response to each animal. Images were now presented in a faster pace (800 ms  
110 per image and 750 ms between images on average), and only on 10% of trials (120 graph trials in  
111 total per participant), ITIs were set to 10 seconds (s). Importantly, the order of the images now  
112 followed a probabilistic transition structure (see below), about which participants were not informed,



113 and no feedback was provided. At the end of the second session, participants completed a post-task  
 114 questionnaire assessing explicit sequence knowledge.

115 The sequential ordering of images during graph trials was determined by either a *unidirectional*  
 116 or *bidirectional* ring-like graph structure with probabilistic transitions (Fig. 2a–b; for details, see  
 117 Methods). In the unidirectional graph condition (Fig. 2a, middle, henceforth *uni*), each image had  
 118 one frequent transition to the clockwise neighboring node (probability of  $p_{ij} = 0.7$ ), never transitioned  
 119 to the counterclockwise neighbor ( $p_{ij} = 0.0$ ), and was followed occasionally by the three other nodes  
 120 ( $p_{ij} = 0.1$  each; Fig. 2b, left). In consequence, stimuli most commonly transitioned in clockwise order  
 121 along the ring shown in Fig. 2a. In the bidirectional graph condition (Fig. 2a, right, henceforth *bi*),  
 122 transitions to both neighboring nodes (clockwise and counterclockwise) were equally likely ( $p_{ij} = 0.35$ ),  
 123 and transitions to all other three nodes occurred with  $p_{ij} = 0.1$  (Fig. 2b, right), as in the unidirectional  
 124 graph. Every participant started the task in one of these conditions (*uni* or *bi*). Halfway through  
 125 the third run, transitions began to be governed by the alternative graph, such that all participants  
 126 experienced both graphs as well as the change between them (Fig. 2c). 12 participants started in the  
 127 unidirectional condition and transitioned to the bidirectional graph (*uni* – *bi*), while 27 participants  
 128 experienced the reverse order (*bi* – *uni*).

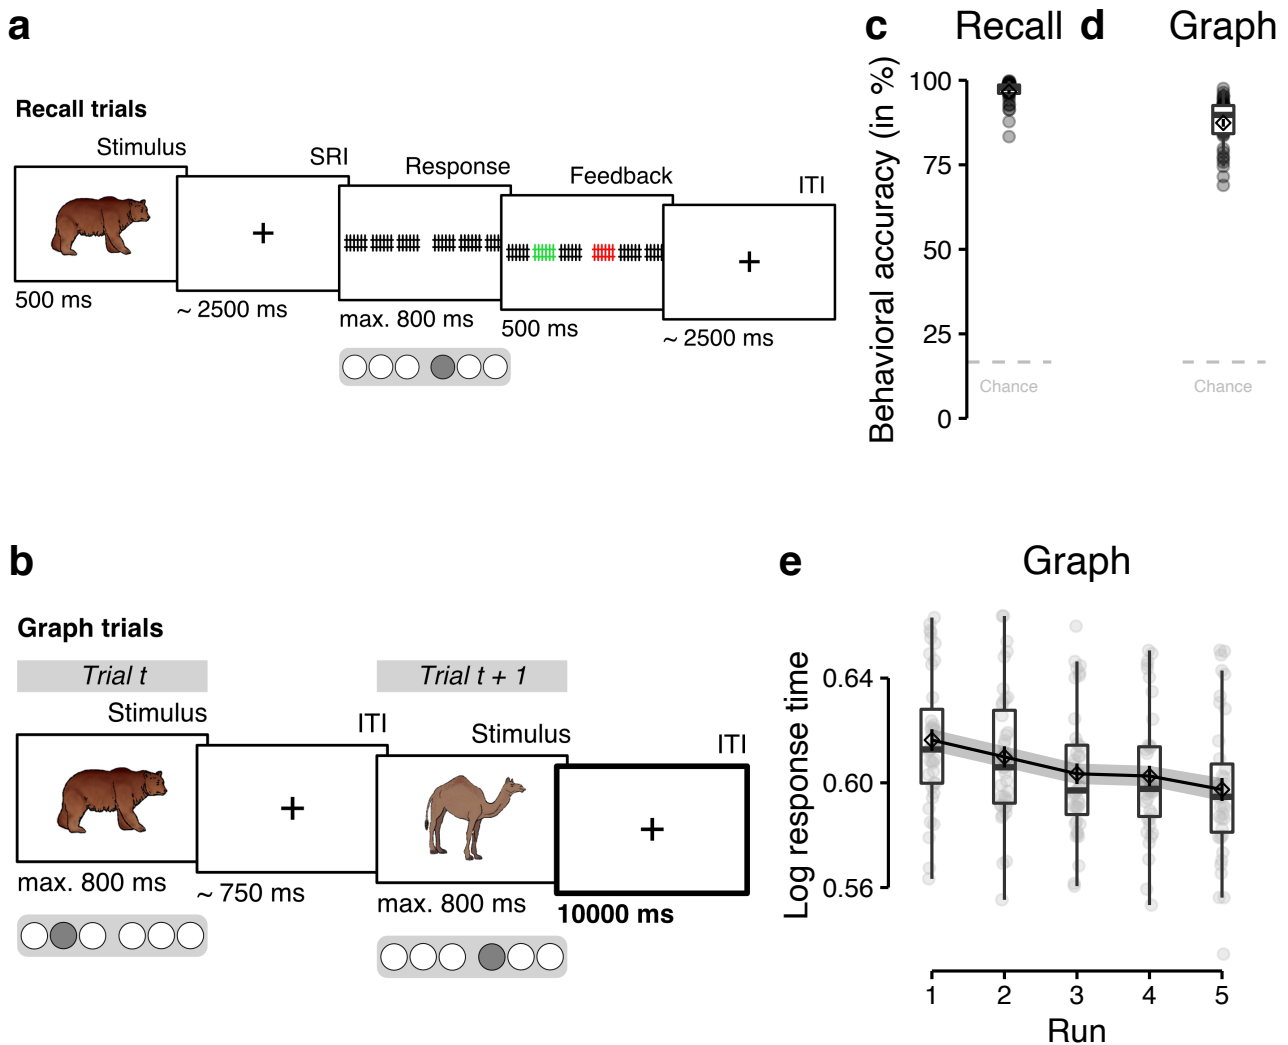


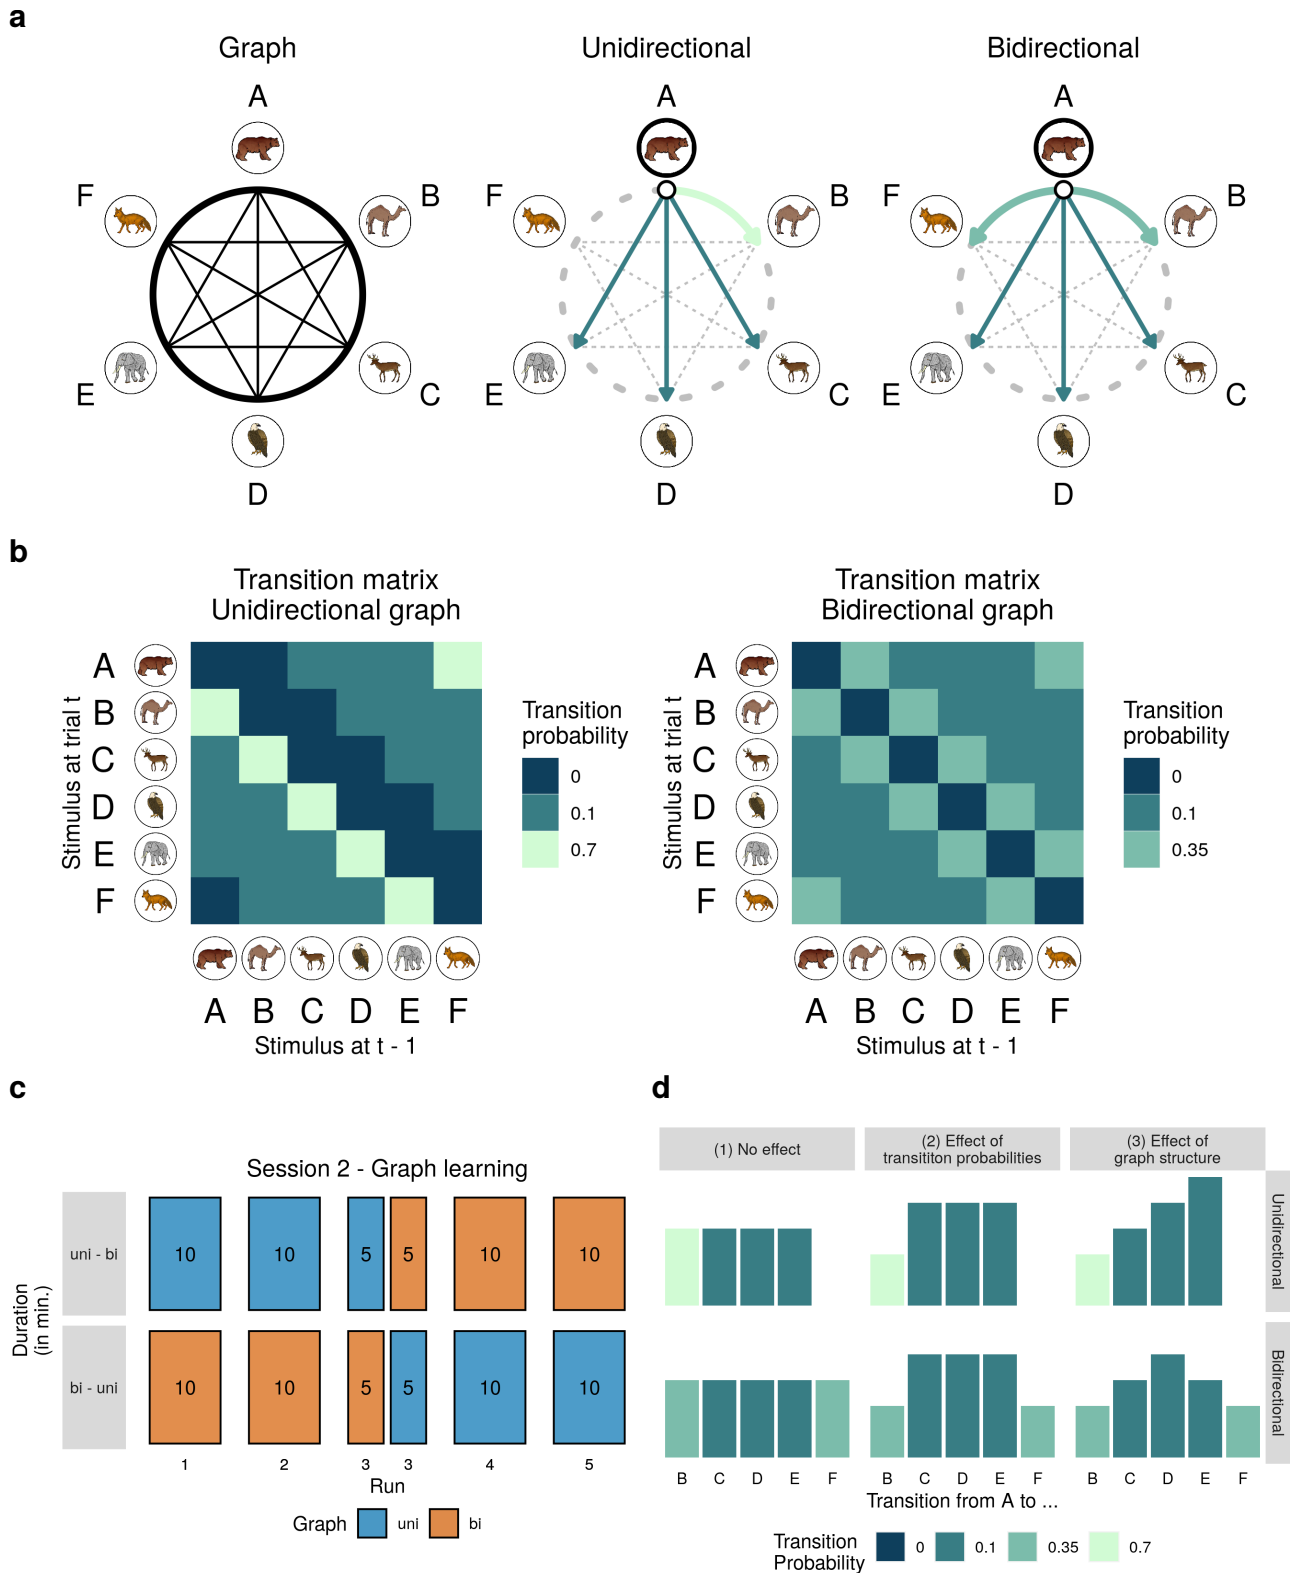
Figure 1: [see caption on the next page]

**Figure 1: Task design and stimulus-response learning.** (a) On recall trials, individual images were presented for 500 ms. Participants were instructed to press the correct response button associated with the stimulus during the response interval (time limit of 800 ms). Stimulus presentations and motor responses were separated by SRIs and ITIs which lasted 2.5 s on average (cf. Wittkuhn and Schuck, 2021). Feedback was only presented on incorrect trials. Classifiers were trained on fMRI data from correct recall trials only. (b) On graph trials, images were presented for 800 ms, separated by only 750 ms on average. Participants were asked to press the correct response button associated with the presented stimulus as quickly and accurately as possible within 800 ms. On 10% of trials, ITIs lasted 10 s (see ITI in trial  $t + 1$ ; highlighted by the thick border, for illustrative purposes only). Classifier trained on fMRI data from correct recall trials were applied to the eight TRs of the 10 s ITIs in graph trials to investigate task-related neural activation patterns during on-task pauses. (c) Mean behavioral accuracy (in %; y-axis) across all nine runs of the recall trials. (d) Mean behavioral accuracy (in %; y-axis) across all five runs of the graph trials. (e) Mean log response time (y-axis) per run (x-axis) in graph trials. Boxplots in (c), (d), and (e) indicate the median and interquartile range (IQR). The lower and upper hinges correspond to the first and third quartiles (the 25<sup>th</sup> and 75<sup>th</sup> percentiles). The upper whisker extends from the hinge to the largest value no further than 1.5\* IQR from the hinge (where IQR is the interquartile range (IQR), or distance between the first and third quartiles). The lower whisker extends from the hinge to the smallest value at most 1.5\* IQR of the hinge. The diamond shapes show the sample mean. Error bars in (c), (d) and shaded areas in (e) indicate  $\pm 1$  standard error of the mean (SEM). Each dot in (c), (d), and (e) corresponds to averaged data from one participant. All statistics have been derived from data of  $n = 39$  human participants who participated in one experiment. The stimulus material (individual images of a bear and a dromedary) shown in (a) and (b) were taken from a set of colored and shaded images commissioned by Rossion and Pourtois (2004), which are loosely based on images from the original Snodgrass and Vanderwart set (Snodgrass and Vanderwart, 1980). The images are freely available from the internet at <https://sites.google.com/andrew.cmu.edu/tarrlab/resources/tarrlab-stimuli> under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported license (CC BY-NC-SA 3.0; for details, see <https://creativecommons.org/licenses/by-nc-sa/3.0/>). Stimulus images courtesy of Michael J. Tarr, Carnegie Mellon University, (for details, see <http://www.tarrlab.org/>).

## 129 Behavioral results

130 We first asked whether participants learned the stimulus-response (S-R) mapping sufficiently well.  
131 Behavioral accuracy on recall trials indeed surpassed chance-level (16.67%) in all runs ( $\bar{x} \geq 86.50\%$ ,  
132 CIs  $[\geq 80.79, +\infty]$ ,  $t_{38} \geq 20.62$ ,  $ps < 0.001$  (corrected),  $ds \geq 3.30$ ; Figs. 1c, S2b-c). Likewise, during  
133 graph trials, participants also performed above chance in all runs ( $\bar{x} \geq 85.12$ , CIs  $[\geq 82.55, +\infty]$ ,  
134  $t_{38} \geq 44.90$ ,  $ps < 0.001$  (corrected),  $ds \geq 7.19$ ; Figs. 1d, S2d), and improved with time (effect of run:  
135  $F_{1.00,38.00} = 7.96$ ,  $p = 0.008$ , Fig. S2d).

136 Next, we investigated sequential knowledge. Although participants were not informed that images  
137 followed a sequential structure during graph trials, we expected that incidental learning would allow  
138 them to anticipate upcoming stimuli during these trials, and thus respond faster with learning. A linear  
139 mixed effects (LME) model that tested the effect of task run on response times was broadly in line  
140 with this assumption as it showed a significant decrease of response times over the course of learning,  
141  $F_{1.00,38.00} = 25.86$ ,  $p < 0.001$  (Figs. 1e, S2e). More directly, we expected that participants would learn  
142 the probabilistic transition structure of images and response buttons during graph trials, including  
143 the change in transition structure in the middle of the third run. Specifically, we hypothesized that  
144 participants would not only learn about one-step transition probabilities, but also form internal maps  
145 of the underlying graphs that reflect the higher-order structure of statistical multi-step relationships  
146 between stimuli, i.e., how likely a particular stimulus will be experienced in two, three, or more steps  
147 from the current time point (cf. Lynn and Bassett, 2020; Lynn et al., 2020a). In our task, this  
148 meant that participants might react differently to the three transitions that all have the same one-  
149 step transition probability, since they differ in how likely they would occur in multi-step trajectories.  
150 For instance, the one-step transition probabilities for A→C, A→D, and A→E were the same in the  
151 unidirectional graph, but the two-step probability of A→C was higher than for the other transitions,  
152 since the most likely two-step path was A→B→C. This means that participants should react faster



**Figure 2:** [see caption on the next page]

153 to  $A \rightarrow C$  transitions if they have multi-step knowledge. For simplicity, we will henceforth refer to the  
 154  $A \rightarrow C$  transition as having a shorter “node distance”, than  $A \rightarrow D$  or  $A \rightarrow E$  (see the rightmost column  
 155 in Fig. 2d, where colors reflect one-step transition probabilities, and the height of the bars indicate  
 156 node distance).

**Figure 2: Graph learning task.** (a) The relationships among the six task stimuli depicted as a ring-like graph structure (left). In the unidirectional graph (middle), stimuli frequently transitioned to the clockwise neighboring node ( $p_{ij} = p_{AB} = 0.7$ ), never to the counterclockwise neighboring node ( $p_{AF} = 0.0$ ), and only occasionally to the three other nodes ( $p_{AC} = p_{AD} = p_{AE} = 0.1$ ). In the bidirectional graph (right), stimuli were equally likely to transition to the clockwise or counterclockwise neighboring node ( $p_{AB} = p_{AF} = 0.35$ ) and only occasionally transitioned to the three other nodes ( $p_{AC} = p_{AD} = p_{AE} = 0.1$ ). Transition probabilities are highlighted for node *A* only, but apply equally to all other nodes. Arrows indicate possible transitions, colors indicate transition probabilities (for a legend, see panel b). (b) Transition matrices of the unidirectional (left) and bidirectional (right) graph structures. Each matrix depicts the probability (colors) of transitioning from the stimulus at the previous trial  $t - 1$  (x-axis) to the current stimulus at trial  $t$  (y-axis). (c) Within-participant order of the two graph structures across the five runs of the graph learning task.  $n = 12$  participants first experienced the unidirectional, then the bidirectional graph structure (uni - bi; top horizontal panel) while  $n = 27$  participants experienced the reverse order (bi - uni; bottom horizontal panel). In both groups of participants, the graph structure was changed without prior announcement halfway through the third task run. Numbers indicate approximate run duration in minutes (min). Colors indicate graph condition (uni vs. bi; see legend). (d) Visualization of the relative magnitude of the outcome variable (e.g., behavioral responses or classifier probabilities; y-axis) for specific transitions between the nodes (x-axis) and the two graph structures (uni vs. bi; horizontal panels) under the three assumptions (vertical panels), (1) that there is no difference between transitions (null hypothesis), (2) that response times are only influenced by the one-step transition probabilities between the nodes (colors), or (3) that response times are influenced by the multi-step relationships between nodes in the graph structure (here indicated by node distance). An effect of unidirectional graph structure would be evident in a linear relationship between node distance and the outcome variable, whereas a bidirectional graph structure would be reflected in a U-shaped relationship between node distance and independent measures (possibly inverted, depending on the measure). The stimulus material (individual images of a bear, a dromedary, a deer, an eagle, an elephant and a fox) shown in (a), and (b) were taken from a set of colored and shaded images commissioned by Rossion and Pourtois (2004), which are loosely based on images from the original Snodgrass and Vanderwart set (Snodgrass and Vanderwart, 1980). The images are freely available from the internet at <https://sites.google.com/andrew.cmu.edu/tarrlab/resources/tarrlab-stimuli> under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported license (CC BY-NC-SA 3.0; for details, see <https://creativecommons.org/licenses/by-nc-sa/3.0/>). Stimulus images courtesy of Michael J. Tarr, Carnegie Mellon University, (for details, see <http://www.tarrlab.org/>).

157 A first analysis revealed that participants reacted faster and more accurately to transitions with  
158 high compared to low one-step probabilities in the unidirectional graph condition ( $p_{ij} = 0.7$  versus  
159  $p_{ij} = 0.1$  transition probabilities,  $ps < 0.001$ ), and in the bidirectional graph condition ( $p_{ij} = 0.35$   
160 versus  $p_{ij} = 0.1$ ,  $ps < 0.001$ , Fig. 3a–b). In order to investigate whether multi-step transition  
161 probabilities also influenced participants’ behavior, we then analyzed response times and error rates  
162 as a function of the node distance (Fig. 2d; for details, see Methods). Using this analysis approach, we  
163 found a significant effect of node distance on response times in both unidirectional,  $F_{1.00,115.78} = 44.34$ ,  
164  $p < 0.001$ , and bidirectional data,  $F_{1.00,38.00} = 57.36$ ,  $p < 0.001$  (Fig. 3c). To further disentangle the  
165 effects of one-step and multi-step knowledge, we excluded data of frequent transitions ( $p_{ij} = 0.7$  and  
166  $p_{ij} = 0.35$  in the uni and bi conditions, respectively). In this case, the effect of node distance on  
167 response times in the unidirectional condition disappeared,  $F_{1.00,72.32} = 0.43$ ,  $p = 0.51$ , but persisted  
168 in bidirectional data,  $F_{1.00,76.98} = 5.52$ ,  $p = 0.02$  (Fig. 3c). No effects on behavioral accuracy were  
169 observed in either of the above analyses (all  $ps > 0.11$ ).

170 While these results offer a first indication of incidental learning of multi-step transitions, node  
171 distance is only an approximate reflection of the graph structure. A more precise way to express  
172 multi-step knowledge is to consider the discounted sum of different  $n$ -step probabilities as experienced  
173 by participants. This is equivalent to successor representation (SR) models (Dayan, 1993), which  
174 assume a representation of each node that reflects the discounted long-term occupation probability  
175 of all other nodes starting from the current node. Notably, recent work has shown that SRs can  
176 be updated through replay, rather than through online experience alone (Russek et al., 2017). We  
177 therefore investigated whether behavior reflected integrated mental SR-based maps of the experienced  
178 graph structure.

179 Specifically, for each node we modeled a vector that reflected the probability that starting from  
180 there a participant would experience any of the other nodes over a future-discounted predictive horizon.  
181 This vector was dynamically updated following the transitions that participants experienced in the  
182 task, using a temporal difference (TD) learning rule as used in SR models (Dayan, 1993; Russek et al.,  
183 2017). After experiencing the transition from image  $s_t$  to  $s_{t+1}$ , the row corresponding to image  $s_t$  of  
184 the successor matrix  $\mathbf{M}$  was updated as

$$\mathbf{M}_{s_t,*} = \mathbf{M}_{s_t,*} + \alpha [\mathbf{1}_{s_{t+1}} + \gamma \mathbf{M}_{s_{t+1},*} - \mathbf{M}_{s_t,*}] \quad (1)$$

185 whereby  $\mathbf{1}_{s_{t+1}}$  is a zero vector with a 1 in the  $s_{t+1}^{\text{th}}$  position, and  $\alpha$  is a learning rate. Crucially, the  
186 discounting parameter  $\gamma$  defined the extent to which multi-step transitions were taken into account,  
187 which we will henceforth refer to as the “predictive horizon” (cf. Gershman et al., 2012; Momennejad,  
188 2020). We computed a series of SR models with different predictive horizons between  $\gamma = 0$  (no  
189 predictive horizon) and  $\gamma = 0.95$  (in steps of 0.05), and asked how well response times could be  
190 predicted from these individually calculated, time-varying SRs (for details, see Methods). We then  
191 compared different LME models of response time, with a Shannon surprise predictor (cf. Shannon,  
192 1948) derived from each participants’ SR model, in addition to fixed effects of task run, graph (uni  
193 vs. bi) and graph order (uni – bi vs. bi – uni) as well as by-participant random intercepts and slopes.  
194 Comparing LME models that contained predictors from SR models with varying predictive horizons  
195 (i.e., levels of  $\gamma$ ) showed that a discount parameter of  $\gamma = 0.3$  resulted in the lowest Akaike information  
196 criterion (AIC) score (Fig. 3d), and models with non-zero  $\gamma$  parameters yielded substantially better  
197 fits than a model which assumed only knowledge of one-step transitions ( $\gamma = 0$ , leftmost data point in  
198 Fig. 3d). Thus, participants’ response times clearly indicated multi-step graph knowledge consistent  
199 with SR models.

200 To investigate if these analyses would differ between the two graph structures (uni vs. bi) and the  
201 two graph orders (uni – bi vs. bi – uni), we split the data according to these two factors and repeated  
202 a similar analysis of LME models (for details, see Methods). These analyses again showed that models  
203 based on a non-zero  $\gamma$  parameter achieved better fits, confirming that participants learned higher-order  
204 relationships among the nodes in the graph structure from experiencing sequences of transitions in the  
205 task (Fig. 3e). Interestingly, data from the first graph structure were fit best by the same  $\gamma$  parameter  
206 ( $\gamma = 0.55$ ), irrespective of graph condition (uni vs. bi; Fig. 3e, left panel column). When considering  
207 data from the second graph structure, in contrast, the depth of integration differed markedly depending  
208 on whether participants learned the uni- or bidirectional graph structure: participants who transitioned  
209 from the uni- to the bidirectional graph condition had a larger predictive horizon ( $\gamma = 0.75$ ; Fig. 3e,  
210 top right panel) in the second graph learning phase compared to participants who transitioned from  
211 a bi- to a unidirectional graph ( $\gamma = 0.3$ ; Fig. 3e, bottom right panel). These results indicated that  
212 the order in which graphs were experienced determined the depth of integration when learning was  
213 updated following a change in transition probabilities.

214 Finally, we assessed whether participants were able to express knowledge of the sequential ordering  
215 of stimuli and graph structures explicitly during a post-task questionnaire. Asked whether they had  
216 noticed any sequential ordering of the stimuli in the preceding graph task,  $n = 19$  participants replied  
217 “yes” and  $n = 20$  replied “no” (Fig. 3f). Of those participants who noticed sequential ordering  
218 ( $n = 19$ ), almost all (18 out of 19) indicated that they had noticed ordering within the first three runs  
219 of the task (Fig. 3g), and more than half of those participants (11 out of 19) indicated that they had  
220 noticed ordering during the third task run, i.e., the run during which the graph structure was changed.

221 Thus, sequential ordering of task stimuli remained at least partially implicit in half of the sample,  
222 and the change in the sequential order halfway through the third run of graph trials seemed to be one  
223 potential cause for the conscious realization of sequential structure. Participants were also asked to rate  
224 the transition probabilities of all pairwise sequential combinations of the six task stimuli (30 ratings in  
225 total). Interestingly, participants on average reported probability ratings that reflected bidirectional  
226 graph structure. Probabilities of transitions to clockwise and counterclockwise neighboring nodes were  
227 rated higher than rarer transitions to intermediate nodes, regardless of the order in which participants  
228 had experienced the two graph structures immediately before the questionnaire (Fig. 3h).

---

**Figure 3: Behavioral responses are modulated by transition probabilities and graph structure.** (a) Behavioral accuracy (y-axis) following transitions with low ( $p_{ij} = 0.1$ ) and high probability ( $p_{ij} = 0.7$  and  $p_{ij} = 0.35$  in the uni and bi conditions, respectively) for both graph structures (panels). Colors as in Fig. 2d. The horizontal dashed lines indicate the chance level (16.67%). (b) Log response time (y-axis) following transitions with low ( $p_{ij} = 0.1$ ) and high probability ( $p_{ij} = 0.7$  and  $p_{ij} = 0.35$  in the uni and bi conditions, respectively) for both graph structures (panels). Colors as in panel (a) and Fig. 2d. (c) Log response times (y-axis) as a function of uni- or bidirectional (u | b) node distance (x-axis) in data from the two graph structures (colors / panels). (d) AIC scores (y-axis) for LME models fit to participants' log response time data using Shannon surprise based on SRs with varying predictive horizons (the discounting parameter  $\gamma$ ; x-axis) as the predictor variable. (e) AIC scores (y-axis) for LME models fit to participants' log response time data using Shannon information based on SRs with varying predictive horizons (the discounting parameter  $\gamma$ ; x-axis) as the predictor variable, separated by graph order (uni – bi vs. bi – uni; horizontal panels) and graph condition (uni vs. bi; panel colors). (f) Number of participants (y-axis) indicating whether they had noticed any sequential ordering during the graph task (“yes” or “no”, x-axis). (g) Number of those participants (y-axis) who had detected sequential ordering indicating in which of the five runs of the graph task (x-axis) they had first noticed sequential ordering. (h) Ratings of pairwise transition probabilities (in %; y-axis) as a function of node distance / transition probability, separately for both graph orderings (uni – bi vs. bi – uni; panels). Boxplots in (a), (b), (c), and (h) indicate the median and IQR. The lower and upper hinges correspond to the first and third quartiles (the 25<sup>th</sup> and 75<sup>th</sup> percentiles). The upper whisker extends from the hinge to the largest value no further than  $1.5 \times$  IQR from the hinge (where IQR is the interquartile range (IQR), or distance between the first and third quartiles). The lower whisker extends from the hinge to the smallest value at most  $1.5 \times$  IQR of the hinge. The diamond shapes in (a), (b), (c), and (h) show the sample mean. Error bars and shaded areas in (a), (b), (c), and (h) indicate  $\pm 1$  SEM. Each dot in (a), (b), (c), and (h) corresponds to averaged data from one participant. Vertical lines in (d) and (e) mark the lowest AIC score. All statistics have been derived from data of  $n = 39$  human participants who participated in one experiment.

---

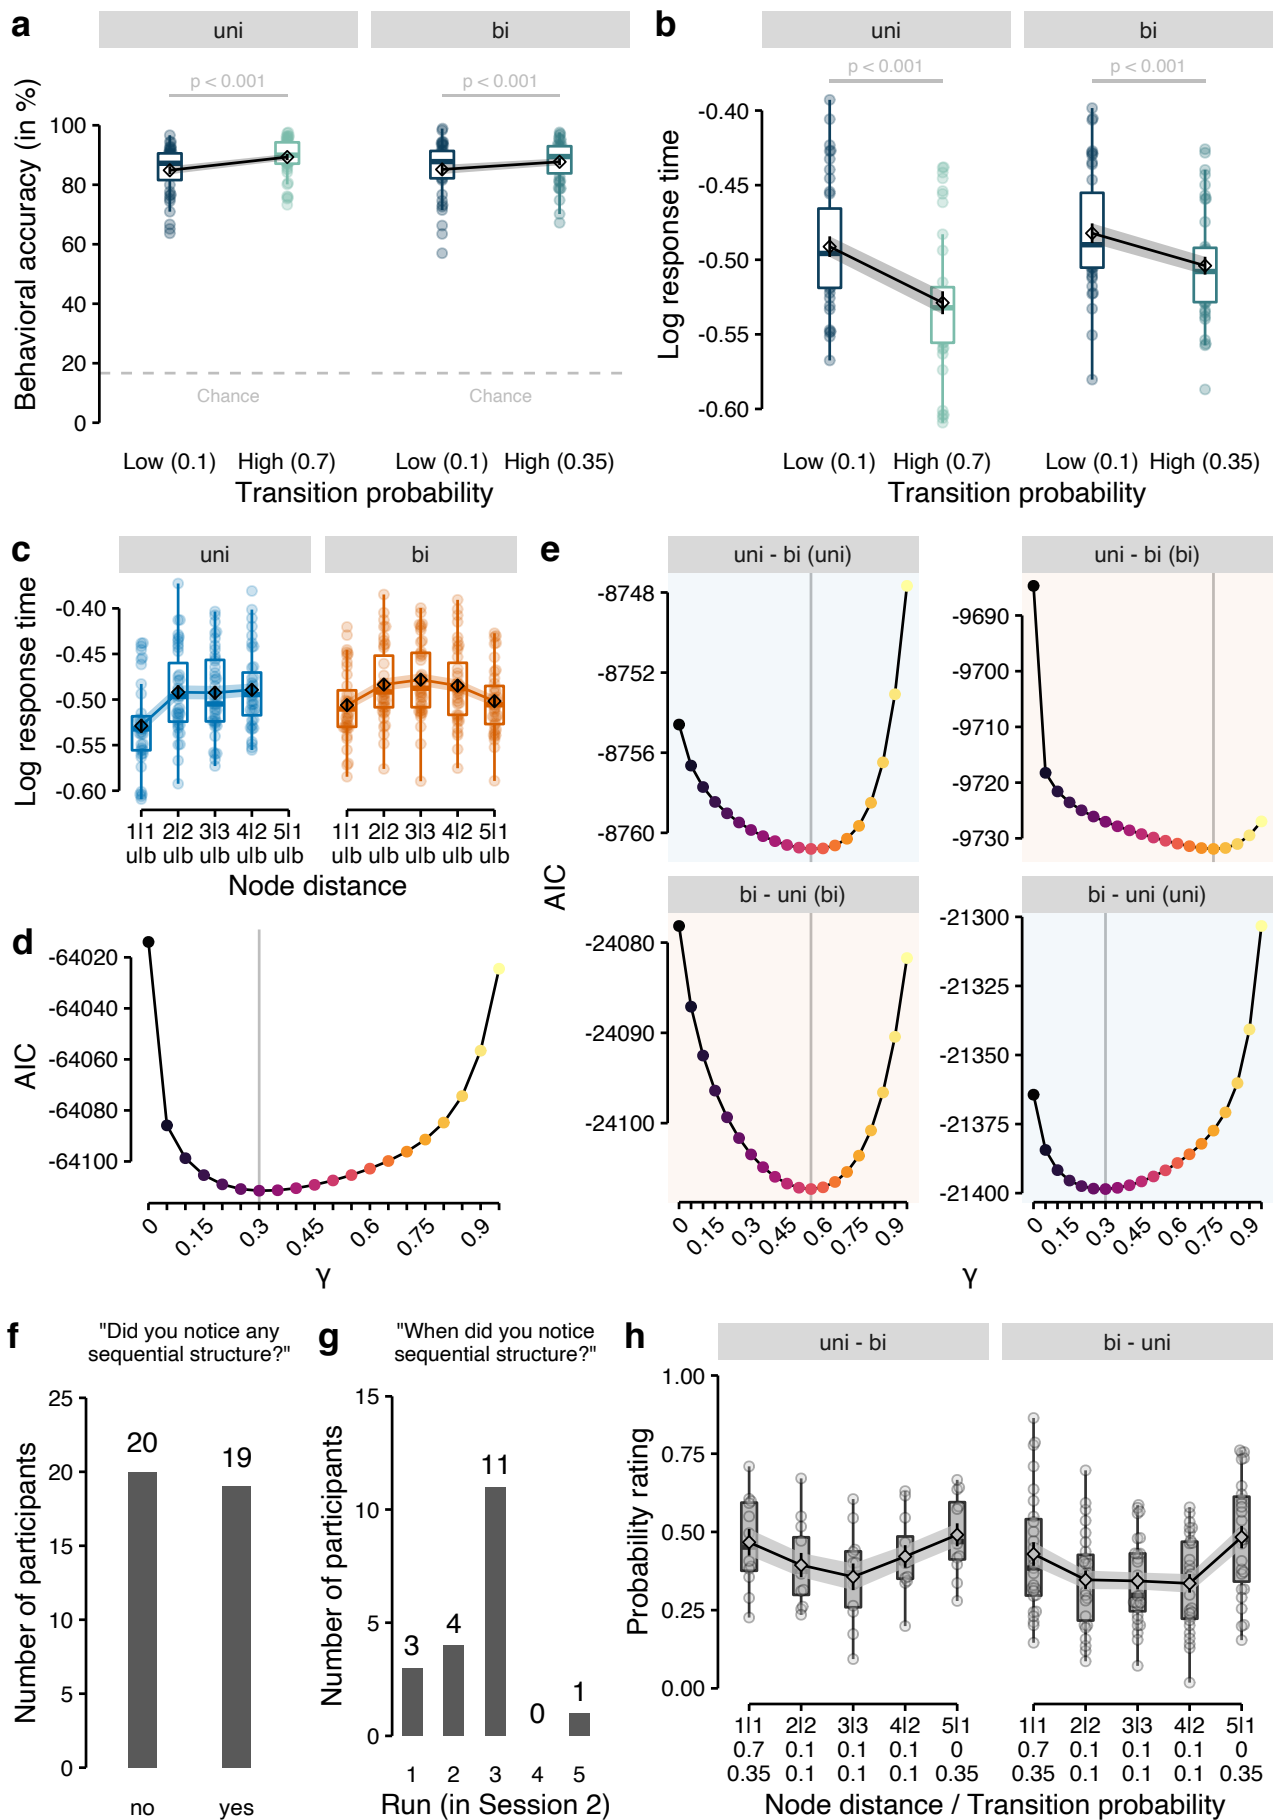


Figure 3: [see caption on the previous page]

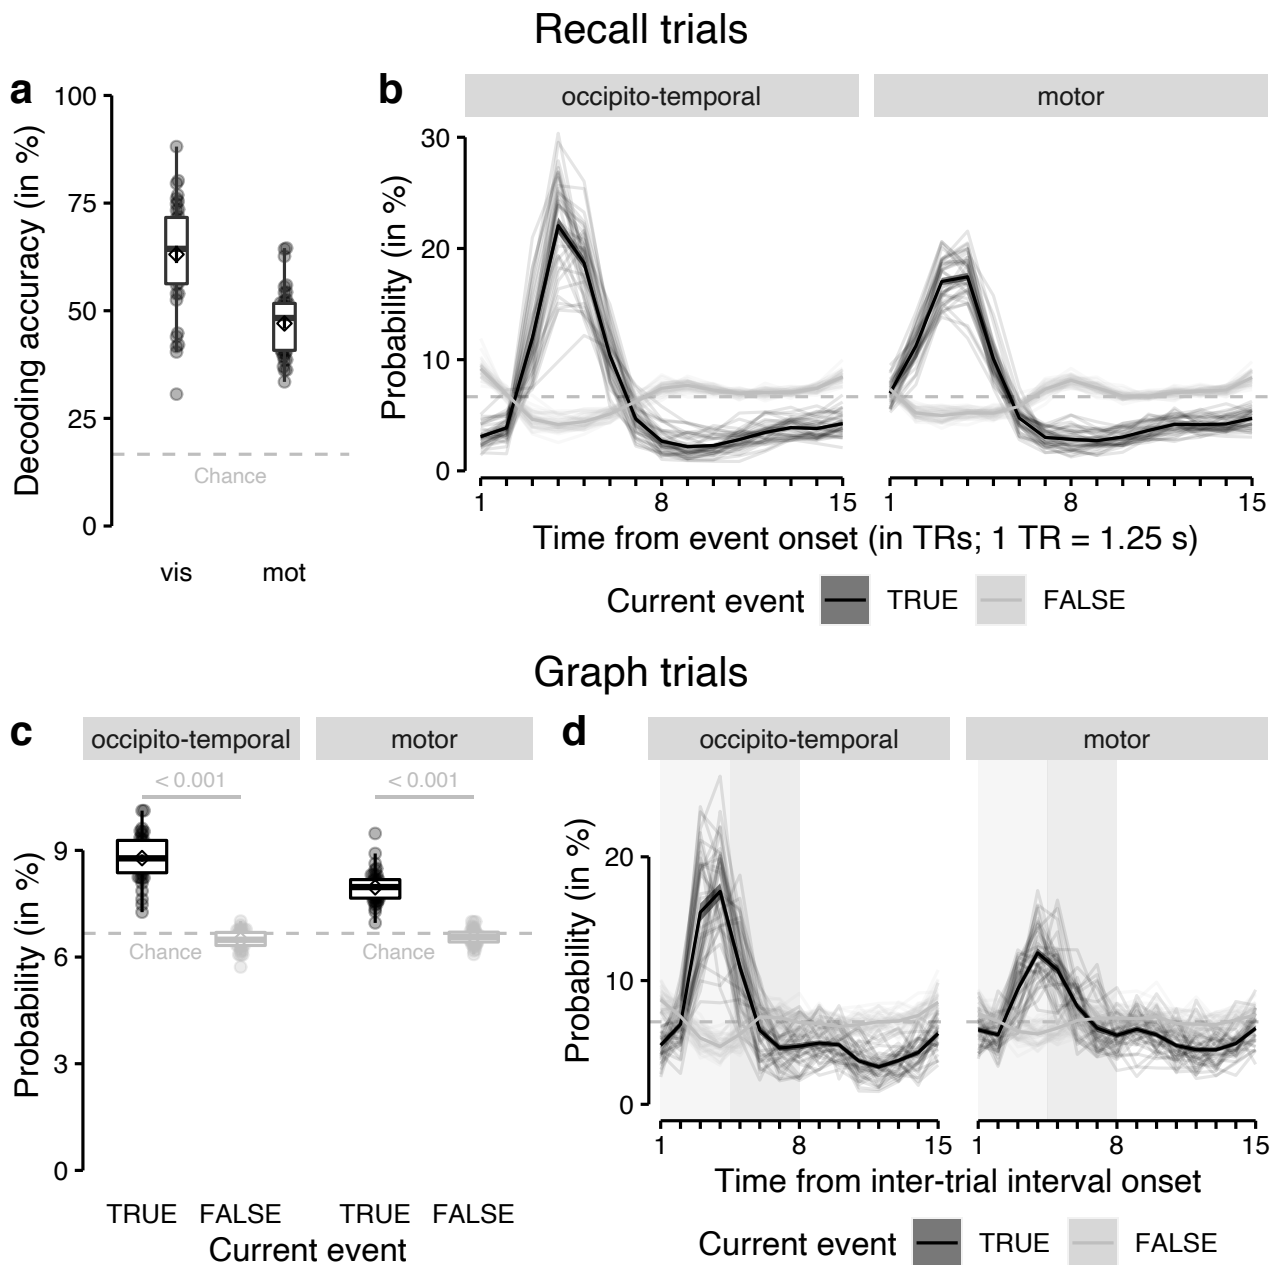
## 229 fMRI results

230 We next asked whether learning of map-like graph representations was accompanied by on-task replay.  
231 First, we trained logistic regression classifiers on fMRI signals related to stimulus and response onsets  
232 in correct recall trials (one-versus-rest training; for details, see Methods; cf. Wittkuhn and Schuck,  
233 2021). Separate classifiers were trained on data from gray-matter-restricted anatomical regions of  
234 interest (ROIs) of (a) occipito-temporal cortex and (b) pre- and postcentral gyri, which reflect visual  
235 object processing (cf. Haxby et al., 2001) and sensorimotor activity (e.g., Kolasinski et al., 2016),  
236 respectively. In each case, a single repetition time (TR) per trial corresponding either to the onset of  
237 the visual stimulus, or to participants' motor response was chosen (accounting for hemodynamic lag,  
238 time points were shifted by roughly 4 s; for details, see Methods). Note, that the order of displayed  
239 animals in recall trials was random, and image displays and motor responses were separated by SRIs  
240 and ITIs of 2500 ms to reduce temporal autocorrelation (cf. Dale, 1999; Wittkuhn and Schuck, 2021).

241 The trained classifiers successfully distinguished between the six animals. Leave-one-run-out clas-  
242 sification accuracy was  $M = 63.08\%$  in occipito-temporal data ( $SD = 12.57$ ,  $t_{38} = 23.06$ , CI [59.69,  
243  $+\infty$ ],  $p < 0.001$ , compared to a chance level of 16.67%,  $d = 3.69$ ) and  $M = 47.05\%$  in motor cortex  
244 data ( $SD = 7.79\%$ ,  $t_{38} = 24.36$ , CI [44.95,  $+\infty$ ],  $p < 0.001$ , compared to a chance level of 16.67%,  
245  $d = 3.90$ , all  $p$ -values Bonferroni-corrected, Fig. 4a). We also tested whether the classifiers successfully  
246 generalized from session 1 (eight recall runs) to session 2 (one recall run), and found no evidence for  
247 diminished cross-session decoding, compared to within-session,  $F_{8,00,655.00} = 0.95$ ,  $p = 0.48$  (for details  
248 see Methods). Next, we examined the sensitivity of the classifiers to pattern activation time courses by  
249 applying them to fifteen TRs following event onsets in recall trials (cf. Wittkuhn and Schuck, 2021).  
250 This analysis showed that the estimated normalized classification probability of the true stimulus class  
251 given the data peaked at the fourth TR as expected (Fig. 4b), where the probability of the true event  
252 was significantly higher than the mean probability of all other events at that time point (difference  
253 between current vs. other events; motor:  $M = 12.24$ ,  $t_{38} = 32.10$ , CI [11.47, 13.01],  $p < 0.001$ ,  
254  $d = 5.14$ ; occipito-temporal:  $M = 17.88$ ,  $t_{38} = 21.72$ , CI [16.22, 19.55],  $p < 0.001$ ,  $d = 3.48$ , all  
255  $p$ -values Bonferroni-corrected; Fig. 4b).

256 To address our main questions concerning on-task neural replay, we applied the classifiers to data  
257 from the graph trials that included 10 s on-task intervals (ITIs) with only a fixation on screen (120 trials  
258 per participant in total; 24 trials per run; 4 trials per stimulus per run; 10 s correspond to 8 TRs). We  
259 expected that participants would replay anticipated upcoming events or recently experienced event  
260 sequences during these on-task intervals, and that such replay would be evident in the ordering of  
261 classification probabilities. Crucially, classifier probabilities should reflect participants' knowledge of  
262 one-step transitions, but also their map-like representations that enabled them to form multi-step  
263 expectations, as described above. For example, in unidirectional graph trials image  $A$  was followed  
264 by image  $B$  with a higher probability than the other images. Therefore, the probability of decoding  
265 image  $B$  during an on-task interval following image  $A$  should be higher than the classifier probabilities  
266 of the other four possible next images (see Fig. 2a). In addition, although images  $C$ ,  $D$ , and  $E$   
267 had equal one-step transition probabilities, we expected the corresponding classifier probabilities to  
268 be ordered such as to reflect the multi-step SR-model described above. Following our previous work  
269 (Wittkuhn and Schuck, 2021), we also assumed that the ordering during the earlier phase of the on-  
270 task interval (TRs 1–4) would reflect the true directionality of the replayed sequence and would be  
271 reversed in the later phase of the interval (TRs 5–8), reflecting the rising and falling slopes of the  
272 underlying hemodynamic response functions (HRFs). As expected, the classifier probability of the





**Figure 4:** [see caption on the next page]

273 animal displayed in the current trial was higher compared to all other classes (Fig. 4c), and rising  
 274 and falling slowly as observed in recall trials (Fig. 4d, Fig. 5a; mean probability of current event vs.  
 275 all others;  $t_s \geq 17.88$ ,  $p_s < .001$ ,  $d_s \geq 3.48$ ,  $p$ -values Bonferroni-corrected). Because stimulus-evoked  
 276 activation was not of interest, we removed probabilities of the current stimulus from all following  
 277 analyses, considering only (normalized) probabilities from the five classes that did not occur on the  
 278 current trial.

279 To investigate replay of experienced or anticipated stimulus sequences, we modeled classifier prob-  
 280 abilities of non-displayed stimuli with LME models. LME models contained predictors that reflected  
 281 node distance, i.e., how likely each stimulus was to appear soon, given either a unidirectional (lin-  
 282 ear node distance) or bidirectional graph (quadratic node distance, see above). Because linear and  
 283 quadratic predictors were collinear, corresponding LME models were run separately. Each model  
 284 included fixed effects of ROIs (occipito-temporal vs. sensorimotor) and ITI phase (early vs. late).

**Figure 4: Classification accuracy and probabilistic classifier time courses on recall and graph trials.** (a) Cross-validated classification accuracy (in %) in decoding the six unique visual objects in occipito-temporal data (“vis”) and six unique motor responses in sensorimotor cortex data (“mot”) during task performance. Chance level is at 16.67% (horizontal dashed line). (b) Time courses (in TRs from stimulus onset; x-axis) of probabilistic classification evidence (in %; y-axis) for the event on the current recall trial (black) compared to all other events (gray), separately for both ROIs (panels). (c) Mean classifier probability (in %; y-axis) for the event that occurred on the current graph trial (black color), shortly before the onset of the on-task interval, compared to all other events (gray color), averaged across all TRs in the on-task interval, separately for each ROI (panels). (d) Time courses (in TRs from on-task interval onset; x-axis) of mean probabilistic classification evidence (in %; y-axis) in graph trials for the event that occurred on the current trial (black) and all other events (gray). Each line in (b) and (c) represents one participant. Classifier probabilities in (b), (c), and (d) were normalized across 15 TRs. The chance level therefore is at  $100/15 = 6.67\%$  (horizontal dashed line). Gray rectangles in (d) indicate the on-task interval (TRs 1–8). The light and dark gray areas in (d) indicate early (TRs 1–4) and late (TRs 5–8) phases, respectively. Boxplots in (a) and (c) indicate the median and IQR. The lower and upper hinges correspond to the first and third quartiles (the 25<sup>th</sup> and 75<sup>th</sup> percentiles). The upper whisker extends from the hinge to the largest value no further than  $1.5 \times$  IQR from the hinge (where IQR is the interquartile range (IQR), or distance between the first and third quartiles). The lower whisker extends from the hinge to the smallest value at most  $1.5 \times$  IQR of the hinge. The diamond shapes in (a) and (c) show the sample mean. Error bars and shaded areas indicate  $\pm 1$  SEM. Each dot corresponds to averaged data from one participant. All statistics have been derived from data of  $n = 39$  human participants who participated in one experiment.

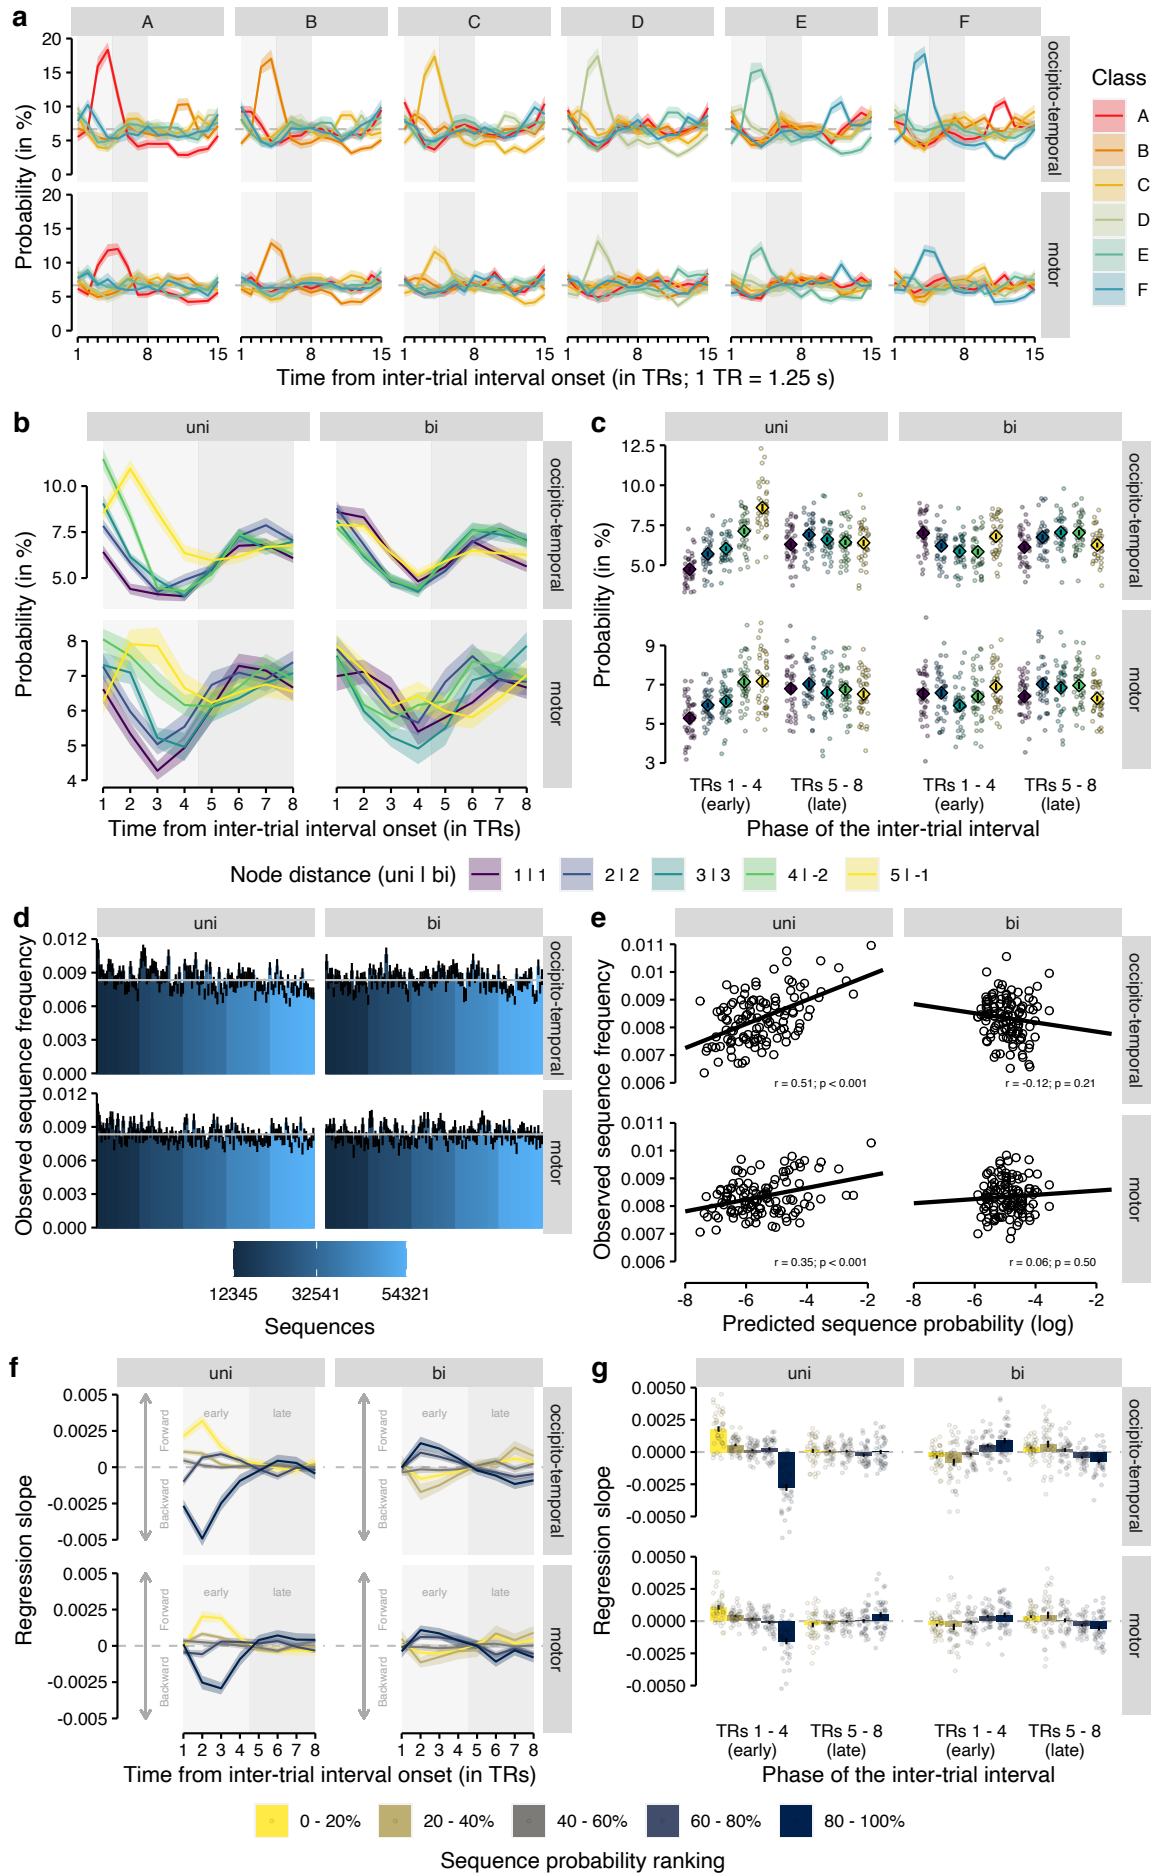
285 Considering data from runs in which stimulus transitions were governed by the unidirectional graph,  
286 an LME model containing the linear node distance predictor indicated a three-way interaction between  
287 node distance, ROI and phase  $F_{1.00,852.00} = 7.21$ ,  $p = 0.007$ . Post-hoc tests revealed an effect of node  
288 distance on classifier probabilities in unidirectional data in both ROIs in the early phase (TRs 1–4)  
289 of the ITIs,  $F_{1.00,810.00} \geq 78.18$ ,  $ps < 0.001$ , akin to backward replay of recently experienced stimuli.  
290 Effects in the late phase failed to reach significance (TRs 5–8),  $ps \leq 0.11$  (Fig. 5c). Considering  
291 data from the bidirectional run, we found a corresponding three-way interaction between bidirectional  
292 node distance, ROI and phase  $F_{1.00,852.00} = 5.59$ ,  $p = 0.02$ . Again, post-hoc tests revealed an effect  
293 of bidirectional node distance on classifier probabilities in both ROIs, showing a sign reversal when  
294 comparing the early to the late phase of the ITIs,  $F_{1.00,810.00} \geq 7.09$ ,  $ps \leq 0.008$  (Fig. 5c), in line  
295 with our expectations about on-task multi-step replay. Although linear and quadratic node distance  
296 predictors were collinear and therefore difficult to disentangle, we next tried to assess the specificity  
297 of the above effects by testing the linear (unidirectional) node distance on bidirectional data and the  
298 quadratic (bidirectional) node distance on unidirectional data. When a linear predictor was used  
299 in an LME model of bidirectional data, only a main effect of phase (early vs. late) was observed,  
300  $F_{1.00,852.00} = 11.55$ ,  $p < 0.001$ , but no main effect of the linear predictor,  $F_{1.00,852.00} = 0.27$ ,  $p = 0.60$ ,  
301 or any interactions among the predictor variables,  $ps \leq 0.09$ . Importantly, direct model comparison  
302 revealed that the linear model fit better in the unidirectional graph condition and the early phase  
303 of the ITI (see Fig. S6a–b). Using the quadratic predictor in the analysis of unidirectional data,  
304 we observed a three-way interaction between bidirectional node distance, the ROI, and the phase,  
305  $F_{1.00,852.00} = 4.35$ ,  $p = 0.04$ . Post-hoc tests revealed an effect of bidirectional node distance on classi-  
306 fier probabilities in unidirectional data only in the occipito-temporal ROI and only in the early phase  
307 (TRs 1–4) of the ITIs,  $F_{1.00,810.00} \geq 5.56$ ,  $ps < 0.02$  (Fig. 5c). Yet, model comparison again showed  
308 that the the quadratic model fit better in the bidirectional graph condition in both TR phases (dif-  
309 ferences in AICs were between  $-31.02$  and  $162.03$ , see Fig. S6a–b). Hence, these analyses confirmed  
310 that the observed classifier ordering was specific to the currently experienced graph.

311 The above analysis assumed that replayed sequences would always follow the most likely transitions  
312 (assuming a fixed ordering of replay sequences according to the multi-step graph structure). Yet, replay

313 might correspond more closely to a mental simulation of several possible sequences that are generated  
314 from a mental model. Consistent with this idea, the distribution of the observed sequential orders  
315 of classifier probabilities indicated a wide variety of replayed sequences (Fig. 5d, distribution over  
316 the entire ITI of 8 TRs). We next quantified how likely each possible sequential ordering of 5-item  
317 sequences was, based on the transition probabilities estimated by the SR model described above ( $\gamma$   
318 was set to 0.3 in order to approximate to the mean level of planning depth we had estimated based on  
319 the behavioral data, see above). To model measurement noise in the observed relative to the predicted  
320 sequences, we employed a hidden markov model (HMM) with structured emission probabilities (for  
321 details, see Methods). This revealed that during the unidirectional runs, the frequency with which  
322 we observed a sequence in brain data during the on-task pauses, strongly related to the probability  
323 of that sequence given the unidirectional graph structure (occipito-temporal ROI:  $r = .51$ ,  $p < 0.001$ ;  
324 motor ROI:  $r = .35$ ,  $p < 0.001$ ; Fig. 5e). Unexpectedly, this was not the case for the bidirectional  
325 runs ( $p = 0.21$  and  $p = 0.50$ , respectively; Fig. 5e).

326 We then sought to characterize the time courses of evidence for replay of sequences most likely  
327 to occur when mentally simulating a given sequence in the two graph structures. To this end, we  
328 calculated TR-wise linear regression slopes between the classifier probabilities and the 24 most likely  
329 sequences (top 20% of the  $5! = 120$  possible permutations), which resulted in an average sequentiality  
330 metric for each TR, similar to our previous work (Wittkuhn and Schuck, 2021). This analysis revealed  
331 significant backward sequentiality in the earlier phase (TRs 1–4) of the ITIs based on data from the  
332 unidirectional graph structure in both ROIs specifically for those sequences that were most likely  
333 given the unidirectional graph structure,  $t_{38}$ 's  $\leq -7.51$ ,  $ps < 0.001$ ,  $p$ -values Bonferroni-corrected (80  
334 – 100%; Fig. 5e). We did not find evidence for sequentiality in the late phase of the interval (TRs  
335 5–8) for either ROI in the unidirectional condition ( $ps > 0.97$ ). These findings mirror the results from  
336 the analysis of classification probabilities (see above) in showing that classifier probabilities in earlier  
337 TRs of fMRI data with unidirectional graph structure are ordered backward relative to the sequential  
338 ordering implied by the graph structure. In the bidirectional condition, we found forward sequentiality  
339 in the earlier phase (TRs 1–4;  $t_{38}$ 's  $\geq 3.90$ ,  $ps < 0.02$ ,  $ds \geq 0.63$ ) of the ITI and backward sequentiality  
340 in the later phase (TRs 5–8;  $t_{38}$ 's  $\leq -4.31$ ,  $ps < 0.001$ ,  $ds \leq -0.69$ ), in occipito-temporal data for the  
341 top 40% most likely sequences (i.e., both 80–100% and 60–80%,  $p$ -values Bonferroni-corrected, Fig.  
342 5e). Again, these results were in line with the analyses of classification probabilities, that found an  
343 influence on bidirectional graph structure in both early and late TRs.

344 Together, these results provide evidence that classifier probabilities in ITIs of graph trials are  
345 modulated by the multi-step distances between nodes in the graph structure. These effects of multi-  
346 step distances are in line with the idea that participants replayed multi-step sequences during brief  
347 on-task pauses, which could provide the basis for participants' map-like knowledge of incidentally  
348 experienced graph structures. When transition probabilities among stimuli in the task followed a  
349 unidirectional graph structure, classifier probabilities are influenced by a linear ordering of nodes  
350 that scales with the distance among the nodes in a unidirectional ordering, albeit only in earlier  
351 TRs following ITI onset (Fig. 5). When classifier probabilities from trials of the bidirectional graph  
352 structure are considered, classifier probabilities are influenced by a quadratic relationship to node  
353 distance (modeling a bidirectional ordering of nodes), in both the early (TRs 1–4) and late (TRs 5–8)  
354 phases of the ITIs and in both ROIs (Fig. 5). The graph distance effect appeared more pronounced  
355 in earlier compared to later TRs, but was present in both occipito-temporal and motor ROIs and  
356 followed a similar dynamic with respect to early and late phases of the ITI in both ROIs.



**Figure 5:** [see caption on the next page]

---

**Figure 5: Classifier probabilities during inter-trial intervals (ITIs) of graph trials are modulated by node distances in the graph structure.** (a) Time courses (in TRs from ITI onset; x-axis) of mean probabilistic classification evidence (in %; y-axis) for each of the six classes (colors) depending on the event of the current trial (vertical panels) and the anatomical ROI (horizontal panels). The event of the current trial (stimulus presentation or motor response) happened a few hundred ms before the onset of the ITI (for the trial procedure of graph trials, see Fig. 1b). (b) Time courses (in TRs from ITI onset; x-axis) of mean probabilistic classification evidence (in %; y-axis) for each of the five classes that were not presented on the current trial, colored by node distance in the two graph structures (vertical panels) for both anatomical ROI (horizontal panels). (c) Mean probabilistic classification evidence (in %; y-axis) for each node distance (colors) in the unidirectional (left vertical panel) and bidirectional (right vertical panel) graph structures averaged across TRs in the early (TRs 1–4) or late (TRs 5–8) phase (x-axis) for data in the occipito-temporal (top horizontal panels) and motor (bottom horizontal panels) ROIs. (d) Relative frequencies (y-axis) of all 120 permutations of probability-ordered 5-item sequences within each TR observed during on-task intervals, separately for both graph structures (vertical panels) and anatomical ROIs (horizontal panels). The horizontal gray line indicates the expected frequency if all sequences would occur equally often ( $1/120 = 0.008$ ). Colors indicate sequence ordering from forward (e.g., 12345; dark blue) to backward (e.g., 54321; light blue) sequences. (e) Correlations (Pearson’s  $r$ ) between the predicted sequence probability and the observed sequence frequency (120 5-item sequences per correlation), separately for both graph structures (vertical panels) and anatomical ROIs (horizontal panels). Each dot represents one 5-item sequence. (f) Regression slopes (y-axis) relating classifier probabilities to sequential positions for both graph structures (vertical panels) and anatomical ROIs (horizontal panels). Sequential orderings were determined based on a hidden markov model (HMM) identifying the most likely sequences based on the two graph structures (colors). Positive and negative slopes indicate forward and backward sequentiality, respectively (cf. Wittkuhn and Schuck, 2021). (g) Mean classifier probabilities averaged across all TRs in the early and late phase (x-axis) of the ITIs, separately for both graph structures (vertical panels) and anatomical ROIs (horizontal panels). Each dot in (c) and (g) corresponds to averaged data from one participant. Error bars in (c), (d), and (g) and shaded areas in (a), (b), and (f) represent  $\pm 1$  SEM. Gray rectangles in (a), (b), and (d) indicate the on-task interval (TRs 1–8). The light and dark gray areas in (a), (b), and (f) indicate early (TRs 1–4) and late (TRs 5–8) interval phases, respectively. 1 TR in (a), (b), and (f) = 1.25 s. All statistics have been derived from data of  $n = 39$  human participants who participated in one experiment.

---

## 357 Discussion

358 We present results showing on-task cortical replay of future sequences simulated from a mental model  
359 of an experienced graph in humans. Replay was detected in visual and sensorimotor cortex while  
360 participants briefly paused during an incidental statistical learning task. Statistical regularities in our  
361 main task were governed by two graph structures, one of which determined transitions in the first half  
362 of the experiment, while the other one determined transitions in the second half. We demonstrate that  
363 participants' response times reflect continuous learning of future-discounted predictive expectations  
364 that go beyond knowledge of one-step transitions and are captured by temporal difference (TD)  
365 learning of a successor representation (SR) model (cf. [Dayan, 1993](#)). These behavioral effects are  
366 in line with our neural results which indicate on-task replay consistent with sampling from such an  
367 SR model. Participants did not receive explicit instructions to learn and about half of participants  
368 reported no explicit knowledge of the experienced sequentiality. Learning was therefore automatic and  
369 partially implicit.

370 Our behavioral results are consistent with previous findings showing that humans learn about  
371 networks of stimuli beyond one-step transitions (e.g., [Schapiro et al., 2013](#); [Karuza et al., 2016, 2017,](#)  
372 [2019](#); [Garvert et al., 2017](#); [Kahn et al., 2018](#); [Lynn and Bassett, 2020](#); [Lynn et al., 2020a,b](#)). Our  
373 computational modeling establishes a link between these behavioral effects and an online temporal  
374 difference (TD) learning mechanism that tracks the long-term visitation probabilities. Our findings  
375 add to a growing set of studies that uses models based on SRs ([Dayan, 1993](#)) to demonstrate the  
376 formation of predictive representations of task structure in human behavioral and neuroimaging data  
377 ([Garvert et al., 2017](#); [Russek et al., 2017](#); [Momennejad et al., 2017](#); [Momennejad, 2020](#); [Russek et al.,](#)  
378 [2021](#)). Through model comparisons between SR models that differed in their discounting parameter  
379  $\gamma$ , i.e., their predictive horizon, we found that behavior overall was best explained by a medium  
380 deep predictive horizon corresponding to  $\gamma = 0.3$  (note, that any model with  $\gamma > 0$  suggests that  
381 participants formed predictive representations). When we separated the analyses by graph condition  
382 and graph order, we found that during learning of the first graph structure, planning depth was  
383 deeper, as indicated by a predictive horizon of  $\gamma = 0.55$ , irrespective of whether transition structure  
384 was governed by the uni- or bidirectional graph condition. This finding suggests that, upon entering a  
385 novel environment with sequential events, humans might integrate multi-step transition probabilities  
386 to a medium depth that is independent from the specific structure of the environment. Interestingly,  
387 after the transition structure changed to the second graph structure halfway through the task, this  
388 also seemed to influence the predictive horizon in a manner that was dependent on the order in which  
389 the two graphs were experienced. In participants who first learned the unidirectional and then the  
390 bidirectional graph, the best fitting model was based on an SR with a higher discount parameter of  
391  $\gamma = 0.75$ . This may indicate a deeper integration of higher-order relationships in the bidirectional  
392 graph structure compared to the unidirectional graph structure. In contrast, in participants who  
393 experienced the reverse order, the best fitting model during the second half of the experiment was  
394 based on an SR with a lower discount parameter of  $\gamma = 0.3$ . This could indicate a reduced predictive  
395 horizon when learning relationships in the unidirectional graph. In sum, these results suggest that  
396 participants' predictive horizon interacts with the structure of the task as well as the learning history  
397 and indicates that the depth of integration could adapt to changes in the task environment. This  
398 idea relates to recent work suggesting that the brain may host SRs at varying predictive horizons in  
399 parallel ([Momennejad and Howard, 2018](#); [Brunec and Momennejad, 2021](#)).

400 Analyzing fMRI data recorded during 10 s pauses in-between performing the main task, we found  
401 evidence that classification probabilities were modulated by the transition probabilities and multi-  
402 step node distances within the two graph structures. Applying our previously developed sequentiality  
403 metric (Schuck and Niv, 2019; Wittkuhn and Schuck, 2021), we found evidence for backward sequen-  
404 tiality in unidirectional data and forward sequentiality in bidirectional data in both occipito-temporal  
405 and motor ROIs. The sequentiality metric was strongest specifically for those sequential orderings of  
406 classification probabilities that were most likely given an SR model of the two graph structures (Fig.  
407 5). Our evidence for on-task replay relates to research in rodents, where time-compressed sequential  
408 place cell activations, called theta sequences, occur during active behavior (Foster and Wilson, 2007)  
409 and reflect multiple potential future trajectories when the animal pauses at a decision point (Johnson  
410 and Redish, 2007), or cycle between future trajectories during movement (Kay et al., 2020) possibly  
411 reflecting an online planning process. Similar relationships between hippocampal theta and planning  
412 have been observed in human magnetoencephalography (MEG) experiments (Kaplan et al., 2020),  
413 which have also yielded evidence for on-task planning in the form of fast sequential neural reactiva-  
414 tion (Kurth-Nelson et al., 2016; Eldar et al., 2020). An fMRI study in humans has related on-task  
415 prospective neural activation to model-based decision-making (Doll et al., 2015), but the temporal  
416 dynamics of the prospective neural representations remained unclear. In contrast to previous studies,  
417 participants in our experiment did not engage in any explicit planning process. As mentioned before,  
418 participants were not instructed to learn about any sequentiality in the task. Moreover, participants  
419 were only told that short pauses may occur during the task, but they were not informed about the  
420 purpose of these pauses, and could not predict when the pauses would occur. It therefore seems likely  
421 that neural representations during on-task pauses reflect ongoing task representations similar to theta  
422 sequences in rodents.

423 One important aspect of our work is that we focused on cortical replay of predictive representations  
424 in visual (occipito-temporal) and sensorimotor (pre- and postcentral gyri) cortex. Previous work has  
425 largely focused on the hippocampus as a site of replay and as a potential brain region to host predictive  
426 cognitive maps (Garvert et al., 2017; Stachenfeld et al., 2017), while other studies have also emphasized  
427 the role of the prefrontal cortex (PFC) (Wilson et al., 2014; Schuck et al., 2016; Badre and Nee, 2018).  
428 Several fMRI studies demonstrated that hippocampal activity is modulated by stimulus predictability  
429 in sequential learning tasks (Strange et al., 2005; Harrison et al., 2006; Bornstein and Daw, 2012) and  
430 is related to the reinstatement of cortical task representations in visual cortex (Bosch et al., 2014;  
431 Hindy et al., 2016; Kok and Turk-Browne, 2018). Replay is known to occur throughout the brain (see  
432 e.g., Foster, 2017) but the functions of distributed replay events still remain to be further illuminated.  
433 Our findings shed light on the distribution of predictive representations and replay in the human brain,  
434 and suggest a potential involvement of sensory and motor areas. Yet, which roles the hippocampus  
435 and PFC play in this process remains an open question.

436 Our results suggest that participants formed a predominantly bidirectional representation of the  
437 ring-like graph structure, irrespective of the order in which the two graphs were experienced. The  
438 influence of node distance on response times was more pronounced and the predictive horizon in  
439 SR-based analyses was deeper in bidirectional compared to unidirectional behavioral data. Post-task  
440 ratings of transition probabilities were biased by bidirectional node distance, irrespective of graph  
441 order. The reversal in the directionality of classifier probabilities from early to late TRs, which is  
442 characteristic for sequential neural events in fMRI data (cf. Wittkuhn and Schuck, 2021), was only  
443 observed in on-task intervals during bidirectional but not unidirectional graph trials. This dominance

444 of a bidirectional representation could reflect that transitions in clockwise order in the unidirectional  
445 graph (e.g., from  $A$  to  $B$ ; Fig. 2) still allow to infer an associative relationship in the reverse direction  
446 (i.e., from  $B$  to  $A$ ), even though this transition actually never occurs during the task.

447 One remaining challenge for future research is to better understand the sequentiality of replay. We  
448 have previously shown that, at the level of classifier probabilities, sequences of neural events first elicit  
449 forward followed by backward sequentiality relative to the true sequence of events due to the dynamics  
450 of the HRF (Wittkuhn and Schuck, 2021). The fact that we found backward sequentiality in earlier  
451 TRs relative to an assumed sequential ordering of classifier probabilities in line with the unidirectional  
452 graph structure suggests that the true sequence of neural events at the start of the on-task intervals  
453 was indeed backwards. In the bidirectional graph structure, however, sequences can be expected in  
454 both directions, i.e.,  $A-B-C-D-E$  and  $E-D-C-B-A$  sequences are both very likely. It therefore remains  
455 unclear whether detecting a replayed sequence of  $A-B-C-D-E$  reflects forward replay of this sequence  
456 or backward replay of its reverse ( $E-D-C-B-A$ ). Previous research has found awake replay in both  
457 forward and backward order in rodents (Foster and Wilson, 2006; Diba and Buzsáki, 2007; Gupta  
458 et al., 2010) as well as in humans (Liu et al., 2021), and suggested that the directionality of replay  
459 may be tied to different functions, such as memory consolidation vs. value learning (e.g., Foster and  
460 Wilson, 2006; Ólafsdóttir et al., 2018; Liu et al., 2019; Wittkuhn et al., 2021). Neural sequences that  
461 have been associated with a prospective planning function are typically in forward order relative to the  
462 experienced sequence (Johnson and Redish, 2007; van der Meer and Redish, 2009; Pfeiffer and Foster,  
463 2013; Wikenheiser and Redish, 2015b). However, as others have pointed out before (Kurth-Nelson  
464 et al., 2016), it is plausible to plan backward instead of forward (also see LaValle, 2006), and previous  
465 studies also reported backward sequences during theta in rodents (Wang et al., 2020) as well as during  
466 value learning in humans (Liu et al., 2021).

467 Another challenge will be to better understand the relation between changes in neural representa-  
468 tions and replay. Repeated exposure to sequences of stimuli has been shown to increase the similarity  
469 of neural stimulus representations in the medial temporal lobe (MTL) in both macaques (Miyashita,  
470 1988) and humans (Schapiro et al., 2012). Using fMRI adaptation (cf. Barron et al., 2016), Garvert  
471 et al. (2017) showed that the similarity of neural representations of task stimuli decreases with distance  
472 between stimuli in a graph structure. This may pose a challenge to classifiers trained on individual  
473 stimulus presentations as in the current study, because increases in the similarity of neural represen-  
474 tations could increase the confusability of decoded patterns, which in turn may cause biases in the  
475 measured sequentiality.

476 In conclusion, our results provide insights into how the human brain forms predictive represen-  
477 tations of the structural relationships in the environment from continuous experience and samples  
478 sequences from these internal cognitive maps during on-task replay.



## 479 **Methods**

### 480 **Participants**

481 44 young and healthy adults were recruited from an internal participant database or through local  
482 advertisement and fully completed the experiment. No statistical methods were used to predetermine  
483 the sample size but it was chosen to be larger than similar previous neuroimaging studies (e.g., [Schuck](#)  
484 [and Niv, 2019](#); [Momennejad et al., 2018](#); [Tambini and Davachi, 2013](#)). Five participants were excluded  
485 from further analysis because they viewed different animals in session 1 and 2 due to a programming  
486 error in the behavioral task. Thus, the final sample consisted of 39 participants (mean age = 24.28  
487 years,  $SD = 4.24$  years, age range: 18 - 33 years, 23 female, 16 male). All participants were screened  
488 for MRI eligibility during a telephone screening prior to participation and again at the beginning  
489 of each study session according to standard MRI safety guidelines (e.g., asking for metal implants,  
490 claustrophobia, etc.). None of the participants reported to have any major physical or mental health  
491 problems. All participants were required to be right-handed, to have corrected-to-normal vision,  
492 and to speak German fluently. The ethics commission of the German Psychological Society (DGPs)  
493 approved the study protocol (reference number: SchuckNicolas2020-06-22VA). All volunteers gave  
494 written informed consent prior to the beginning of the experiments. Every participant received 70.00  
495 Euro and a performance-based bonus of up to 5.00 Euro upon completion of the study. None of the  
496 participants reported to have any prior experience with the stimuli or the behavioral task.

### 497 **Task**

### 498 **Stimuli**

499 All visual stimuli were taken from a set of colored and shaded images commissioned by [Rossion](#)  
500 [and Pourtois \(2004\)](#), which are loosely based on images from the original Snodgrass and Vanderwart  
501 set ([Snodgrass and Vanderwart, 1980](#)). The images are freely available on the internet at <https://sites.google.com/andrew.cmu.edu/tarrlab/resources/tarrlab-stimuli> under the terms of  
502 the Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported license (for details, see  
503 <https://creativecommons.org/licenses/by-nc-sa/3.0/>) and have been used in similar previous  
504 studies (e.g., [Garvert et al., 2017](#)). Stimulus images courtesy of Michael J. Tarr at Carnegie Mellon Uni-  
505 versity, (for details, see <http://www.tarrlab.org/>). In total, we selected 24 images which depicted  
506 animals that could be expected in a public zoo. Specifically, the images depicted a bear, a dromedary,  
507 a deer, an eagle, an elephant, a fox, a giraffe, a goat, a gorilla, a kangaroo, a leopard, a lion, an ostrich,  
508 a peacock, a penguin, a raccoon, a rhinoceros, a seal, a skunk, a swan, a tiger, a turtle, and a  
509 zebra (in alphabetical order). For each participant, six task stimuli were randomly selected from the  
510 set of 24 the animal images and each image was randomly assigned to one of six response buttons. This  
511 randomization ensured that any potential systematic differences between the stimuli (e.g., familiarity,  
512 preference, or ability to decode) would not influence the results on a group level (for a similar reasoning,  
513 see e.g., [Liu et al., 2021](#)). Cages were represented by a clipart illustration of a black fence which is freely  
514 available from <https://commons.wikimedia.org/wiki/File:Maki-fence-15.svg>, open-source and  
515 licensed under the Creative Commons CC0 1.0 Universal Public Domain Dedication, allowing further  
516 modification (for details, see <https://creativecommons.org/publicdomain/zero/1.0/>). When  
517 feedback was presented in the training and recall task conditions, correct responses were indicated  
518 by a fence colored in green and incorrect responses were signaled by a fence colored in red. The color  
519

520 of the original image was modified accordingly. All stimuli were presented against a white background.

## 521 **Hardware and software**

522 Behavioral responses were collected using two 4-button inline fiber optic response pads (Current  
523 Designs, Philadelphia, PA, USA), one for each hand, with a linear arrangement of four buttons (buttons  
524 were colored in blue, yellow, green, and red, from left to right). The two response pads were attached  
525 horizontally to a rectangular cushion that was placed in participants' laps such that they could place  
526 their fingers on the response buttons with arms comfortably extended while resting on the scanner  
527 bed. Participants were asked to place their index, middle, and ring finger of their left and right  
528 hand on the yellow, green, and red buttons of the left and right response pads, respectively. The  
529 fourth (blue) button on each response pad was masked with tape and participants were instructed to  
530 never use this response button. Behavioral responses on the response pads were transferred to the  
531 computer running the experimental task and mapped to the keyboard keys **z**, **g**, **r** and **w**, **n**, **d** for  
532 the left and right hand, respectively. The task was programmed in PsychoPy3 (version 3.0.11; [Peirce,](#)  
533 [2007, 2008](#); [Peirce et al., 2019](#)) and run on a Windows 7 computer with a monitor refresh-rate of 16.7  
534 ms. We recorded the presentation time stamps of all task events (onsets of all presentations of the  
535 fixation, stimulus, SRI, response, feedback, and ITI events) and confirmed that all components of the  
536 experimental task procedure were presented as expected.

## 537 **Instructions**

538 After participants entered the MRI scanner during the first study session and completed an anatomical  
539 T1-weighted (T1w) scan and a 5 min fMRI resting-state scan, they read the task instructions while  
540 lying inside the MRI scanner (for an illustration of the study procedure, see Fig. [S1](#)). Participants  
541 were asked to read all task instructions carefully (for the verbatim instructions, see Boxes [S1](#) to [S15](#)).  
542 They were further instructed to clarify any potential questions with the study instructor right away  
543 and to lie as still and relaxed as possible for the entire duration of the MRI scanning procedure. As  
544 part of the instructions, participants were presented with a cover story in order to increase motivation  
545 and engagement (see Box [S1](#)). Participants were told to see themselves in the role of a zookeeper in  
546 training whose main task is to ensure that all animals are in the correct cages. In all task conditions,  
547 participants were asked to always keep their fingers on the response buttons to be able to respond as  
548 quickly and as accurately as possible. The full task instructions can be found in the supplementary  
549 information (SI), translated to English (see SI, starting on page [7](#), Boxes [S1](#) to [S15](#)) from the original  
550 in German (see SI, page [11](#)).

## 551 **Training trials**

552 After participants read the instructions and clarified all remaining questions with the study instructors  
553 via the intercom, they completed the *training* phase of the task. The training condition was designed  
554 to explicitly teach participants the assignment of stimuli to response buttons. Each of the six animal  
555 stimuli selected per participant was randomly assigned to one of six response buttons. For the training  
556 condition, participants were told to see themselves in the role of a zookeeper in training in a public zoo  
557 whose task is to learn which animal belongs in which cage (see Box [S1](#)). During each trial, participants  
558 saw six black cages at the bottom of the screen with each cage belonging to one of the six animals.  
559 On each trial, an animal appeared above one of the six cages. Participants were tasked to press the

560 response button for that cage as fast and accurately as possible and actively remember the cage where  
561 the animal belonged (see Box S3 and Box S4). The task instructions emphasized that it would be very  
562 important for participants to actively remember which animal belonged in which cage and that they  
563 would have the chance to earn a higher bonus if they learned the assignment and responded accurately  
564 (see Box S5).

565 In total, participants completed 30 trials of the training condition. Across all trials, the pairwise  
566 ordering of stimuli was set to be balanced, with each pairwise sequential combination of stimuli  
567 presented exactly once, i.e., with  $n = 6$  stimuli, this resulted in  $n * (n - 1) = 6 * (6 - 1) = 30$  trials.  
568 In this sense, the stimulus order was drawn from a graph with all nodes connected to each other  
569 and an equal probability of  $p_{ij} = 0.2$  of transitioning from one node to any other node in the graph.  
570 This pairwise balancing of sequential combinations was used to ensure that participants would not  
571 learn any particular sequential order among the stimuli. Note, that this procedure only controlled for  
572 sequential order between pairs of consecutive stimuli but not higher-order sequential ordering of two  
573 steps or more.

574 On the first trial of the training condition, participants first saw a small black fixation cross that  
575 was displayed centrally on the screen for a fixed duration of 300 ms and signaled the onset of the  
576 following stimulus. The fixation cross was only shown on the first trial of the training phase, to allow  
577 for a short preparation signal before stimulus presentation began. Following the fixation cross, one of  
578 the animals was presented in the upper half of the screen above one of six cages that referred to the  
579 six response buttons and were presented in the lower half of the screen. The stimuli were shown for a  
580 fixed duration of 800 ms which was also the maximum time allowed for participants to respond. Note,  
581 that the instructions told participants that they would have 1 s to respond (see Box S4), an actual  
582 difference of 200 ms that was likely hardly noticeable. Following the stimulus, participants always  
583 received feedback that was shown for a fixed duration of 500 ms. If participants responded correctly,  
584 the cage corresponding to the correctly pressed response button, was shown in green. If participants  
585 did not respond correctly, the cage referring to the correct response button was shown in green and the  
586 cage referring to the incorrectly pressed response button was shown in red. If participants responded  
587 too late, the cage referring to the correct response button was shown in green and the German words  
588 “Zu langsam” (in English: “Too slow”) appeared in large red letters in the upper half of the screen.  
589 Finally, a small black fixation cross was shown during an ITI with a variable duration of  $M = 1500$   
590 ms. The ITIs were drawn from a truncated exponential distribution with a mean of  $M = 1.5$  s, a  
591 lower bound of  $x_1 = 1.0$  s and an upper bound of  $x_2 = 10.0$  s. To this end, we used the `truncexpon`  
592 distribution from the SciPy package (Virtanen et al., 2020) implemented in Python 3 (Van Rossum  
593 and Drake, 2009). The `truncexpon` distribution is described by three parameters, the shape  $b$ , the  
594 location  $\mu$  and the scale  $\beta$ . The support of the distribution is defined by the lower and upper bounds,  
595  $[x_1, x_2]$ , where  $x_1 = \mu$  and  $x_2 = b * \beta + \mu$ . We solved the latter equation for the shape  $b$  to get  
596  $b = (x_2 - x_1)/\beta$ . We chose the scale parameter  $\beta$  such that the mean of the distribution would be  
597  $M = 2.5$ . To this end, we applied `scipy.optimize.fsolve` (Virtanen et al., 2020) to a function of  
598 the scale  $\beta$  that becomes zero when  $truncexpon.mean((x_2 - x_1)/\beta, \mu, \beta) - M = 2.5$ . In total, the  
599 training phase took approximately 2 min to complete.

## 600 Recall trials

601 After participants finished the training phase of the task in the first experimental session, they com-  
602 pleted eight runs of the *recall* condition and another ninth run at the beginning of the second session

603 (for an illustration of the study procedure, see Fig. S1). The recall condition of the task mainly served  
604 two purposes: First, the recall condition was used to further train participants on the associations  
605 between animal stimuli and response keys. Second, the recall condition was designed to elicit object-  
606 specific neural activation patterns of the presented visual animal stimuli and the following motor  
607 response. The resulting neural activation patterns were later used to train the probabilistic classifiers.  
608 The cover story of the instructions told participants that they would be tested on how well they have  
609 learned the association between animals and response keys during the training phase (see Box S6).

610 In total, participants completed nine runs of the recall condition. Eight runs were completed during  
611 session 1 and an additional ninth run was completed at the beginning of session 2 in order to remind  
612 participants about the S-R mappings (for an illustration of the study procedure, see Fig. S1). Each  
613 run consisted of 60 trials. As in the training phase, the proportion of pairwise sequential combinations  
614 of stimuli was balanced within a run. Across all trials, each pairwise sequential combination of stimuli  
615 was presented twice, i.e., with  $n = 6$  stimuli, this results in  $n * (n - 1) * 2 = 6 * (6 - 1) * 2 = 60$   
616 trials. As for the training trials, the sequential ordering of stimuli was drawn from a graph with all  
617 nodes connected to each other and an equal probability of  $p_{ij} = 0.2$  of transitioning from one node  
618 to any other node in the graph. With 60 trials per run, each of the six animal stimuli was shown  
619 10 times per run. Given nine runs of the recall condition in total, this amounted to a maximum of  
620 90 trials per stimulus per participant of training examples for the classifiers. Including a ninth run  
621 at the beginning of session 2 offered two advantages. First, participants were reminded about the  
622 associations between the stimuli and response keys that they had learned extensively during session 1.  
623 Second, the ninth run allowed to investigate decoding performance across session boundaries. Note,  
624 that the two experimental sessions were separated by about one week. Although the pre-processing  
625 of fMRI data (for details, see section on fMRI pre-processing below) should align the data of the two  
626 sessions, remaining differences between the two sessions (e.g., positioning of the participant in the MRI  
627 scanner) could lead to a decrement in decoding accuracy when testing classifiers that were trained  
628 on session 1 data to data from session 2. Our decoding approach was designed such that pattern  
629 classifiers would be mainly trained on neural data from recall trials in session 1 but then applied to  
630 data from session 2.

631 As in training trials, the first trial of each run in the recall phase started with a black fixation  
632 cross on a white background that was presented for a fixed duration of 300 ms. Only the first trial of  
633 a run contained a fixation cross, to provide a preparatory signal for participants which would later be  
634 substituted for by the ITI. Participants were then presented with one of the six animal stimuli that  
635 was presented centrally on the screen for a fixed duration of 500 ms. Participants were instructed  
636 to not respond to the stimulus (see instructions in Box S7). To check if participants indeed did not  
637 respond during the stimulus or the following SRI, we also recorded responses during these trial events.  
638 During the breaks between task runs, participants received feedback about the proportion of trials  
639 on which they responded too early. If participants responded too early, they were reminded by the  
640 study instructors to not respond before the response screen. A variable SRI followed the stimulus  
641 presentation during which a fixation cross was presented again. Including a jittered SRI ensured that  
642 the neural responses to the visual stimulus and the motor response could be separated in time and  
643 reduce temporal autocorrelation. Following the SRI, the cages indicating the response buttons were  
644 displayed centrally on the screen for a fixed duration of 800 ms, which was also the response time  
645 limit for participants. If participants responded incorrectly, the cage referring to the correct response  
646 button was shown in green and the cage referring to the incorrectly pressed response key was shown

647 in red. If participants responded too late, the cage referring to the correct response button was shown  
648 in green and the German words “Zu langsam” (in English: “Too slow”) appeared in large red letters  
649 in the upper half of the screen. If participants responded correctly, the feedback screen was skipped.  
650 Each trial ended with an ITI with a variable duration of  $M = 2.5$  s. Both SRIs and ITIs were drawn  
651 from a truncated exponential distribution as on training trials (for details, see description of training  
652 trials above).

## 653 Graph trials

654 Following the ninth run of the recall condition in session 2, participants completed five runs of the *graph*  
655 condition (for an illustration of the study procedure, see Fig. S1). During graph trials, participants  
656 were exposed to a fast-paced stream of the same six animal stimuli as in the training and recall phase.  
657 Unbeknownst to participants, the sequential ordering of animal stimuli followed particular transition  
658 probabilities.

659 During the graph task, the sequential order of stimuli across trials was determined by two graph  
660 structures with distinct transition probabilities. In the first graph structure, each node had a high  
661 probability ( $p_{ij} = 0.7$ ) of transitioning to the next neighboring (i.e., transitioning from  $A$  to  $B$ ,  $B$  to  
662  $C$ ,  $C$  to  $D$ ,  $D$  to  $E$ ,  $E$  to  $F$ , and  $F$  to  $A$ ). Transitions to all other nodes (except the previous node)  
663 happened with equal probability of 0.1. Transitions to the previous node never occurred (transition  
664 probability of  $p_{ij} = 0.0$ ). These transition probabilities resulted in a sequential ordering of stimuli  
665 that can be characterized by a continuous progression in a unidirectional (i.e., clockwise) order around  
666 the ring-like graph structure. We therefore termed this graph structure the *unidirectional graph*  
667 (or *uni* in short). The second graph structure allowed sequential ordering that could also progress  
668 in counterclockwise order. To this end, stimuli were now equally likely to transition to the next  
669 neighboring but also the previous node (probability of  $p_{ij} = 0.35$ , i.e., splitting up the probability of  
670  $p_{ij} = 0.7$  of transitioning to the next neighboring node only in the unidirectional graph structure). As  
671 in the unidirectional graph, transitions to all other nodes happened with equal probability of  $p_{ij} = 0.1$ .  
672 Given that stimuli could follow a sequential ordering in both directions of the ring, we refer to this  
673 graph structure as the *bidirectional graph* (or *bi* in short).

674 Participants completed five runs of the graph task condition. Each run consisted of 240 trials.  
675 Each stimulus was shown 40 times per run. In the unidirectional graph, for each stimulus the most  
676 likely transitions (probability of  $p_{ij} = 0.7$ ) to the next neighboring node occurred 28 times per partic-  
677 ipant. Per stimulus and participant, 4 transitions to the other three possible nodes (low probability  
678 of  $p_{ij} = 0.1$ ) happened. No transitions to the previous node happened when stimulus transitions were  
679 drawn from a unidirectional graph structure. Together, this resulted in  $28 + 4 * 3 = 40$  presentations  
680 per stimulus, run and participant. For the bidirectional graph structure, transitions to the next neigh-  
681 boring and the previous node occurred 14 times per stimulus and to all other nodes 4 times as for  
682 the unidirectional graph structure. Together, this resulted in  $14 + 14 + 4 * 3 = 40$  presentations per  
683 stimulus, run and participant.

684 As for the other task conditions, only the first trial of the graph phase started with the presentation  
685 of a small black fixation cross that was presented centrally on the screen for a fixed duration of 300  
686 ms. Then, an animal stimulus was presented centrally on the screen for a fixed duration of 800 ms,  
687 which also constituted the time limit in which participants could respond with the correct response  
688 button. Participants did not receive feedback during the graph phase of the task in order to avoid any  
689 influence of feedback on graph learning. The stimulus was followed by an ITI with a mean duration

690 of 750 ms. The ITI in the graph trial phase was also drawn from a truncated exponential distribution  
691 with a mean of  $M = 750$  ms, a lower bound of  $x_1 = 500$  ms and an upper bound of  $x_2 = 5000$  ms.

692 Importantly, during the graph task, we also included long ITIs of 10 s in order to investigate  
693 on-task replay. As stated above, participants completed 240 trials of the main task per run. In each  
694 run, each stimulus was shown on a total of 40 trials. For each stimulus, every 10<sup>th</sup> trial on average  
695 was selected to be followed by a long ITI of 10 s. This meant that in each of the five main task runs,  
696 4 trials per stimulus were followed by a long ITI. In total, each participant experienced 24 long ITI  
697 trials per run and 120 long ITI trials across the entire experiment. The duration of 10 s (roughly  
698 corresponding to eight TRs at a repetition time (TR) of 1.25 s) was chosen based on our previous  
699 results showing that the large majority of sequential fMRI signals can be captured within this time  
700 period (cf. [Wittkuhn and Schuck, 2021](#), their Fig. 3).

## 701 Post-task questionnaire

702 After participants left the scanner in session 2, they were asked to complete a computerized post-task  
703 questionnaire consisting of four parts. First, participants were asked to report their handedness by  
704 selecting from three alternative options, “left”, “right” or “both”, in a forced-choice format. Note,  
705 that participants were required to be right-handed to participate in the study, hence this question  
706 merely served to record the self-reported handedness in addition to the participant details acquired  
707 as part of the recruitment procedure and demographic questionnaire assessment. Second, participants  
708 were asked whether they noticed any sequential order among the animal stimuli in the main task and  
709 could respond either “yes” or “no” in a forced-choice format. Third, if participants indicated that they  
710 noticed a sequential order of the stimuli (selecting “yes” on the previous question), they were asked  
711 to indicate during which run of the main task they had started to notice the ordering (selecting from  
712 run “1” to “5”). In case participants indicated that they did not notice a sequential ordering, they  
713 were asked to select “None” when asked about the run. Fourth, participants were presented with all  
714 sequential combinations of pairs of the animal stimuli and asked to indicate how likely animal A (on  
715 the left) was followed by animal B (on the right) during the Main condition of the task. Participants  
716 were instructed to follow their gut feeling in case they were uncertain about the probability ratings.  
717 With  $n = 6$  stimuli, this resulted in  $n * (n - 1) = 6 * (6 - 1) = 30$  trials. Participants indicated their  
718 response using a slider on a continuous scale from 0% to 100%. We recorded participants probability  
719 rating and response time on each trial. There was no time limit for any of the assessments in the  
720 questionnaire. Participants took  $M = 5.49$  min ( $SD = 2.38$  min; range: 2.23 to 12.63 min) to complete  
721 the questionnaire. The computerized questionnaire was programmed in PsychoPy3 (version 3.0.11;  
722 [Peirce, 2007, 2008; Peirce et al., 2019](#)) and run on the same Windows 7 computer that was used for  
723 the main experimental task.

## 724 Study procedure

725 All participants were screened for study and MRI eligibility during a telephone screening prior to  
726 participation. The study consisted of two experimental sessions. Upon arrival at the study center in  
727 both sessions, participants were first asked about any symptoms that could indicate an infection with  
728 the SARS-CoV-2 virus. The study instructors then measured participants’ body temperature which  
729 was required to not be higher than 37.5°C. Participants were asked to read and sign all the relevant  
730 study documents at home prior to their arrival at the study center.

731 **Session 1** The first MRI session started with a short localizer sequence of ca. 1 min during which  
732 participants were asked to rest calmly, close their eyes and move as little as possible. Once the  
733 localizer data was acquired, the study personnel aligned the field of view (FOV) for the acquisition  
734 of the T1w sequence. The acquisition of the T1w sequence took about 4 min to complete. Using the  
735 anatomical precision of the T1w images, the study personnel then aligned the FOV of the functional  
736 MRI sequences. Here, the lower edge of the FOV was first aligned to the visually identified anterior  
737 commissure - posterior commissure (AC-PC) line of the participant's brain. The FOV was then  
738 manually tilted by 20 degrees forwards relative to the rostro-caudal axis (positive tilt; for details see  
739 the section on "[MRI data acquisition](#)" on page 26). Shortly before the functional MRI sequences  
740 were acquired, we performed Advanced Shimming. During the shimming period, which took ca. 2  
741 min, participants were again instructed to move as little as possible and additionally asked to avoid  
742 swallowing to further reduce any potential movements. Next, we acquired functional MRI data during  
743 a resting-state period of 5 min. For this phase, participants were instructed to keep their eyes open  
744 and fixate a white fixation cross that was presented on a black background. Acquiring fMRI resting-  
745 state data before participants had any exposure to the task allowed us to record a resting-state period  
746 that was guaranteed to be free of any task-related neural activation or reactivation. Following this  
747 pre-task resting-state scan, participants read the task instructions inside the MRI scanner and were  
748 able to clarify any questions with the study instructions via the intercom system. Participants then  
749 performed the training phase of the task (for details, see the section "[Training trials](#)" on page 21)  
750 while undergoing acquisition of functional MRI data. The training phase took circa 2 min to complete.  
751 Following the training phase, participants performed eight runs of the recall phase of the task of circa 6  
752 min each while fMRI data was recorded. Before participants left the scanner, field maps were acquired.

753 **Session 2** At the beginning of the second session, participants first completed the questionnaire for  
754 MRI eligibility and the questionnaire on COVID-19 symptoms before entering the MRI scanner again.  
755 As in the first session, the second MRI session started with the acquisition of a short localizer sequence  
756 and a T1w sequence followed by the orientation of the FOV for the functional acquisitions and the  
757 Advanced Shimming. Participants were asked to rest calmly and keep their eyes closed during this  
758 period. Next, during the first functional sequence of the second study session, participants performed  
759 a ninth run of the recall phase of the task in order to remind them about the correct response buttons  
760 associated with each of the six stimuli. We then acquired functional resting-state scans of 3 min each  
761 and functional task scans of 10 min each in an interleaved fashion, starting with a resting-state scan.  
762 During the acquisition of functional resting-state data, participants were asked to rest calmly and  
763 fixate a small white cross on a black background that was presented on the screen. During each of  
764 the functional task scans, participants performed the graph learning phase of the task (for details, see  
765 section "[Graph trials](#)" on page 24). Importantly, half-way through the third block of the main task, the  
766 graph structure was changed without prior announcement towards the second graph structure. After  
767 the sixth resting-state acquisition, field maps were acquired and participants left the MRI scanner.

## 768 **MRI data acquisition**

769 All MRI data were acquired using a 32-channel head coil on a research-dedicated 3-Tesla Siemens  
770 Magnetom TrioTim MRI scanner (Siemens, Erlangen, Germany) located at the Max Planck Institute  
771 for Human Development in Berlin, Germany.

772 At the beginning of each of the two MRI recording sessions, high-resolution T1w anatomical Mag-

773 netization Prepared Rapid Gradient Echo (MPRAGE) sequences were obtained from each participant  
774 to allow co-registration and brain surface reconstruction (sequence specification: 256 slices; TR =  
775 1900 ms; echo time (TE) = 2.52 ms; flip angle (FA) = 9 degrees; inversion time (TI) = 900 ms; matrix  
776 size = 192 x 256; FOV = 192 x 256 mm; voxel size = 1 x 1 x 1 mm).

777 For the functional scans, whole-brain images were acquired using a segmented k-space and steady  
778 state T2\*-weighted multi-band (MB) echo-planar imaging (EPI) single-echo gradient sequence that is  
779 sensitive to the blood-oxygen-level dependent (BOLD) contrast. This measures local magnetic changes  
780 caused by changes in blood oxygenation that accompany neural activity (sequence specification: 64  
781 slices in interleaved ascending order; anterior-to-posterior (A-P) phase encoding direction; TR = 1250  
782 ms; TE = 26 ms; voxel size = 2 x 2 x 2 mm; matrix = 96 x 96; FOV = 192 x 192 mm; FA = 71  
783 degrees; distance factor = 0%; MB acceleration factor 4). Slices were tilted for each participant by 20  
784 degrees forwards relative to the rostro-caudal axis (positive tilt) to improve the quality of fMRI signal  
785 from the hippocampus (cf. [Weiskopf et al., 2006](#)) while preserving good coverage of occipito-temporal  
786 and motor brain regions. The same sequence parameters were used for all acquisitions of fMRI data.  
787 For each functional task run, the task began after the acquisition of the first four volumes (i.e., after  
788 5.00 s) to avoid partial saturation effects and allow for scanner equilibrium.

789 The first MRI session included nine functional task runs in total (for the study procedure, see  
790 Fig. S1). After participants read the task instructions inside the MRI scanner, they completed the  
791 training trials of the task which explicitly taught participants the correct mapping between stimuli  
792 and response keys. During this task phase, 80 volumes of fMRI were collected, which were not used  
793 in any further analysis. The other eight functional task runs during session 1 consisted of eight runs  
794 of the recall condition. Each run of the recall task was about 6 min in length, during which 320  
795 functional volumes were acquired. We also recorded two functional runs of resting-state fMRI data,  
796 one before and one after the task runs. Each resting-state run was about 5 min in length, during  
797 which 233 functional volumes were acquired.

798 The second MRI session included six functional task runs in total (for the study procedure, see  
799 Fig. S1). After participants entered the MRI scanner, they completed a ninth run of the recall task.  
800 As before, this run of the recall task was also about 6 min in length, during which 320 functional  
801 volumes were acquired. Participants then completed five runs of the graph learning task. Each run of  
802 the five graph learning runs was about 10 min in length, during which 640 functional volumes were  
803 acquired. The five runs of the graph learning task were interleaved with six recordings of resting-state  
804 fMRI data, each about 3 min in length, during which 137 functional volumes were acquired.

805 At the end of each scanning session, two short acquisitions with six volumes each were collected  
806 using the same sequence parameters as for the functional scans but with varying phase encoding  
807 polarities, resulting in pairs of images with distortions going in opposite directions between the two  
808 acquisitions (also known as the *blip-up / blip-down* technique). From these pairs the displacement  
809 maps were estimated and used to correct for geometric distortions due to susceptibility-induced field  
810 inhomogeneities as implemented in the *fMRIPrep* preprocessing pipeline ([Esteban et al., 2018](#)) (see  
811 details below). In addition, a whole-brain spoiled gradient recalled (GR) field map with dual echo-time  
812 images (sequence specification: 36 slices; A-P phase encoding direction; TR = 400 ms; TE1 = 4.92  
813 ms; TE2 = 7.38 ms; FA = 60 degrees; matrix size = 64 x 64; FOV = 192 x 192 mm; voxel size = 3  
814 x 3 x 3.75 mm) was obtained as a potential alternative to the blip-up / blip-down method described  
815 above.

816 We also measured respiration during each scanning session using a pneumatic respiration belt as



817 part of the Siemens Physiological Measurement Unit (PMU). Pulse data could not be recorded as the  
818 recording device could not be attached to the participants' index finger as it would have otherwise  
819 interfered with the motor responses.

## 820 MRI data preparation

821 **Conversion of data to the brain imaging data structure (BIDS) standard** The majority  
822 of the steps involved in preparing and preprocessing the MRI data employed recently developed tools  
823 and workflows aimed at enhancing standardization and reproducibility of task-based fMRI studies  
824 (for a similar data processing pipeline, see e.g., [Esteban et al., 2019a](#); [Wittkuhn and Schuck, 2021](#)).  
825 Version-controlled data and code management was performed using `DataLad` (version 0.13.0; [Halchenko](#)  
826 [et al., 2019, 2021](#)), supported by the `DataLad` handbook ([Wagner et al., 2020](#)). Following success-  
827 ful acquisition, all study data were arranged according to the brain imaging data structure (BIDS)  
828 specification ([Gorgolewski et al., 2016](#)) using the `HeuDiConv` tool (version 0.8.0.2; freely available  
829 from <https://github.com/ReproNim/reproin> or <https://hub.docker.com/r/repronim/reproin>)  
830 in combination with the `ReproIn` heuristic ([Visconti di Oleggio Castello et al., 2020](#)) (version 0.6.0)  
831 that allows automated creation of BIDS data sets from the acquired Digital Imaging and Commu-  
832 nications in Medicine (DICOM) images. To this end, the sequence protocol of the MRI data ac-  
833 quisition was set up to conform with the specification required by the `ReproIn` heuristic (for details  
834 of the heuristic, see [https://github.com/nipy/heudiconv/blob/master/heudiconv/heuristics/](https://github.com/nipy/heudiconv/blob/master/heudiconv/heuristics/reproin.py)  
835 [reproin.py](#)). `HeuDiConv` was run inside a `Singularity` container ([Kurtzer et al., 2017](#); [Sochat et al.,](#)  
836 [2017](#)) that was built from the most recent version (at the time of access) of a `Docker` container (tag  
837 0.8.0.2), available from <https://hub.docker.com/r/repronim/reproin/tags>. DICOMs were con-  
838 verted to the NIfTI-1 format using `dcm2niix` (version 1.0.20190410GCC6.3.0; [Li et al., 2016](#)). In  
839 order to make personal identification of study participants unlikely, we eliminated facial features from  
840 all high-resolution structural images using `pydeface` (version 2.0.0; [Gulban et al., 2019](#), available  
841 from <https://github.com/poldracklab/pydeface> or [https://hub.docker.com/r/poldracklab/](https://hub.docker.com/r/poldracklab/pydeface)  
842 [pydeface](#)). `pydeface` ([Gulban et al., 2019](#)) was run inside a `Singularity` container ([Kurtzer et al.,](#)  
843 [2017](#); [Sochat et al., 2017](#)) that was built from the most recent version (at the time of access) of a `Docker`  
844 container (tag 37-2e0c2d), available from <https://hub.docker.com/r/poldracklab/pydeface/tags>  
845 and used `Nipype`, version 1.3.0-rc1 ([Gorgolewski et al., 2011, 2019](#)). During the process of convert-  
846 ing the study data to BIDS the data set was queried using `pybids` (version 0.12.1; [Yarkoni et al.,](#)  
847 [2019a,b](#)), and validated using the `bids-validator` (version 1.5.4; [Gorgolewski et al., 2020](#)). The  
848 `bids-validator` ([Gorgolewski et al., 2020](#)) was run inside a `Singularity` container ([Kurtzer et al.,](#)  
849 [2017](#); [Sochat et al., 2017](#)) that was built from the most recent version (at the time of access) of a  
850 `Docker` container (tag v1.5.4), available from <https://hub.docker.com/r/bids/validator/tags>.

851 **MRI data quality control** The data quality of all functional and structural acquisitions were  
852 evaluated using the automated quality assessment tool `MRIQC`, version 0.15.2rc1 (for details, see [Es-](#)  
853 [teban et al., 2017](#), and the `MRIQC` documentation, available at [https://mriqc.readthedocs.io/en/](https://mriqc.readthedocs.io/en/stable/)  
854 [stable/](#)). The visual group-level reports of the estimated image quality metrics confirmed that the  
855 overall MRI signal quality of both anatomical and functional scans was highly consistent across par-  
856 ticipants and runs within each participant.

## 857 MRI data preprocessing

858 Preprocessing of MRI data was performed using fMRIPrep 20.2.0 (long-term support (LTS) release;  
859 [Esteban et al., 2018, 2019b](#), RRID:SCR\_016216), which is based on Nipype 1.5.1 ([Gorgolewski et al.,  
860 2011, 2019](#), RRID:SCR\_002502). Many internal operations of fMRIPrep use Nilearn 0.6.2 ([Abraham  
861 et al., 2014](#), RRID:SCR\_001362), mostly within the functional processing workflow. For more details  
862 of the pipeline, see the section corresponding to workflows in fMRIPrep's documentation at [https:  
863 //fmriprep.readthedocs.io/en/latest/workflows.html](https://fmriprep.readthedocs.io/en/latest/workflows.html). Note, that version 20.2.0 of fMRIPrep  
864 is a long-term support (LTS) release, offering long-term support and maintenance for four years.

865 **Preprocessing of anatomical MRI data using fMRIPrep** A total of two T1w images were found  
866 within the input BIDS data set, one from each study session. All of them were corrected for inten-  
867 sity non-uniformity (INU) using `N4BiasFieldCorrection` ([Tustison et al., 2010](#)), distributed with  
868 Advanced Normalization Tools (ANTs) 2.3.3 ([Avants et al., 2008](#), RRID:SCR\_004757). The T1w-  
869 reference was then skull-stripped with a Nipype implementation of the `antsBrainExtraction.sh`  
870 workflow (from ANTs), using `OASIS30ANTs` as target template. Brain tissue segmentation of cere-  
871 brospinal fluid (CSF), white-matter (WM) and gray-matter (GM) was performed on the brain-  
872 extracted T1w using `fast` (FMRIB Software Library (FSL) 5.0.9, RRID:SCR\_002823, [Zhang et al.,  
873 2001](#)). A T1w-reference map was computed after registration of two T1w images (after INU-correction)  
874 using `mri_robust_template` (FreeSurfer 6.0.1, [Reuter et al., 2010](#)). Brain surfaces were reconstructed  
875 using `recon-all` (FreeSurfer 6.0.1, RRID:SCR\_001847, [Dale et al., 1999](#)), and the brain mask es-  
876 timated previously was refined with a custom variation of the method to reconcile ANTs-derived  
877 and FreeSurfer-derived segmentations of the cortical GM of Mindboggle (RRID:SCR\_002438, [Klein  
878 et al., 2017](#)). Volume-based spatial normalization to two standard spaces (MNI152NLin6Asym,  
879 MNI152NLin2009cAsym) was performed through nonlinear registration with `antsRegistration` (ANTs  
880 2.3.3), using brain-extracted versions of both T1w reference and the T1w template. The following  
881 templates were selected for spatial normalization: FSL's MNI ICBM 152 non-linear 6<sup>th</sup> Generation  
882 Asymmetric Average Brain Stereotaxic Registration Model ([Evans et al., 2012](#), RRID:SCR\_002823;  
883 TemplateFlow ID: MNI152NLin6Asym), ICBM 152 Nonlinear Asymmetrical template version 2009c  
884 ([Fonov et al., 2009](#), RRID:SCR\_008796; TemplateFlow ID: MNI152NLin2009cAsym).

885 **Preprocessing of functional MRI data using fMRIPrep** For each of the BOLD runs found per  
886 participant (across all tasks and sessions), the following preprocessing was performed. First, a refer-  
887 ence volume and its skull-stripped version were generated using a custom methodology of fMRIPrep.  
888 A B0-nonuniformity map (or fieldmap) was estimated based on two (or more) echo-planar imaging  
889 (EPI) references with opposing phase-encoding directions, with `3dQwarp` ([Cox and Hyde, 1997](#), AFNI  
890 20160207). Based on the estimated susceptibility distortion, a corrected echo-planar imaging (EPI)  
891 reference was calculated for a more accurate co-registration with the anatomical reference. The BOLD  
892 reference was then co-registered to the T1w reference using `bbregister` (FreeSurfer) which implements  
893 boundary-based registration ([Greve and Fischl, 2009](#)). Co-registration was configured with six degrees  
894 of freedom. Head-motion parameters with respect to the BOLD reference (transformation matrices,  
895 and six corresponding rotation and translation parameters) are estimated before any spatiotemporal  
896 filtering using `mcflirt` (FSL 5.0.9, [Jenkinson et al., 2002](#)). BOLD runs were slice-time corrected using  
897 `3dTshift` from AFNI 20160207 ([Cox and Hyde, 1997](#), RRID:SCR\_005927). The BOLD time-series  
898 were resampled onto the following surfaces (FreeSurfer reconstruction nomenclature): `fsnative`. The

899 BOLD time-series (including slice-timing correction) were resampled onto their original, native space  
900 by applying a single, composite transform to correct for head-motion and susceptibility distortions.  
901 These resampled BOLD time-series will be referred to as preprocessed BOLD in original space, or just  
902 preprocessed BOLD. The BOLD time-series were resampled into standard space, generating a prepro-  
903 cessed BOLD run in MNI152NLin6Asym space. First, a reference volume and its skull-stripped version  
904 were generated using a custom methodology of `fMRIPrep`. Several confounding time-series were calcu-  
905 lated based on the preprocessed BOLD: framewise displacement (FD), DVARS and three region-wise  
906 global signals. FD was computed using two formulations following [Power et al.](#) (absolute sum of  
907 relative motions, 2014) and [Jenkinson et al.](#) (relative root mean square displacement between affines,  
908 2002). FD and DVARS are calculated for each functional run, both using their implementations in  
909 `Nipype` (following the definitions by [Power et al., 2014](#)). The three global signals are extracted within  
910 the CSF, the WM, and the whole-brain masks. Additionally, a set of physiological regressors were  
911 extracted to allow for component-based noise correction (`CompCor`, [Behzadi et al., 2007](#)). Principal  
912 components are estimated after high-pass filtering the preprocessed BOLD time-series (using a discrete  
913 cosine filter with 128s cut-off) for the two `CompCor` variants: temporal (`tCompCor`) and anatomical  
914 (`aCompCor`). `tCompCor` components are then calculated from the top 2% variable voxels within the  
915 brain mask. For `aCompCor`, three probabilistic masks (CSF, WM and combined CSF+WM) are gener-  
916 ated in anatomical space. The implementation differs from that of [Behzadi et al. \(2007\)](#) in that instead  
917 of eroding the masks by 2 pixels on BOLD space, the `aCompCor` masks are subtracted from a mask  
918 of pixels that likely contain a volume fraction of GM. This mask is obtained by dilating a GM mask  
919 extracted from the FreeSurfer’s `aseg` segmentation, and it ensures components are not extracted from  
920 voxels containing a minimal fraction of GM. Finally, the masks are resampled into BOLD space and  
921 binarized by thresholding at 0.99 (as in the original implementation). Components are also calculated  
922 separately within the WM and CSF masks. For each `CompCor` decomposition, the  $k$  components with  
923 the largest singular values are retained, such that the retained components’ time series are sufficient  
924 to explain 50 percent of variance across the nuisance mask (CSF, WM, combined, or temporal). The  
925 remaining components are dropped from consideration. The head-motion estimates calculated in the  
926 correction step were also placed within the corresponding confounds file. The confound time series  
927 derived from head motion estimates and global signals were expanded with the inclusion of temporal  
928 derivatives and quadratic terms for each ([Satterthwaite et al., 2013](#)). Frames that exceeded a threshold  
929 of 0.5 mm FD or 1.5 standardized DVARS were annotated as motion outliers. All resamplings can be  
930 performed with a single interpolation step by composing all the pertinent transformations (i.e. head-  
931 motion transform matrices, susceptibility distortion correction when available, and co-registrations  
932 to anatomical and output spaces). Gridded (volumetric) resamplings were performed using `antsAp-  
933 plyTransforms` (ANTs), configured with Lanczos interpolation to minimize the smoothing effects of  
934 other kernels ([Lanczos, 1964](#)). Non-gridded (surface) resamplings were performed using `mri_vol2surf`  
935 (`FreeSurfer`).

936 **Additional preprocessing of functional MRI data following `fMRIPrep`** Following preprocess-  
937 ing using `fMRIPrep`, the fMRI data were spatially smoothed using a Gaussian mask with a standard  
938 deviation (Full Width at Half Maximum (FWHM) parameter) set to 4 mm using an example `Nipype`  
939 smoothing workflow (see the [Nipype documentation](#) for details) based on the Smallest Univalued Seg-  
940 ment Assimilating Nucleus (SUSAN) algorithm as implemented in FSL ([Smith and Brady, 1997](#)). In  
941 this workflow, each run of fMRI data is separately smoothed using FSL’s SUSAN algorithm with the

942 brightness threshold set to 75% of the median value of each run and a mask constituting the mean  
943 functional image of each run.

## 944 **Multi-variate fMRI pattern analysis**

945 All fMRI pattern classification analyses were conducted using the open-source Python (Python Soft-  
946 ware Foundation, Python Language Reference, version 3.8.6) packages `Nilearn` (version 0.7.0; [Abra-  
947 ham et al., 2014](#)) and `scikit-learn` (version 0.24.1; [Pedregosa et al., 2011](#)). In all classification  
948 analyses, we trained an ensemble of six independent classifiers, one for each of the six event classes.  
949 Depending on the analysis, these six classes either referred to the identity of the six visual animal  
950 stimuli or the identity of the participant’s motor response, when training the classifiers with respect  
951 to the stimulus or the motor onset, respectively. For each class-specific classifier, labels of all other  
952 classes in the data were relabeled to a common “other” category. In order to ensure that the classifier  
953 estimates were not biased by relative differences in class frequency in the training set, the weights  
954 associated with each class were adjusted inversely proportional to the class frequencies in each train-  
955 ing fold. Given that there were six classes to decode, the frequencies used to adjust the classifiers’  
956 weights were  $\frac{1}{6}$  for the class of interest, and  $\frac{5}{6}$  for the “other” class, comprising any other classes.  
957 Adjustments to minor imbalances caused by the exclusion of erroneous trials were performed in the  
958 same way. We used separate logistic regression classifiers with identical parameter settings. All classi-  
959 fiers were regularized using L2 regularization. The  $C$  parameter of the cost function was fixed at the  
960 default value of  $C = 1.0$  for all participants. The classifiers employed the `lbfgs` algorithm to solve the  
961 multi-class optimization problem and were allowed to take a maximum of 4,000 iterations to converge.  
962 Pattern classification was performed within each participant separately, never across participants. For  
963 each example in the training set, we added 4 s to the event onset and chose the volume closest to  
964 that time point (i.e., rounding to the nearest volume) to center the classifier training on the expected  
965 peaks of the BOLD response (for a similar approach, see e.g., [Deuker et al., 2013](#)). At a TR of 1.25  
966 s this corresponded roughly to the fourth MRI volume which thus compromised a time window of  
967 3.75 s to 5.0 s after each event onset. We detrended the fMRI data separately for each run across all  
968 task conditions to remove low frequency signal intensity drifts in the data due to noise from the MRI  
969 scanner. For each classifier and run, the features were standardized ( $z$ -scored) by removing the mean  
970 and scaling to unit variance separately for each training and test set.

971 **Classification procedures** First, in order to assess the ability of the classifiers to decode the correct  
972 class from fMRI patterns, we conducted a leave-one-run-out cross-validation procedure for which data  
973 from seven task runs of the recall phase in session 1 were used for training and data from the left-out  
974 run (i.e., the eighth run) from session 1 was used for testing the classification performance. This  
975 procedure was repeated eight times so that each task run served as the testing set once. Classifier  
976 training was performed on data from all correct recall trials of the seven runs in the respective cross-  
977 validation fold. In each iteration of the leave-one-run-out procedure, the classifiers trained on seven out  
978 of eight runs were then applied separately to the data from the left-out run. Specifically, the classifiers  
979 were applied to (1) data from the recall trials of the left-out run, selecting volumes capturing the  
980 expected activation peaks to determine classification accuracy, and (2) data from the recall trials of  
981 the left-out run, selecting all volumes from the volume closest to the stimulus or response onset and  
982 the next seven volumes to characterize temporal dynamics of probabilistic classifier predictions on a  
983 single trial basis.

984 Second, we assessed decoding performance on recall trials across the two experimental sessions.  
985 The large majority of fMRI data that was used to train the classifiers was collected in session 1 (eight of  
986 nine runs of the recall task), but the trained classifiers were mainly applied to fMRI data from session  
987 2 (i.e., on-task intervals during graph trials). At the beginning of the second experimental session,  
988 participants completed another run of the recall task (i.e., a ninth run; for the study procedure, see  
989 Fig. S1). This additional task run mainly served the two purposes of (1) reminding participants about  
990 the correct S-R mapping that they had learned in session 1, and (2) to investigate the ability of the  
991 classifiers to correctly decode fMRI patterns in session 2 when they were only trained on session 1  
992 data. This second aspect is crucial, as the main focus of investigation is the potential reactivation of  
993 neural task representations in session 2 fMRI data. Thus, it is important to demonstrate that this  
994 ability is not influenced by losses in decoding performance due to decoding across session boundaries.  
995 In order to test cross-session decoding, we thus trained the classifiers on all eight runs of the recall  
996 condition in session 1 and tested their decoding performance on the ninth run of the recall condition  
997 in session 2. Classifiers trained on data from all nine runs of the recall task were subsequently applied  
998 to data from on-task intervals in graph trials in session 2. For the classification analyses in on-task  
999 intervals of the graph task, classifiers were trained on the peak activation patterns from all correct  
1000 recall trials (including session 1 and session 2 data) and then tested on all TR corresponding to the  
1001 graph task ITIs.

1002 **Feature selection** All participant-specific anatomical masks were created based on automated  
1003 anatomical labeling of brain surface reconstructions from the individual T1w reference image cre-  
1004 ated with Freesurfer’s `recon-all` (Dale et al., 1999) as part of the `fMRIPrep` workflow (Esteban et al.,  
1005 2018), in order to account for individual variability in macroscopic anatomy and to allow reliable la-  
1006 beling (Fischl et al., 2004; Poldrack, 2007). For the anatomical masks of occipito-temporal regions we  
1007 selected the corresponding labels of the cuneus, lateral occipital sulcus, pericalcarine gyrus, superior  
1008 parietal lobule, lingual gyrus, inferior parietal lobule, fusiform gyrus, inferior temporal gyrus, parahip-  
1009 pocampal gyrus, and the middle temporal gyrus (cf. Haxby et al., 2001; Wittkuhn and Schuck, 2021).  
1010 For the anatomical ROI of motor cortex, we selected the labels of the left and right gyrus precentralis  
1011 as well as gyrus postcentralis. The labels of each ROI are listed in Table 1. Only gray-matter voxels  
1012 were included in the generation of the masks as BOLD signal from non-gray-matter voxels cannot be  
1013 generally interpreted as neural activity (Kunz et al., 2018). Note, however, that due to the whole-brain  
1014 smoothing performed during preprocessing, voxel activation from brain regions outside the anatomical  
1015 mask but within the sphere of the smoothing kernel might have entered the anatomical mask (thus,  
1016 in principle, also including signal from surrounding non-gray-matter voxels).

ROI	Freesurfer labels (brain region)
Occipito-temporal	1005, 2005 (cuneus); 1011, 2011 (lateral occipital sulcus); 1021, 2021 (pericalcarine gyrus); 1029, 2029 (superio parietal lobule); 1013, 2013 (lingual gyrus); 1008, 2008 (inferior parietal lobule); 1007, 2007 (fusiform gyrus); 1009, 2009 (inferior temporal gyrus); 1016, 2016 (parahippocampal gyrus); 1015, 2015 (middle temporal gyrus)
Motor	1024, 2024 (left and right gyrus precentralis); 1022, 2022 (left and right gyrus postcentralis)

**Table 1:** Labels used to index brain regions to create participant-specific anatomical masks of selected ROIs based on Freesurfer’s `recon-all` labels (Dale et al., 1999)

## 1017 Statistical analyses

1018 All statistical analyses were run inside a Docker software container or, if analyses were executed on  
1019 a high performance computing (HPC), a Singularity version of the same container (Kurtzer et al.,  
1020 2017; Sochat et al., 2017). All main statistical analyses were conducted using LME models employing  
1021 the `lmer` function of the `lme4` package (version 1.1.27.1, Bates et al., 2015) in R (version 4.1.2, R  
1022 Core Team, 2019). If not stated otherwise, all models were fit with participants considered as a  
1023 random effect on both the intercept and slopes of the fixed effects, in accordance with results from  
1024 Barr et al. (2013) who recommend to fit the most complex model consistent with the experimental  
1025 design. If applicable, explanatory variables were standardized to a mean of zero and a standard  
1026 deviation of one before they entered the models. If necessary, we removed by-participant slopes  
1027 from the random effects structure to achieve a non-singular fit of the model (Barr et al., 2013).  
1028 Models were fitted using the Bound Optimization BY Quadratic Approximation (BOBYQA) optimizer  
1029 (Powell, 2007, 2009) with a maximum of 500,000 function evaluations and no calculation of gradient  
1030 and Hessian of nonlinear optimization solution. The likelihoods of the fitted models were assessed  
1031 using Type III analysis of variance (ANOVA) with Satterthwaite’s method. A single-step multiple  
1032 comparison procedure between the means of the relevant factor levels was conducted using Tukey’s  
1033 honest significant difference (HSD) test (Tukey, 1949), as implemented in the `emmeans` package in R  
1034 (version 1.7.0, Lenth, 2019; R Core Team, 2019). In all other analyses, we used one-sample *t*-tests  
1035 if group data was compared to a baseline or paired *t*-tests if two samples from the same population  
1036 were compared. If applicable, correction for multiple hypothesis testing was performed using the false  
1037 discovery rate (FDR) (Benjamini and Hochberg, 1995) or Bonferroni (Bonferroni, 1936) correction  
1038 method. If not stated otherwise, the  $\alpha$ -level was set to  $\alpha = 0.05$ , and analyses of response times  
1039 included data from correct trials only. When effects of stimulus transitions were analyzed, data from  
1040 the first trial of each run and the first trial after the change in transition structure were removed.

1041 **Statistical analyses of behavioral data** In order to test the a-priori hypothesis that behavioral  
1042 accuracy in each of the nine runs of the recall trials and five runs of the graph trials would be higher  
1043 than the chance-level, we performed a series of one-sided one-sample *t*-tests that compared partici-  
1044 pants’ mean behavioral accuracy per run against the chance level of  $100\%/6 = 16.67\%$ . Participants’  
1045 behavioral accuracy was calculated as the proportion of correct responses per run (in %). The effect  
1046 sizes (Cohen’s *d*) were calculated as the difference between the mean of behavioral accuracy scores  
1047 across participants and the chance baseline (16.67%), divided by the standard deviation of the data  
1048 (Cohen, 1988). The resulting *p*-values were adjusted for multiple comparisons using the Bonferroni  
1049 correction (Bonferroni, 1936).

1050 To examine the effect of task run on behavioral accuracy and response times in recall and graph  
1051 trials, we conducted an LME model that included all nine task runs of the recall trials (or five runs  
1052 of graph trials) as a numeric predictor variable (runs 1 to 9 and 1 to 5, respectively) as the main  
1053 fixed effect of interest as well as random intercepts and slopes for each participant. We also conceived  
1054 separate LME models that did not include data from the first task run of each task condition. These  
1055 models only included eight task runs of the recall trials (or four runs of the graph trials) as a numeric  
1056 predictor variable (runs 2 to 9 and 2 to 5, respectively) as the main fixed effect of interest as well as  
1057 by-participant random intercepts and slopes.

1058 Analyzing the effect of one-step transition probabilities on behavioral accuracy and response times,  
1059 we conducted two-sided paired *t*-tests comparing the effect of high vs. low transition probability

1060 separately for both unidirectional ( $p_{ij} = 0.7$  vs.  $p_{ij} = 0.1$ ) and bidirectional ( $p_{ij} = 0.35$  vs.  $p_{ij} = 0.1$ )  
1061 data. Effect sizes (Cohen’s  $d$ ) were calculated by dividing the mean difference of the paired samples  
1062 by the standard deviation of the difference (Cohen, 1988) and  $p$ -values were adjusted for multiple  
1063 comparisons across both graph conditions and response variables using the Bonferroni correction  
1064 (Bonferroni, 1936).

1065 In order to examine the effect of node distance on response times in graph trials, we conducted  
1066 separate LME models for data from the unidirectional and bidirectional graph structures. For LME  
1067 models of response time in unidirectional data, we included a linear predictor variable of node distance  
1068 (assuming a linear increase of response time with node distance; see Fig. 2d top right) as well as random  
1069 intercepts and slopes for each participant. The linear predictor variable was coded such that the node  
1070 distance linearly increased from  $-2$  to  $+2$  in steps of 1, modeling the hypothesized increase of response  
1071 time with node distance from 1 to 5 (centered on the node distance of 3). For LME models of response  
1072 time in bidirectional data, we included a quadratic predictor variable of node distance (assuming an  
1073 inverted U-shaped relationship between node distance and response time; see Fig. 2d bottom right) as  
1074 well as by-participant random intercepts and slopes. The quadratic predictor variable of node distance  
1075 was obtained by squaring the linear predictor variable. We also conducted separate LME models, that  
1076 did not include data of the most frequent transitions in both the uni- and bi-directional data, but  
1077 were otherwise specified in the same fashion.

1078 **Behavioral modeling based on the successor representation** We modeled successor represen-  
1079 tations (SRs) for each participant depending on the transitions they experienced in the task, including  
1080 training and recall trials. Specifically, each of the six stimuli was associated with a vector that reflected  
1081 a *running* estimate of the long-term visitation probability of all six stimuli, starting from the present  
1082 node. The successor matrix  $\mathbf{M}^t$  was therefore a 6-by-6 matrix that contained six predictive vectors,  
1083 one for each stimulus, and changed over time (hence the index  $t$ ). The SR matrix on the first trial was  
1084 initialized with a baseline expectation of  $\frac{1}{36}$  for each node. After a transition between stimuli  $s_t$  and  
1085  $s_{t+1}$ , the matrix row corresponding to  $s_t$  was updated following a temporal difference (TD) learning  
1086 rule (Dayan, 1993; Russek et al., 2017) as follows:

$$\mathbf{M}_{s_t,*}^t = \mathbf{M}_{s_t,*}^t + \alpha \left[ \mathbf{1}_{s_{t+1}} + \gamma \mathbf{M}_{s_{t+1},*}^t - \mathbf{M}_{s_t,*}^t \right] \quad (2)$$

1087 whereby  $\mathbf{1}_{s_{t+1}}$  is a zero vector with a 1 in the  $s_{t+1}$ <sup>th</sup> position,  $\mathbf{M}_{s_t,*}^t$  is the row corresponding to  
1088 stimulus  $s_t$  of matrix  $\mathbf{M}$ . The learning rate  $\alpha$  was arbitrarily set to a fixed value of 0.1, and the  
1089 discount parameter  $\gamma$  was varied in increments of 0.05 from 0 to 0.95, as described in the main text.  
1090 This meant that the SR matrix would change throughout the task to reflect the experienced transitions  
1091 of each participant, first reflecting the random transitions experienced during the training and recall  
1092 trials, then adapting to the first experienced graph structure and later to the second graph structure.  
1093 In order to relate the SR models to participants’ response times, we calculated how surprising each  
1094 transition in the graph learning task was – assuming participants’ expectations were based on the  
1095 current SR on the given trial,  $\mathbf{M}^t$ . To this end, we normalized  $\mathbf{M}^t$  to sum to 1, and then calculated  
1096 the Shannon information (Shannon, 1948) for each trial, reflecting how surprising the just observed  
1097 transition from stimulus  $i$  to  $j$  was given the history of previous transitions up to time point  $t$ :

$$I(j) = -\log_2(\tilde{m}_{i,j}^t) \quad (3)$$

1098 where  $\tilde{m}_{i,j}^t$  is the normalized  $(i, j)^{\text{th}}$  entry of SR matrix  $\mathbf{M}^t$ . Using the base-2 logarithm allowed  
1099 to express the units of information in bits (binary digits) and the negative sign ensured that the  
1100 information measure was always positive or zero.

1101 The final step in our analysis was to estimate LME models that tested how strongly this trial-wise  
1102 measure of SR-based surprise was related to participants' response times in the graph learning task,  
1103 for each level of the discount parameter  $\gamma$ . LME models therefore included fixed effects of the SR-  
1104 based Shannon surprise, in addition to factors of task run, graph order (uni – bi vs. bi – uni) and  
1105 graph structure (uni vs. bi) of the current run, as well as by-participant random intercepts and slopes.  
1106 Separate LME models were conducted for each level of  $\gamma$ , and model comparison of the twenty models  
1107 was performed using AIC, as reported in the main text. To independently investigate the effects of  
1108 graph condition (uni vs. bi) and graph order (uni – bi vs. bi – uni), we analyzed separate LME models  
1109 for each combination of the two factors, using only SR-based Shannon surprise as the main fixed effect  
1110 of interest, and including by-participant random intercepts and slopes.

1111 **Statistical analysis of classification accuracy and single-trial decoding time courses** In  
1112 order to assess the classifiers' ability to differentiate between the neural activation patterns of individ-  
1113 ual visual objects and motor responses, we compared the predicted visual object or motor response  
1114 of each example in the test set to the visual object or motor response that actually occurred on the  
1115 corresponding trial. We obtained an average classification accuracy score for each participant by cal-  
1116 culating the mean proportion of correct classifier predictions across all correctly answered recall trials  
1117 in session 1 (Fig. 4a). The mean decoding accuracy scores of all participants were then compared  
1118 to the chance baseline of  $100\%/6 = 16.67\%$  using a one-sided one-sample  $t$ -test, testing the a-priori  
1119 hypothesis that mean classification accuracy would be higher than the chance baseline. The effect  
1120 size (Cohen's  $d$ ) was calculated as the difference between the mean of accuracy scores and the chance  
1121 baseline, divided by the standard deviation of the data (Cohen, 1988). These calculations were per-  
1122 formed separately for each ROI and the resulting  $p$ -values were adjusted for multiple comparisons  
1123 using Bonferroni correction (Bonferroni, 1936).

1124 We also examined the effect of task run on classification accuracy in recall trials. To this end,  
1125 we conducted an LME model including the task run as the main fixed effect of interest as well as  
1126 by-participant random intercepts and slopes (Fig. 4c). We then assessed whether performance was  
1127 above the chance level for all nine task runs and conducted nine separate one-sided one-sample  $t$ -tests  
1128 separately per ROIs, testing the a-priori hypothesis that mean decoding accuracy would be higher  
1129 than the 16.67% chance-level in each task run. All  $p$ -values were adjusted for 18 multiple comparisons  
1130 (across nine runs and two ROIs) using the Bonferroni-correction (Bonferroni, 1936).

1131 Furthermore, we assessed the classifiers' ability to accurately detect the presence of visual objects  
1132 and motor responses on a single trial basis. For this analysis we applied the trained classifiers to fifteen  
1133 volumes from the volume closest to the event onset and examined the time courses of the probabilistic  
1134 classification evidence in response to the event on a single trial basis (Fig. 4b). In order to test if  
1135 the time series of classifier probabilities reflected the expected increase of classifier probability for  
1136 the event occurring on a given trial, we compared the time series of classifier probabilities related to  
1137 the classified class with the mean time courses of all other classes using a two-sided paired  $t$ -test at  
1138 the fourth TR from event onset. Classifier probabilities were normalized by dividing each classifier  
1139 probability by the sum of the classifier probabilities across all fifteen TRs of a given trial. Here,  
1140 we used the Bonferroni-correction method (Bonferroni, 1936) to adjust for multiple comparisons of



1141 two observations. In the main text, we report the results for the peak in classification probability  
1142 of the true class, corresponding to the fourth TR after stimulus onset. The effect size (Cohen’s  $d$ )  
1143 was calculated as the difference between the means of the probabilities of the current versus all other  
1144 stimuli, divided by the standard deviation of the difference (Cohen, 1988).

1145 **Statistical analyses of classifier time courses on graph trials** Classifier probabilities on graph  
1146 trials indicated that the fMRI signal was strongly dominated by the activation of the event on the  
1147 current trial. In order to test this effect, we calculated the mean classifier probabilities for the current  
1148 and all other five events of the current trial across all eight TRs in the ITIs. The mean classifier prob-  
1149 abilities of the current event were then compared to the mean classifier probabilities of all other events  
1150 using two two-sided paired  $t$ -tests, one for each ROI. The Bonferroni-correction method Bonferroni  
1151 (1936) was used to correct the  $p$ -values for two comparisons. The effect size (Cohen’s  $d$ ) was calculated  
1152 as the difference between the means of the probabilities of the current versus all other events, divided  
1153 by the standard deviation of the difference Cohen (1988).

1154 After excluding data from the event of the current trial, we analyzed the effect of node distance on  
1155 classifier probabilities for all non-displayed items using separate LME models for each graph structure,  
1156 similar to the analysis of response times described above. Based on our previous findings indicating  
1157 that the ordering of sequential neural events unfolds in the same order in earlier TRs and in reverse  
1158 order in later TRs (cf. Wittkuhn and Schuck, 2021), we also included a fixed effect of interval phase  
1159 (early TRs 1–4 vs. late TRs 5–8). In addition, each model included a fixed effect of ROI (occipito-  
1160 temporal vs. sensorimotor). As for response times (see above), LME models of classifier probabilities  
1161 in unidirectional or bidirectional data included a linear or quadratic predictor variable of node distance,  
1162 respectively, as well as random intercepts and slopes for each participant. In order to examine the effect  
1163 of a linear predictor in bidirectional data and the effect of the quadratic predictor in unidirectional  
1164 data, predictor variables were switched accordingly, but otherwise the LME were conducted as before.  
1165 Finally, we also directly compared the fits of a linear and quadratic model for each graph condition,  
1166 ROI, and interval phase and quantified the model comparison using AIC.

1167 **Predicting sequence probability during on-task intervals** We computed how likely it was  
1168 to observe each 5-item sequence of stimuli under the assumption that participants were internally  
1169 sampling from an SR model of the unidirectional or bidirectional graph structure. This was done in  
1170 two steps.

1171 First, we computed an ideal SR representation based on the true transition probabilities for each  
1172 graph structure. Specifically, we defined the true transition function  $\mathbf{T}$ , as given by a graph, such that  
1173 each entry  $t_{ij}$  reflected the true probability of transitioning from image  $i$  to  $j$ . Following the main  
1174 ideas of the SR, we then calculated the long-term visitation probabilities as the time-discounted 5-step  
1175 probabilities following the Chapman-Kolmogorov Equation:

$$\hat{\mathbf{M}} = \mathbf{T} + \gamma\mathbf{T}^2 + \gamma^2\mathbf{T}^3 + \gamma^3\mathbf{T}^4 + \gamma^4\mathbf{T}^5 \quad (4)$$

1176 The discount rate  $\gamma$  was set to 0.3. We used five steps since more steps make little practical dif-  
1177 ference given the exponential discounting. The theoretical sequence probabilities for a given sequence  
1178  $\mathbf{s}$  were then computed as the product of probabilities for all pairwise transitions  $(i, j)$  in the sequence,  
1179 according to the approximated and normalized SR matrix:

$$p(\mathbf{s}) = \prod_{i,j \in \mathbf{s}} \tilde{m}_{i,j} \quad (5)$$

1180 Second, we approximated how likely it was to observe a sequence in the fMRI signal, given a  
1181 particular sequence event in the brain. Our previous work has investigated which sequences are  
1182 observed in classifier probabilities for a known true sequence (Wittkuhn and Schuck, 2021), and found  
1183 that random reordering of items (induced by noise) was most prominent for the middle sequence items,  
1184 and less severe for the start and end items. To model this effect, we set up a hidden markov model  
1185 (HMM) in which the emission probabilities for the items that came first or last in a sequence were  
1186 tuned sharply, sampled from a Gaussian distribution with a standard deviation of 0.5. This meant  
1187 that the probability to observe the true item was 79%, and the probabilities to observe other items  
1188 decreased sharply with distance from the true sequence position. The intermediate items had emission  
1189 probabilities sampled from a Gaussian with a larger standard deviation of 2, yielding a much flatter  
1190 distribution (probability to observe the true item at these positions was merely 19.9%). Using the  
1191 HMM framework, we then computed the “forward” probabilities to observe a specific sequence given  
1192 the transitions of a true sequence and the specified emission probabilities.

1193 Finally, we combined the two probabilities that resulted from steps 1 and 2: (1) how likely a given  
1194 sequence was to have resulted from a sample of an SR-based internal model of a graph structure,  
1195 and (2) how likely it was to *observe* a sequence in the fMRI signal, given a specific sequence has  
1196 been reactivated in the brain. To obtain our final estimates, we multiplied these probabilities for  
1197 each sequence. This yielded the total probability to *observe* each sequence, assuming a true sequence  
1198 distribution that results from sampling from the SR model, and a noise model that relates true to  
1199 observed sequences.

1200 To examine the relationship between predicted sequences based on this approach and observed  
1201 sequences in fMRI during on-task intervals, we ordered the classes by their classifier probabilities  
1202 within each TR (removing the class of the stimulus shown on the current trial) to obtain the observed  
1203 frequencies for each of the possible 120 5-item sequences across all TRs of the on-task intervals during  
1204 the graph learning task, separately for each participant, ROI and graph condition. The resulting  
1205 distribution indicated how often classifier probabilities within TRs were ordered according to the 120  
1206 sequential 5-item combinations. This distribution was then averaged across participants for each of the  
1207 120 sequences and correlated with the sequence probability based on the HMM approach described  
1208 above, separately for each ROI and graph condition (using Pearson’s correlation across 120 data  
1209 points).

1210 **Calculating the TR-wise sequentiality metric** To analyze evidence for sequential replay during  
1211 on-task intervals in graph trials, we calculated a sequentiality metric quantified by the slope of a linear  
1212 regression between the classifier probabilities and each of the  $5! = 120$  possible sequential orderings  
1213 of a 5-item sequence in each TR, similar to our previous work (Wittkuhn and Schuck, 2021). We  
1214 next separated the regression slope data based on how likely the permuted sequences were given the  
1215 transition probabilities of the two graph structures in our experiment. To determine the probabilities of  
1216 each possible sequential ordering of the 5-item sequences, we used the HMM approach described above  
1217 to obtain the probability of all the  $5! = 120$  sequences, assuming a particular starting position (i.e.,  
1218 the event on the current trial). Next, we ranked the permuted sequences according to their probability  
1219 given the graph structures which allowed us to separately investigate sequentiality for the most and

1220 the least likely sequences based on the graph structure. We then separated the ranked sequences into  
1221 quintiles, i.e., five groups of ranked sequences from the least likely to the most likely 20%. Finally,  
1222 we averaged the regression slopes separately for both ROIs, the two graph structures and the early  
1223 and late TRs and compared the average slope against zero (the assumption of no sequentiality). The  
1224 mean slope coefficients of all participants were compared to zero using a series of two-sided one-sample  
1225 *t*-test, one for each graph condition, ROI, interval phase and sequence ranking bracket. *p*-values were  
1226 adjusted for multiple comparisons using Bonferroni correction (Bonferroni, 1936). The effect size  
1227 (Cohen's *d*) was calculated as the difference between the mean of slope coefficients and the baseline,  
1228 divided by the standard deviation of the data (Cohen, 1988).

## 1229 Data and code availability statement

1230 Behavioral and MRI data as well as custom code used in this study will be made available upon  
1231 publication in a peer-reviewed journal.

## 1232 Acknowledgments

1233 This work was supported by an Independent Max Planck Research Group grant awarded to N.W.S  
1234 by the Max Planck Society (M.TN.A.BILD0004), and a Starting Grant awarded to N.W.S by the  
1235 European Union (ERC-2019-StG REPLAY-852669). L.W. is a pre-doctoral fellow of the Interna-  
1236 tional Max Planck Research School on Computational Methods in Psychiatry and Ageing Research  
1237 (IMPRS COMP2PSYCH). The participating institutions are the Max Planck Institute for Human De-  
1238 velopment, Berlin, Germany, and University College London, London, UK. For more information, see  
1239 <https://www.mps-ucl-centre.mpg.de/en/comp2psych>. We also thank Leonardo Pettini for help  
1240 with task development, Gregor Caregnato for help with participant recruitment and study coordi-  
1241 nation, Sonali Beckmann, Sam Chien (<https://orcid.org/0000-0003-4306-1308>), Theresa Fox,  
1242 Sam Hall-McMaster (<https://orcid.org/0000-0003-1641-979X>), Nir Moneta (<https://orcid.org/0000-0001-6125-4117>),  
1243 Liliana Polyanska (<https://orcid.org/0000-0002-0842-8787>), Na-  
1244 dine Taube, and Kateryna Yasynska – in alphabetical order of last names – for assistance with MRI  
1245 data acquisition, Anika Löwe (<https://orcid.org/0000-0003-3132-5767>) for help with MRI data  
1246 collection and comments on a previous version of this manuscript, Ondřej Zíka (<https://orcid.org/0000-0003-0483-4443>)  
1247 for help with MRI data collection and statistical analyses, Michael Krause  
1248 for help with high performance computing (HPC), all members of the Max Planck Research Group  
1249 NeuroCode for helpful discussions about the contents of this manuscript, and all participants for their  
1250 participation.

## 1251 Author Contributions

1252 The following list of author contributions is based on the CRediT taxonomy (Brand et al., 2015). For  
1253 details on each type of author contribution, please see Brand et al. (2015).

1254 Conceptualization: L.W., N.W.S.; Methodology: L.W., L.M.K., N.W.S.; Software: L.W., L.M.K.,  
1255 N.W.S.; Validation: L.W.; Formal analysis: L.W., N.W.S.; Investigation: L.W., L.M.K.; Resources:  
1256 L.W., N.W.S.; Data curation: L.W., L.M.K.; Writing - original draft: L.W.; Writing - review &  
1257 editing: L.W., L.M.K., N.W.S.; Visualization: L.W.; Supervision: N.W.S.; Project administration:  
1258 L.W., N.W.S.; Funding acquisition: N.W.S.

1259 **Competing Interests**

1260 The authors declare no competing interests.

## 1261 References

- 1262 Alexandre Abraham, Fabian Pedregosa, Michael Eickenberg, Philippe Gervais, Andreas Mueller,  
1263 Jean Kossaifi, Alexandre Gramfort, Bertrand Thirion, and Gaël Varoquaux. Machine learning  
1264 for neuroimaging with scikit-learn. *Frontiers in Neuroinformatics*, 8, Feb 2014. ISSN 1662-5196.  
1265 doi:[10.3389/fninf.2014.00014](https://doi.org/10.3389/fninf.2014.00014). URL <http://dx.doi.org/10.3389/fninf.2014.00014>.
- 1266 B. Avants, C Epstein, M. Grossman, and J. Gee. Symmetric diffeomorphic image registration with  
1267 cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical*  
1268 *Image Analysis*, 12(1):26–41, Feb 2008. ISSN 1361-8415. doi:[10.1016/j.media.2007.06.004](https://doi.org/10.1016/j.media.2007.06.004). URL  
1269 <http://dx.doi.org/10.1016/j.media.2007.06.004>.
- 1270 David Badre and Derek Evan Nee. Frontal cortex and the hierarchical control of  
1271 behavior. *Trends in Cognitive Sciences*, 22(2):170–188, 2018. ISSN 1364-6613.  
1272 doi:<https://doi.org/10.1016/j.tics.2017.11.005>. URL [https://www.sciencedirect.com/science/](https://www.sciencedirect.com/science/article/pii/S1364661317302450)  
1273 [article/pii/S1364661317302450](https://www.sciencedirect.com/science/article/pii/S1364661317302450).
- 1274 Jan Balaguer, Hugo Spiers, Demis Hassabis, and Christopher Summerfield. Neural mechanisms of  
1275 hierarchical planning in a virtual subway network. *Neuron*, 90(4):893 – 903, 2016. ISSN 0896-  
1276 6273. doi:<https://doi.org/10.1016/j.neuron.2016.03.037>. URL [http://www.sciencedirect.com/](http://www.sciencedirect.com/science/article/pii/S0896627316300575)  
1277 [science/article/pii/S0896627316300575](http://www.sciencedirect.com/science/article/pii/S0896627316300575).
- 1278 Dale J. Barr, Roger Levy, Christoph Scheepers, and Harry J. Tily. Random effects structure for  
1279 confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3):255–  
1280 278, 2013. ISSN 0749596X. doi:[10.1016/j.jml.2012.11.001](https://doi.org/10.1016/j.jml.2012.11.001). URL <http://dx.doi.org/10.1016/j.jml.2012.11.001>.
- 1282 Helen C. Barron, Mona M. Garvert, and Timothy E. J. Behrens. Repetition suppression: a means  
1283 to index neural representations using bold? *Philosophical Transactions of the Royal Society B:*  
1284 *Biological Sciences*, 371(1705):20150355, Oct 2016. ISSN 1471-2970. doi:[10.1098/rstb.2015.0355](https://doi.org/10.1098/rstb.2015.0355).  
1285 URL <http://dx.doi.org/10.1098/rstb.2015.0355>.
- 1286 Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. Fitting linear mixed-effects  
1287 models using lme4. *Journal of Statistical Software*, 67(1):1–48, 2015. ISSN 1548-7660.  
1288 doi:[10.18637/jss.v067.i01](https://doi.org/10.18637/jss.v067.i01). URL <https://www.jstatsoft.org/v067/i01>.
- 1289 Timothy E.J. Behrens, Timothy H. Muller, James C.R. Whittington, Shirley Mark, Alon B.  
1290 Baram, Kimberly L. Stachenfeld, and Zeb Kurth-Nelson. What is a cognitive map? orga-  
1291 nizing knowledge for flexible behavior. *Neuron*, 100(2):490–509, Oct 2018. ISSN 0896-6273.  
1292 doi:[10.1016/j.neuron.2018.10.002](https://doi.org/10.1016/j.neuron.2018.10.002). URL <http://dx.doi.org/10.1016/j.neuron.2018.10.002>.
- 1293 Yashar Behzadi, Khaled Restom, Joy Liau, and Thomas T. Liu. A component based noise correc-  
1294 tion method (CompCor) for BOLD and perfusion based fMRI. *NeuroImage*, 37(1):90–101, Aug  
1295 2007. ISSN 1053-8119. doi:[10.1016/j.neuroimage.2007.04.042](https://doi.org/10.1016/j.neuroimage.2007.04.042). URL [http://dx.doi.org/10.1016/](http://dx.doi.org/10.1016/j.neuroimage.2007.04.042)  
1296 [j.neuroimage.2007.04.042](http://dx.doi.org/10.1016/j.neuroimage.2007.04.042).
- 1297 Jacob L. S. Bellmund, William de Cothi, Tom A. Ruiter, Matthias Nau, Caswell Barry, and Chris-  
1298 tian F. Doeller. Deforming the metric of cognitive maps distorts memory. *Nature Human Behaviour*,  
1299 4(2):177–188, 2020. doi:[10.1038/s41562-019-0767-3](https://doi.org/10.1038/s41562-019-0767-3). URL [https://doi.org/10.1038/s41562-](https://doi.org/10.1038/s41562-019-0767-3)  
1300 [019-0767-3](https://doi.org/10.1038/s41562-019-0767-3).

- 1301 Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful  
1302 approach to multiple testing. *Journal of the Royal Statistical Society*, 57(1):289–300, 1995. ISSN  
1303 00359246. URL <http://www.jstor.org/stable/2346101>.
- 1304 Carlo Emilio Bonferroni. Teoria statistica delle classi e calcolo delle probabilità. *Pubblicazioni del R*  
1305 *Istituto Superiore di Scienze Economiche e Commerciali di Firenze*, 8:3–62, 1936.
- 1306 Aaron M. Bornstein and Nathaniel D. Daw. Dissociating hippocampal and striatal contributions  
1307 to sequential prediction learning. *European Journal of Neuroscience*, 35(7):1011–1023, 2012.  
1308 doi:<https://doi.org/10.1111/j.1460-9568.2011.07920.x>. URL [https://onlinelibrary.wiley.com/](https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1460-9568.2011.07920.x)  
1309 [doi/abs/10.1111/j.1460-9568.2011.07920.x](https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1460-9568.2011.07920.x).
- 1310 Sander E. Bosch, Janneke F. M. Jehee, Guillén Fernández, and Christian F. Doeller. Reinstatement of  
1311 associative memories in early visual cortex is signaled by the hippocampus. *Journal of Neuroscience*,  
1312 34(22):7493–7500, 2014. ISSN 0270-6474. doi:[10.1523/JNEUROSCI.0805-14.2014](https://doi.org/10.1523/JNEUROSCI.0805-14.2014). URL [https:](https://www.jneurosci.org/content/34/22/7493)  
1313 [//www.jneurosci.org/content/34/22/7493](https://www.jneurosci.org/content/34/22/7493).
- 1314 Amy Brand, Liz Allen, Micah Altman, Marjorie Hlava, and Jo Scott. Beyond authorship: at-  
1315 tribution, contribution, collaboration, and credit. *Learned Publishing*, 28(2):151–155, apr 2015.  
1316 doi:[10.1087/20150211](https://doi.org/10.1087/20150211). URL <https://doi.org/10.1087/20150211>.
- 1317 Iva K. Brunec and Ida Momennejad. Predictive representations in hippocampal and prefrontal hier-  
1318 archies. *Journal of Neuroscience*, 2021. ISSN 0270-6474. doi:[10.1523/JNEUROSCI.1327-21.2021](https://doi.org/10.1523/JNEUROSCI.1327-21.2021).  
1319 URL <https://www.jneurosci.org/content/early/2021/11/17/JNEUROSCI.1327-21.2021>.
- 1320 Margaret F Carr, Shantanu P Jadhav, and Loren M Frank. Hippocampal replay in the awake state:  
1321 a potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, 14(2):147–153,  
1322 Jan 2011. ISSN 1546-1726. doi:[10.1038/nn.2732](https://doi.org/10.1038/nn.2732). URL <http://dx.doi.org/10.1038/nn.2732>.
- 1323 Jacob Cohen. Statistical power analysis for the behavioral sciences. *Lawrence Erlbaum Associates*,  
1324 1988.
- 1325 Alexandra O. Constantinescu, Jill X. O'Reilly, and Timothy E. J. Behrens. Organizing conceptual  
1326 knowledge in humans with a gridlike code. *Science*, 352(6292):1464–1468, Jun 2016. ISSN 1095-  
1327 9203. doi:[10.1126/science.aaf0941](https://doi.org/10.1126/science.aaf0941). URL <http://dx.doi.org/10.1126/science.aaf0941>.
- 1328 Robert W. Cox and James S. Hyde. Software tools for analysis and visualization of fmri data.  
1329 *NMR in Biomedicine*, 10(4-5):171–178, Jun 1997. ISSN 1099-1492. doi:[10.1002/\(sici\)1099-1492](https://doi.org/10.1002/(sici)1099-1492(199706/08)10:4/5<171::aid-nbm453>3.0.co;2-l)-  
1330 [1492\(199706/08\)10:4/5<171::aid-nbm453>3.0.co;2-l](https://doi.org/10.1002/(sici)1099-1492(199706/08)10:4/5<171::aid-nbm453>3.0.co;2-l). URL [http://dx.doi.org/10.1002/\(SICI\)](http://dx.doi.org/10.1002/(SICI)1099-1492(199706/08)10:4/5<171::aid-nbm453>3.0.co;2-l)  
1331 [1099-1492\(199706/08\)10:4/5<171::aid-nbm453>3.0.co;2-l](http://dx.doi.org/10.1002/(SICI)1099-1492(199706/08)10:4/5<171::aid-nbm453>3.0.co;2-l).
- 1332 Anders M. Dale. Optimal experimental design for event-related fmri. *Human Brain Mapping*, 8(2-3):  
1333 109–114, 1999. ISSN 1097-0193. doi:[10.1002/\(sici\)1097-0193\(1999\)8:2/3<109::aid-hbm7>3.0.co;2-](https://doi.org/10.1002/(sici)1097-0193(1999)8:2/3<109::aid-hbm7>3.0.co;2-w)  
1334 [w](https://doi.org/10.1002/(sici)1097-0193(1999)8:2/3<109::aid-hbm7>3.0.co;2-w). URL [http://dx.doi.org/10.1002/\(SICI\)1097-0193\(1999\)8:2/3<109::aid-hbm7>3.0.co;](http://dx.doi.org/10.1002/(SICI)1097-0193(1999)8:2/3<109::aid-hbm7>3.0.co;2-w)  
1335 [2-w](http://dx.doi.org/10.1002/(SICI)1097-0193(1999)8:2/3<109::aid-hbm7>3.0.co;2-w).
- 1336 Anders M. Dale, Bruce Fischl, and Martin I. Sereno. Cortical surface-based analysis. *NeuroImage*, 9  
1337 (2):179–194, Feb 1999. ISSN 1053-8119. doi:[10.1006/nimg.1998.0395](https://doi.org/10.1006/nimg.1998.0395). URL [http://dx.doi.org/](http://dx.doi.org/10.1006/nimg.1998.0395)  
1338 [10.1006/nimg.1998.0395](http://dx.doi.org/10.1006/nimg.1998.0395).

- 1339 Peter Dayan. Improving generalization for temporal difference learning: The successor representation.  
1340 *Neural Computation*, 5(4):613–624, Jul 1993. ISSN 1530-888X. doi:[10.1162/neco.1993.5.4.613](https://doi.org/10.1162/neco.1993.5.4.613). URL  
1341 <http://dx.doi.org/10.1162/neco.1993.5.4.613>.
- 1342 Lorena Deuker, J. Olligs, J. Fell, T. A. Kranz, F. Mormann, C. Montag, M. Reuter, C. E. Elger, and  
1343 Nikolai Axmacher. Memory consolidation by replay of stimulus-specific neural activity. *Journal of*  
1344 *Neuroscience*, 33(49):19373–19383, Dec 2013. ISSN 1529-2401. doi:[10.1523/jneurosci.0414-13.2013](https://doi.org/10.1523/jneurosci.0414-13.2013).  
1345 URL <http://dx.doi.org/10.1523/JNEUROSCI.0414-13.2013>.
- 1346 Kamran Diba and György Buzsáki. Forward and reverse hippocampal place-cell sequences during  
1347 ripples. *Nature Neuroscience*, 10(10):1241–1242, Sep 2007. ISSN 1546-1726. doi:[10.1038/nn1961](https://doi.org/10.1038/nn1961).  
1348 URL <http://dx.doi.org/10.1038/nn1961>.
- 1349 Bradley B Doll, Katherine D Duncan, Dylan A Simon, Daphna Shohamy, and Nathaniel D Daw.  
1350 Model-based choices involve prospective neural activity. *Nature Neuroscience*, 18(5):767–772, Mar  
1351 2015. ISSN 1546-1726. doi:[10.1038/nn.3981](https://doi.org/10.1038/nn.3981). URL <http://dx.doi.org/10.1038/nn.3981>.
- 1352 S. L. Eagleman and V. Dragoi. Image sequence reactivation in awake V4 networks. *Proceed-*  
1353 *ings of the National Academy of Sciences*, 109(47):19450–19455, Nov 2012. ISSN 1091-6490.  
1354 doi:[10.1073/pnas.1212059109](https://doi.org/10.1073/pnas.1212059109). URL <http://dx.doi.org/10.1073/pnas.1212059109>.
- 1355 Matthias Ekman, Peter Kok, and Floris P. de Lange. Time-compressed preplay of anticipated events  
1356 in human primary visual cortex. *Nature Communications*, 8(15276):1–9, May 2017. ISSN 2041-1723.  
1357 doi:[10.1038/ncomms15276](https://doi.org/10.1038/ncomms15276).
- 1358 Eran Eldar, Gaëlle Lièvre, Peter Dayan, and Raymond J Dolan. The roles of online and offline  
1359 replay in planning. *eLife*, 9, Jun 2020. ISSN 2050-084X. doi:[10.7554/elife.56911](https://doi.org/10.7554/elife.56911). URL <http://dx.doi.org/10.7554/eLife.56911>.  
1360 <http://dx.doi.org/10.7554/eLife.56911>.
- 1361 Oscar Esteban, Daniel Birman, Marie Schaer, Oluwasanmi O. Koyejo, Russell A. Poldrack,  
1362 and Krzysztof J. Gorgolewski. MRIQC: Advancing the automatic prediction of image qual-  
1363 ity in MRI from unseen sites. *PLoS ONE*, 12(9):e0184661, Sep 2017. ISSN 1932-6203.  
1364 doi:[10.1371/journal.pone.0184661](https://doi.org/10.1371/journal.pone.0184661). URL <http://dx.doi.org/10.1371/journal.pone.0184661>.
- 1365 Oscar Esteban, Christopher J. Markiewicz, Ross W. Blair, Craig A. Moodie, A. Ilkay Isik, Asier  
1366 Erramuzpe, James D. Kent, Mathias Goncalves, Elizabeth DuPre, Madeleine Snyder, and et al.  
1367 fMRIPrep: A robust preprocessing pipeline for functional MRI. *Nature Methods*, 16(1):111–116,  
1368 Dec 2018. ISSN 1548-7105. doi:[10.1038/s41592-018-0235-4](https://doi.org/10.1038/s41592-018-0235-4). URL <http://dx.doi.org/10.1038/s41592-018-0235-4>.  
1369 <http://dx.doi.org/10.1038/s41592-018-0235-4>.
- 1370 Oscar Esteban, Rastko Ciric, Karolina Finc, Ross Blair, Christopher J. Markiewicz, Craig A. Moodie,  
1371 James D. Kent, Mathias Goncalves, Elizabeth DuPre, Daniel E. P. Gomez, Zhifang Ye, Taylor Salo,  
1372 Romain Valabregue, Inge K. Amlien, Franziskus Liem, Nir Jacoby, Hrvoje Stojić, Matthew Cieslak,  
1373 Sebastian Urchs, Yaroslav O. Halchenko, Satrajit S. Ghosh, Alejandro De La Vega, Tal Yarkoni,  
1374 Jessey Wright, William H. Thompson, Russell A. Poldrack, and Krzysztof J. Gorgolewski. Analysis  
1375 of task-based functional MRI data preprocessed with fMRIPrep. *bioRxiv*, 2019a. doi:[10.1101/694364](https://doi.org/10.1101/694364).  
1376 URL <https://www.biorxiv.org/content/early/2019/07/08/694364>.

- 1377 Oscar Esteban, Christopher J. Markiewicz, Ross W. Blair, Craig A. Moodie, A. Ilkay Isik, Asier  
1378 Erramuzpe, James D. Kent, Mathias Goncalves, Elizabeth DuPre, Madeleine Snyder, and et al.  
1379 fMRIPrep 1.2.2., 2019b.
- 1380 Alan C. Evans, Andrew L. Janke, D. Louis Collins, and Sylvain Baillet. Brain  
1381 templates and atlases. *NeuroImage*, 62(2):911–922, 2012. ISSN 1053-8119.  
1382 doi:<https://doi.org/10.1016/j.neuroimage.2012.01.024>. URL <https://www.sciencedirect.com/science/article/pii/S1053811912000419>.
- 1384 Bruce Fischl, André van der Kouwe, Christophe Destrieux, Eric Halgren, Florent Ségonne, David H.  
1385 Salat, Evelina Busa, Larry J. Seidman, Jill Goldstein, David Kennedy, Verne Caviness, Nikos Makris,  
1386 Bruce Rosen, and Anders M. Dale. Automatically parcellating the human cerebral cortex. *Cerebral*  
1387 *Cortex*, 14(1):11–22, Jan 2004. ISSN 1460-2199. doi:[10.1093/cercor/bhg087](https://doi.org/10.1093/cercor/bhg087). URL <http://dx.doi.org/10.1093/cercor/bhg087>.
- 1389 VS Fonov, AC Evans, RC McKinstry, CR Almli, and DL Collins. Unbiased nonlinear average age-  
1390 appropriate brain templates from birth to adulthood. *NeuroImage*, 47:S102, Jul 2009. ISSN 1053-  
1391 8119. doi:[10.1016/s1053-8119\(09\)70884-5](https://doi.org/10.1016/s1053-8119(09)70884-5). URL [http://dx.doi.org/10.1016/S1053-8119\(09\)](http://dx.doi.org/10.1016/S1053-8119(09)70884-5)  
1392 [70884-5](http://dx.doi.org/10.1016/S1053-8119(09)70884-5).
- 1393 David J. Foster. Replay comes of age. *Annual Review of Neuroscience*, 40(1):581–602, 2017.  
1394 doi:[10.1146/annurev-neuro-072116-031538](https://doi.org/10.1146/annurev-neuro-072116-031538). URL [https://doi.org/10.1146/annurev-neuro-](https://doi.org/10.1146/annurev-neuro-072116-031538)  
1395 [072116-031538](https://doi.org/10.1146/annurev-neuro-072116-031538).
- 1396 David J. Foster and Matthew A. Wilson. Reverse replay of behavioural sequences in hippocam-  
1397 pal place cells during the awake state. *Nature*, 440(7084):680–683, Feb 2006. ISSN 1476-4687.  
1398 doi:[10.1038/nature04587](https://doi.org/10.1038/nature04587). URL <http://dx.doi.org/10.1038/nature04587>.
- 1399 David J. Foster and Matthew A. Wilson. Hippocampal theta sequences. *Hippocampus*, 17(11):  
1400 1093–1099, 2007. doi:[10.1002/hipo.20345](https://doi.org/10.1002/hipo.20345). URL [https://onlinelibrary.wiley.com/doi/abs/](https://onlinelibrary.wiley.com/doi/abs/10.1002/hipo.20345)  
1401 [10.1002/hipo.20345](https://onlinelibrary.wiley.com/doi/abs/10.1002/hipo.20345).
- 1402 Mona M Garvert, Raymond J Dolan, and Timothy EJ Behrens. A map of abstract relational  
1403 knowledge in the human hippocampal–entorhinal cortex. *eLife*, 6, Apr 2017. ISSN 2050-084X.  
1404 doi:[10.7554/elife.17086](https://doi.org/10.7554/elife.17086). URL <http://dx.doi.org/10.7554/eLife.17086>.
- 1405 Jeffrey P Gavornik and Mark F Bear. Learned spatiotemporal sequence recognition and prediction  
1406 in primary visual cortex. *Nature Neuroscience*, 17(5):732–737, 2014. doi:[10.1038/nn.3683](https://doi.org/10.1038/nn.3683). URL  
1407 <https://doi.org/10.1038/nn.3683>.
- 1408 Samuel J. Gershman, Christopher D. Moore, Michael T. Todd, Kenneth A. Norman, and Per B.  
1409 Sederberg. The successor representation and temporal context. *Neural Computation*, 24(6):1553–  
1410 1568, Jun 2012. ISSN 1530-888X. doi:[10.1162/neco\\_a\\_00282](https://doi.org/10.1162/neco_a_00282). URL [http://dx.doi.org/10.1162/](http://dx.doi.org/10.1162/NECO_a_00282)  
1411 [NECO\\_a\\_00282](http://dx.doi.org/10.1162/NECO_a_00282).
- 1412 Chris Gorgolewski, Nell Hardcastle, Teal Hobson-Lowther, David Nishikawa, Ross Blair, Stefan Ap-  
1413 pelhoff, Suyash, Constellates, Mainak Jas, Chris Holdgraf, Alexander Jones, Rohan Goyal, Robert  
1414 Oostenveld, Chris Markiewicz, Gregory Noack, Matthew Zito, Joke Durnez, Nicolas Traut, Mikael



- 1415 Naveau, Parul Sethi, Yaroslav Halchenko, Taylor Salo, Michael Hanke, Dimitri Papadopoulos Or-  
1416 fanos, Horea Christian, Franklin Feingold, Duncan Macleod, Dewarrn1, Brian Grass, and Adam  
1417 Thomas. bids-standard/bids-validator: 1.4.3, 2020. URL <https://zenodo.org/record/3688707>.
- 1418 Krzysztof J. Gorgolewski, Christopher D. Burns, Cindee Madison, Dav Clark, Yaroslav O. Halchenko,  
1419 Michael L. Waskom, and Satrajit S. Ghosh. Nipype: A flexible, lightweight and extensible neu-  
1420 roimaging data processing framework in Python. *Frontiers in Neuroinformatics*, 5, 2011. ISSN  
1421 1662-5196. doi:10.3389/fninf.2011.00013. URL <http://dx.doi.org/10.3389/fninf.2011.00013>.
- 1422 Krzysztof J. Gorgolewski, Tibor Auer, Vince D. Calhoun, R. Cameron Craddock, Samir Das, Eugene P.  
1423 Duff, Guillaume Flandin, Satrajit S. Ghosh, Tristan Glatard, Yaroslav O. Halchenko, and et al.  
1424 The brain imaging data structure, a format for organizing and describing outputs of neuroimaging  
1425 experiments. *Scientific Data*, 3(160044), Jun 2016. ISSN 2052-4463. doi:10.1038/sdata.2016.44.  
1426 URL <http://dx.doi.org/10.1038/sdata.2016.44>.
- 1427 Krzysztof J. Gorgolewski, Christopher D. Burns, Cindee Madison, Dav Clark, Yaroslav O. Halchenko,  
1428 Michael L. Waskom, and Satrajit S. Ghosh. Nipype, 2019.
- 1429 Douglas N. Greve and Bruce Fischl. Accurate and robust brain image alignment us-  
1430 ing boundary-based registration. *NeuroImage*, 48(1):63–72, Oct 2009. ISSN 1053-8119.  
1431 doi:10.1016/j.neuroimage.2009.06.060. URL <http://dx.doi.org/10.1016/j.neuroimage.2009.06.060>.
- 1432 06.060.
- 1433 Omer Faruk Gulban, Dylan Nielson, Russ Poldrack, John Lee, Chris Gorgolewski, Vanessasaurus, and  
1434 Satrajit Ghosh. poldracklab/pydeface: v2.0.0, 2019. URL <https://zenodo.org/record/3524400>.
- 1435 Anoopum S. Gupta, Matthijs A.A. van der Meer, David S. Touretzky, and Aaron David Redish. Hip-  
1436 pocampal replay is not a simple function of experience. *Neuron*, 65(5):695 – 705, 2010. ISSN  
1437 0896-6273. doi:10.1016/j.neuron.2010.01.034. URL <http://www.sciencedirect.com/science/article/pii/S0896627310000607>.
- 1438 article/pii/S0896627310000607.
- 1439 Yaroslav O. Halchenko, Michael Hanke, Benjamin Poldrack, Kyle Meyer, Debanjum Singh Solanky,  
1440 Gergana Alteva, Jason Gors, Dave MacFarlane, Christian Olaf Häusler, Taylor Olson, Alex Waite,  
1441 Alejandro De La Vega, Vanessa Sochat, Anisha Keshavan, Feilong Ma, Horea Christian, Jorrit  
1442 Poelen, Kusti Skytén, Matteo Visconti di Oleggio Castello, Nell Hardcastle, Torsten Stoeter, Vicky  
1443 C Lau, and Christopher J. Markiewicz. datalad/datalad 0.11.5, 2019. URL <https://zenodo.org/record/3233911>.
- 1444 record/3233911.
- 1445 Yaroslav O. Halchenko, Kyle Meyer, Benjamin Poldrack, Debanjum Singh Solanky, Adina S. Wag-  
1446 ner, Jason Gors, Dave MacFarlane, Dorian Pustina, Vanessa Sochat, Satrajit S. Ghosh, Christian  
1447 Mönch, Christopher J. Markiewicz, Laura Waite, Ilya Shlyakhter, Alejandro de la Vega, Soichi  
1448 Hayashi, Christian Olaf Häusler, Jean-Baptiste Poline, Tobias Kadelka, Kusti Skytén, Dorota  
1449 Jarecka, David Kennedy, Ted Strauss, Matt Cieslak, Peter Vavra, Horea-Ioan Ioanas, Robin Schnei-  
1450 der, Mika Pflüger, James V. Haxby, Simon B. Eickhoff, and Michael Hanke. DataLad: distributed  
1451 system for joint management of code, data, and their relationship. *Journal of Open Source Software*,  
1452 6(63):3262, 2021. doi:10.21105/joss.03262. URL <https://doi.org/10.21105/joss.03262>.

- 1453 L.M. Harrison, A. Duggins, and K.J. Friston. Encoding uncertainty in the hippocampus. *Neural*  
1454 *Networks*, 19(5):535–546, 2006. ISSN 0893-6080. doi:<https://doi.org/10.1016/j.neunet.2005.11.002>.  
1455 URL <https://www.sciencedirect.com/science/article/pii/S0893608006000025>.
- 1456 James V. Haxby, M. Ida Gobbini, Maura L. Furey, Alumit Ishai, Jennifer L. Schouten, and Pietro  
1457 Pietrini. Distributed and overlapping representations of faces and objects in ventral temporal cortex.  
1458 *Science*, 293(5539):2425–2430, Sep 2001. ISSN 1095-9203. doi:[10.1126/science.1063736](https://doi.org/10.1126/science.1063736). URL <http://dx.doi.org/10.1126/science.1063736>.  
1459
- 1460 Nicholas C Hindy, Felicia Y Ng, and Nicholas B Turk-Browne. Linking pattern completion in the  
1461 hippocampus to predictive coding in visual cortex. *Nature Neuroscience*, 19(5):665–667, Apr 2016.  
1462 ISSN 1546-1726. doi:[10.1038/nn.4284](https://doi.org/10.1038/nn.4284). URL <http://dx.doi.org/10.1038/nn.4284>.
- 1463 L. T. Hunt, N. D. Daw, P. Kaanders, M. A. MacIver, U. Mugan, E. Procyk, A. D. Redish, E. Russo,  
1464 J. Scholl, K. Stachenfeld, C. R. E. Wilson, and N. Kolling. Formalizing planning and information  
1465 search in naturalistic decision-making. *Nature Neuroscience*, 2021. doi:[10.1038/s41593-021-00866-w](https://doi.org/10.1038/s41593-021-00866-w).  
1466 URL <https://doi.org/10.1038/s41593-021-00866-w>.
- 1467 Mark Jenkinson, Peter Bannister, Michael Brady, and Stephen Smith. Improved optimization for the  
1468 robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2):  
1469 825–841, Oct 2002. ISSN 1053-8119. doi:[10.1006/nimg.2002.1132](https://doi.org/10.1006/nimg.2002.1132). URL <http://dx.doi.org/10.1006/nimg.2002.1132>.  
1470
- 1471 Daoyun Ji and Matthew A Wilson. Coordinated memory replay in the visual cortex and hippocampus  
1472 during sleep. *Nature Neuroscience*, 10(1):100–107, Dec 2006. ISSN 1546-1726. doi:[10.1038/nm1825](https://doi.org/10.1038/nm1825).  
1473 URL <http://dx.doi.org/10.1038/nm1825>.
- 1474 Adam Johnson and Aaron David Redish. Neural ensembles in CA3 transiently encode paths forward of  
1475 the animal at a decision point. *Journal of Neuroscience*, 27(45):12176–12189, Nov 2007. ISSN 1529-  
1476 2401. doi:[10.1523/jneurosci.3761-07.2007](https://doi.org/10.1523/jneurosci.3761-07.2007). URL <http://dx.doi.org/10.1523/JNEUROSCI.3761-07.2007>.  
1477
- 1478 Ari E. Kahn, Elisabeth A. Karuza, Jean M. Vettel, and Danielle S. Bassett. Network constraints on  
1479 learnability of probabilistic motor sequences. *Nature Human Behaviour*, 2(12):936–947, Nov 2018.  
1480 ISSN 2397-3374. doi:[10.1038/s41562-018-0463-8](https://doi.org/10.1038/s41562-018-0463-8). URL <http://dx.doi.org/10.1038/s41562-018-0463-8>.  
1481
- 1482 Raphael Kaplan, Adrià Tauste Campo, Daniel Bush, John King, Alessandro Principe, Raphael Koster,  
1483 Miguel Ley Nacher, Rodrigo Rocamora, and Karl J. Friston. Human hippocampal theta oscillations  
1484 reflect sequential dependencies during spatial planning. *Cognitive Neuroscience*, 11(3):122–  
1485 131, 2020. doi:[10.1080/17588928.2019.1676711](https://doi.org/10.1080/17588928.2019.1676711). URL <https://doi.org/10.1080/17588928.2019.1676711>. PMID: 31617790.  
1486
- 1487 Elisabeth A. Karuza, Sharon L. Thompson-Schill, and Danielle S. Bassett. Local patterns to  
1488 global architectures: Influences of network topology on human learning. *Trends in Cognitive*  
1489 *Sciences*, 20(8):629–640, 2016. ISSN 1364-6613. doi:[10.1016/j.tics.2016.06.003](https://doi.org/10.1016/j.tics.2016.06.003). URL <https://www.sciencedirect.com/science/article/pii/S1364661316300717>.  
1490

- 1491 Elisabeth A. Karuza, Ari E. Kahn, Sharon L. Thompson-Schill, and Danielle S. Bassett. Process  
1492 reveals structure: How a network is traversed mediates expectations about its architecture. *Scien-*  
1493 *tific Reports*, 7(1):12733, 2017. doi:[10.1038/s41598-017-12876-5](https://doi.org/10.1038/s41598-017-12876-5). URL <https://doi.org/10.1038/s41598-017-12876-5>.
- 1495 Elisabeth A. Karuza, Ari E. Kahn, and Danielle S. Bassett. Human sensitivity to community  
1496 structure is robust to topological variation. *Complexity*, 2019:1–8, Feb 2019. ISSN 1099-0526.  
1497 doi:[10.1155/2019/8379321](https://doi.org/10.1155/2019/8379321). URL <http://dx.doi.org/10.1155/2019/8379321>.
- 1498 Kenneth Kay, Jason E. Chung, Marielena Sosa, Jonathan S. Schor, Mattias P. Karlsson, Margaret C.  
1499 Larkin, Daniel F. Liu, and Loren M. Frank. Constant sub-second cycling between representations  
1500 of possible futures in the hippocampus. *Cell*, 180(3):552–567.e25, Jan 2020. ISSN 0092-8674.  
1501 doi:[10.1016/j.cell.2020.01.014](https://doi.org/10.1016/j.cell.2020.01.014). URL <http://dx.doi.org/10.1016/j.cell.2020.01.014>.
- 1502 Arno Klein, Satrajit S. Ghosh, Forrest S. Bao, Joachim Giard, Yrjö Häme, Eliezer Stavsky, Noah  
1503 Lee, Brian Rossa, Martin Reuter, Elias Chaibub Neto, and et al. Mindboggling morphometry  
1504 of human brains. *PLOS Computational Biology*, 13(2):e1005350, Feb 2017. ISSN 1553-7358.  
1505 doi:[10.1371/journal.pcbi.1005350](https://doi.org/10.1371/journal.pcbi.1005350). URL <http://dx.doi.org/10.1371/journal.pcbi.1005350>.
- 1506 Peter Kok and Nicholas B. Turk-Browne. Associative prediction of visual shape in the hip-  
1507 pocampus. *The Journal of Neuroscience*, 38(31):6888–6899, Jul 2018. ISSN 1529-2401.  
1508 doi:[10.1523/JNEUROSCI.0163-18.2018](https://doi.org/10.1523/JNEUROSCI.0163-18.2018). URL <http://dx.doi.org/10.1523/JNEUROSCI.0163-18.2018>.
- 1510 Peter Kok, Janneke F.M. Jehee, and Floris P. de Lange. Less is more: Expectation sharp-  
1511 ens representations in the primary visual cortex. *Neuron*, 75(2):265–270, 2012. ISSN 0896-  
1512 6273. doi:<https://doi.org/10.1016/j.neuron.2012.04.034>. URL <https://www.sciencedirect.com/science/article/pii/S0896627312004382>.
- 1514 Peter Kok, Michel F. Failing, and Floris P. de Lange. Prior expectations evoke stimulus templates  
1515 in the primary visual cortex. *Journal of Cognitive Neuroscience*, 26(7):1546–1554, 07 2014. ISSN  
1516 0898-929X. doi:[10.1162/jocn\\_a.00562](https://doi.org/10.1162/jocn_a.00562). URL [https://doi.org/10.1162/jocn\\_a.00562](https://doi.org/10.1162/jocn_a.00562).
- 1517 James Kolasinski, Tamar R. Makin, Saad Jbabdi, Stuart Clare, Charlotte J. Stagg, and  
1518 Heidi Johansen-Berg. Investigating the stability of fine-grain digit somatotopy in individ-  
1519 ual human participants. *Journal of Neuroscience*, 36(4):1113–1127, 2016. ISSN 0270-6474.  
1520 doi:[10.1523/JNEUROSCI.1742-15.2016](https://doi.org/10.1523/JNEUROSCI.1742-15.2016). URL <https://www.jneurosci.org/content/36/4/1113>.
- 1521 Lukas Kunz, Lorena Deuker, Hui Zhang, and Nikolai Axmacher. Chapter 26 - tracking human engrams  
1522 using multivariate analysis techniques. In Denise Manahan-Vaughan, editor, *Handbook of in Vivo*  
1523 *Neural Plasticity Techniques*, volume 28 of *Handbook of Behavioral Neuroscience*, chapter 26, pages  
1524 481–508. Elsevier, 2018. doi:<https://doi.org/10.1016/B978-0-12-812028-6.00026-4>. URL <https://www.sciencedirect.com/science/article/pii/B9780128120286000264>.
- 1526 Zeb Kurth-Nelson, Marcos Economides, Raymond J. Dolan, and Peter Dayan. Fast sequences  
1527 of non-spatial state representations in humans. *Neuron*, 91(1):194–204, 2016. ISSN 10974199.  
1528 doi:[10.1016/j.neuron.2016.05.028](https://doi.org/10.1016/j.neuron.2016.05.028). URL <http://dx.doi.org/10.1016/j.neuron.2016.05.028>.

- 1529 Gregory M. Kurtzer, Vanessa Sochat, and Michael W. Bauer. Singularity: Scientific containers for  
1530 mobility of compute. *PLoS ONE*, 12(5):e0177459, May 2017. doi:[10.1371/journal.pone.0177459](https://doi.org/10.1371/journal.pone.0177459).  
1531 URL <http://dx.doi.org/10.1371/journal.pone.0177459>.
- 1532 C. Lanczos. Evaluation of noisy data. *Journal of the Society for Industrial and Applied Mathematics*  
1533 *Series B Numerical Analysis*, 1(1):76–85, Jan 1964. ISSN 0887-459X. doi:[10.1137/0701007](https://doi.org/10.1137/0701007). URL  
1534 <http://dx.doi.org/10.1137/0701007>.
- 1535 Steven M. LaValle. *Planning Algorithms*. Cambridge University Press, Cambridge, 2006. ISBN  
1536 9780521862059. doi:[10.1017/CBO9780511546877](https://doi.org/10.1017/CBO9780511546877). URL [https://www.cambridge.org/core/  
1537 books/planning-algorithms/FC9CC7E67E851E40E3E45D6FE328B768](https://www.cambridge.org/core/books/planning-algorithms/FC9CC7E67E851E40E3E45D6FE328B768).
- 1538 Russell Lenth. emmeans: Estimated marginal means, aka least-squares means. 2019. URL [https:  
1539 //CRAN.R-project.org/package=emmeans](https://CRAN.R-project.org/package=emmeans). R package version 1.3.4.
- 1540 Xiangrui Li, Paul S. Morgan, John Ashburner, Jolinda Smith, and Christopher Rorden. The first step  
1541 for neuroimaging data analysis: Dicom to nifti conversion. *Journal of Neuroscience Methods*, 264:  
1542 47–56, May 2016. ISSN 0165-0270. doi:[10.1016/j.jneumeth.2016.03.001](https://doi.org/10.1016/j.jneumeth.2016.03.001). URL [http://dx.doi.org/  
1543 10.1016/j.jneumeth.2016.03.001](http://dx.doi.org/10.1016/j.jneumeth.2016.03.001).
- 1544 Yunzhe Liu, Raymond J. Dolan, Zeb Kurth-Nelson, and Timothy E.J. Behrens. Human re-  
1545 play spontaneously reorganizes experience. *Cell*, 178(3):640–652, Jul 2019. ISSN 0092-8674.  
1546 doi:[10.1016/j.cell.2019.06.012](https://doi.org/10.1016/j.cell.2019.06.012). URL <http://dx.doi.org/10.1016/j.cell.2019.06.012>.
- 1547 Yunzhe Liu, Marcelo G. Mattar, Timothy E. J. Behrens, Nathaniel D. Daw, and Raymond J. Dolan.  
1548 Experience replay is associated with efficient nonlocal learning. *Science*, 372(6544), 2021. ISSN  
1549 0036-8075. doi:[10.1126/science.abf1357](https://doi.org/10.1126/science.abf1357). URL [https://science.sciencemag.org/content/372/  
1550 6544/eabf1357](https://science.sciencemag.org/content/372/6544/eabf1357).
- 1551 Christopher W. Lynn and Danielle S. Bassett. How humans learn and represent networks. *Pro-  
1552 ceedings of the National Academy of Sciences*, 117(47):29407–29415, 2020. ISSN 0027-8424.  
1553 doi:[10.1073/pnas.1912328117](https://doi.org/10.1073/pnas.1912328117). URL <https://www.pnas.org/content/117/47/29407>.
- 1554 Christopher W. Lynn, Ari E. Kahn, Nathaniel Nyema, and Danielle S. Bassett. Abstract representa-  
1555 tions of events arise from mental errors in learning and memory. *Nature Communications*, 11(1),  
1556 May 2020a. ISSN 2041-1723. doi:[10.1038/s41467-020-15146-7](https://doi.org/10.1038/s41467-020-15146-7). URL [http://dx.doi.org/10.1038/  
1557 s41467-020-15146-7](http://dx.doi.org/10.1038/s41467-020-15146-7).
- 1558 Christopher W. Lynn, Lia Papadopoulos, Ari E. Kahn, and Danielle S. Bassett. Human information  
1559 processing in complex networks. *Nature Physics*, Jun 2020b. ISSN 1745-2481. doi:[10.1038/s41567-  
1560 020-0924-7](https://doi.org/10.1038/s41567-020-0924-7). URL <http://dx.doi.org/10.1038/s41567-020-0924-7>.
- 1561 Kevin J Miller and Sarah Jo C Venditto. Multi-step planning in the brain. *Current Opinion in  
1562 Behavioral Sciences*, 38:29–39, 2021. ISSN 2352-1546. doi:[10.1016/j.cobeha.2020.07.003](https://doi.org/10.1016/j.cobeha.2020.07.003). URL  
1563 <http://www.sciencedirect.com/science/article/pii/S2352154620301054>.
- 1564 Yasushi Miyashita. Neuronal correlate of visual associative long-term memory in the primate temporal  
1565 cortex. *Nature*, 335(6193):817–820, Oct 1988. ISSN 1476-4687. doi:[10.1038/335817a0](https://doi.org/10.1038/335817a0). URL [http:  
1566 //dx.doi.org/10.1038/335817a0](http://dx.doi.org/10.1038/335817a0).

- 1567 Ida Momennejad. Learning structures: Predictive representations, replay, and generaliza-  
1568 tion. *Current Opinion in Behavioral Sciences*, 32:155–166, Apr 2020. ISSN 2352-1546.  
1569 doi:10.1016/j.cobeha.2020.02.017. URL <http://dx.doi.org/10.1016/j.cobeha.2020.02.017>.
- 1570 Ida Momennejad and Marc W. Howard. Predicting the future with multi-scale successor representa-  
1571 tions. *bioRxiv*, 2018. doi:10.1101/449470. URL [https://www.biorxiv.org/content/early/2018/](https://www.biorxiv.org/content/early/2018/10/22/449470)  
1572 [10/22/449470](https://www.biorxiv.org/content/early/2018/10/22/449470).
- 1573 Ida Momennejad, Evan M. Russek, J. H. Cheong, Matthew M. Botvinick, Nathaniel D. Daw, and  
1574 Samuel J. Gershman. The successor representation in human reinforcement learning. *Nature Human*  
1575 *Behaviour*, 1(9):680–692, Aug 2017. ISSN 2397-3374. doi:10.1038/s41562-017-0180-8. URL [http://](http://dx.doi.org/10.1038/s41562-017-0180-8)  
1576 [dx.doi.org/10.1038/s41562-017-0180-8](http://dx.doi.org/10.1038/s41562-017-0180-8).
- 1577 Ida Momennejad, A Ross Otto, Nathaniel D Daw, and Kenneth A Norman. Offline replay supports  
1578 planning in human reinforcement learning. *eLife*, 7:e32548, Dec 2018. doi:10.7554/eLife.32548. URL  
1579 <https://doi.org/10.7554/eLife.32548>.
- 1580 H. Freyja Ólafsdóttir, Daniel Bush, and Caswell Barry. The role of hippocampal replay  
1581 in memory and planning. *Current Biology*, 28(1):R37–R50, Jan 2018. ISSN 0960-9822.  
1582 doi:10.1016/j.cub.2017.10.073. URL <http://dx.doi.org/10.1016/j.cub.2017.10.073>.
- 1583 Fabian Pedregosa, Gael Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier  
1584 Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas,  
1585 Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Edouard Duchesnay.  
1586 Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830,  
1587 2011.
- 1588 Jonathan Peirce, Jeremy R. Gray, Sol Simpson, Michael MacAskill, Richard Höchenberger, Hiroyuki  
1589 Sogo, Erik Kastman, and Jonas Kristoffer Lindeløv. Psychopy2: Experiments in behavior made  
1590 easy. *Behavior Research Methods*, 51(1):195–203, Feb 2019. ISSN 1554-3528. doi:10.3758/s13428-  
1591 [018-01193-y](http://dx.doi.org/10.3758/s13428-018-01193-y). URL <http://dx.doi.org/10.3758/s13428-018-01193-y>.
- 1592 Jonathan W. Peirce. PsychoPy—psychophysics software in python. *Journal of Neuroscience Methods*,  
1593 162(1-2):8–13, may 2007. doi:10.1016/j.jneumeth.2006.11.017. URL [https://doi.org/10.1016%](https://doi.org/10.1016%2Fj.jneumeth.2006.11.017)  
1594 [2Fj.jneumeth.2006.11.017](https://doi.org/10.1016%2Fj.jneumeth.2006.11.017).
- 1595 Jonathan W Peirce. Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformat-*  
1596 *ics*, 2, 2008. doi:10.3389/neuro.11.010.2008. URL [https://doi.org/10.3389%](https://doi.org/10.3389%2Fneuro.11.010.2008)  
1597 [2Fneuro.11.010.2008](https://doi.org/10.3389%2Fneuro.11.010.2008).
- 1598 Brad E. Pfeiffer and David J. Foster. Hippocampal place-cell sequences depict future paths to remem-  
1599 bered goals. *Nature*, 497(7447):74–79, Apr 2013. ISSN 1476-4687. doi:10.1038/nature12112. URL  
1600 <http://dx.doi.org/10.1038/nature12112>.
- 1601 Russell A. Poldrack. Region of interest analysis for fMRI. *Social Cognitive and Affective Neuroscience*,  
1602 2(1):67–70, Mar 2007. ISSN 1749-5024. doi:10.1093/scan/nsm006. URL [http://dx.doi.org/10.](http://dx.doi.org/10.1093/scan/nsm006)  
1603 [1093/scan/nsm006](http://dx.doi.org/10.1093/scan/nsm006).
- 1604 Michael J. D. Powell. Developments of newuoa for unconstrained minimization without derivatives.  
1605 *Department of Applied Mathematics and Theoretical Physics*, 2007.

- 1606 Michael J. D. Powell. The bobyqa algorithm for bound constrained optimization without derivatives.  
1607 *Department of Applied Mathematics and Theoretical Physics*, pages 26–46, 2009.
- 1608 Jonathan D. Power, Anish Mitra, Timothy O. Laumann, Abraham Z. Snyder, Bradley L. Schlaggar,  
1609 and Steven E. Petersen. Methods to detect, characterize, and remove motion artifact in resting state  
1610 fmri. *NeuroImage*, 84:320–341, Jan 2014. ISSN 1053-8119. doi:[10.1016/j.neuroimage.2013.08.048](https://doi.org/10.1016/j.neuroimage.2013.08.048).  
1611 URL <http://dx.doi.org/10.1016/j.neuroimage.2013.08.048>.
- 1612 R Core Team. R: A language and environment for statistical computing, 2019. URL [https://www.R-](https://www.R-project.org/)  
1613 [project.org/](https://www.R-project.org/).
- 1614 Arthur S. Reber. Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*,  
1615 118(3):219–235, 1989. ISSN 0096-3445. doi:[10.1037/0096-3445.118.3.219](https://doi.org/10.1037/0096-3445.118.3.219). URL [http://dx.doi.](http://dx.doi.org/10.1037/0096-3445.118.3.219)  
1616 [org/10.1037/0096-3445.118.3.219](http://dx.doi.org/10.1037/0096-3445.118.3.219).
- 1617 Martin Reuter, H. Diana Rosas, and Bruce Fischl. Highly accurate inverse consistent reg-  
1618 istration: A robust approach. *NeuroImage*, 53(4):1181–1196, Dec 2010. ISSN 1053-8119.  
1619 doi:[10.1016/j.neuroimage.2010.07.020](https://doi.org/10.1016/j.neuroimage.2010.07.020). URL [http://dx.doi.org/10.1016/j.neuroimage.2010.](http://dx.doi.org/10.1016/j.neuroimage.2010.07.020)  
1620 [07.020](http://dx.doi.org/10.1016/j.neuroimage.2010.07.020).
- 1621 Bruno Rossion and Gilles Pourtois. Revisiting Snodgrass and Vanderwart’s object pictorial set: The  
1622 role of surface detail in basic-level object recognition. *Perception*, 33(2):217–236, Feb 2004. ISSN  
1623 1468-4233. doi:[10.1068/p5117](https://doi.org/10.1068/p5117). URL <http://dx.doi.org/10.1068/p5117>.
- 1624 Evan M. Russek, Ida Momennejad, Matthew M. Botvinick, Samuel J. Gershman, and Nathaniel D.  
1625 Daw. Predictive representations can link model-based reinforcement learning to model-free  
1626 mechanisms. *PLoS Computational Biology*, 13(9):e1005768, Sep 2017. ISSN 1553-7358.  
1627 doi:[10.1371/journal.pcbi.1005768](https://doi.org/10.1371/journal.pcbi.1005768). URL <http://dx.doi.org/10.1371/journal.pcbi.1005768>.
- 1628 Evan M. Russek, Ida Momennejad, Matthew M. Botvinick, Samuel J. Gershman, and Nathaniel D.  
1629 Daw. Neural evidence for the successor representation in choice evaluation. *bioRxiv*, 2021.  
1630 doi:[10.1101/2021.08.29.458114](https://doi.org/10.1101/2021.08.29.458114). URL [https://www.biorxiv.org/content/early/2021/08/31/](https://www.biorxiv.org/content/early/2021/08/31/2021.08.29.458114)  
1631 [2021.08.29.458114](https://www.biorxiv.org/content/early/2021/08/31/2021.08.29.458114).
- 1632 J. R. Saffran, R. N. Aslin, and E. L. Newport. Statistical learning by 8-month-old infants. *Science*,  
1633 274(5294):1926–1928, Dec 1996. ISSN 1095-9203. doi:[10.1126/science.274.5294.1926](https://doi.org/10.1126/science.274.5294.1926). URL [http:](http://dx.doi.org/10.1126/science.274.5294.1926)  
1634 [//dx.doi.org/10.1126/science.274.5294.1926](http://dx.doi.org/10.1126/science.274.5294.1926).
- 1635 Theodore D. Satterthwaite, Mark A. Elliott, Raphael T. Gerraty, Kosha Ruparel, James Loughhead,  
1636 Monica E. Calkins, Simon B. Eickhoff, Hakon Hakonarson, Ruben C. Gur, Raquel E. Gur, and  
1637 Daniel H. Wolf. An improved framework for confound regression and filtering for control of motion  
1638 artifact in the preprocessing of resting-state functional connectivity data. *NeuroImage*, 64:240–  
1639 256, 2013. ISSN 1053-8119. doi:<https://doi.org/10.1016/j.neuroimage.2012.08.052>. URL [https:](https://www.sciencedirect.com/science/article/pii/S1053811912008609)  
1640 [//www.sciencedirect.com/science/article/pii/S1053811912008609](https://www.sciencedirect.com/science/article/pii/S1053811912008609).
- 1641 Anna C. Schapiro and N. Turk-Browne. Statistical learning. In Arthur W. Toga, editor, *Brain*  
1642 *Mapping*, volume 3, pages 501–506. Elsevier, 2015. ISBN 9780123973160. doi:[10.1016/B978-0-12-](https://doi.org/10.1016/B978-0-12-397025-1.00276-1)  
1643 [397025-1.00276-1](https://doi.org/10.1016/B978-0-12-397025-1.00276-1). URL <http://dx.doi.org/10.1016/B978-0-12-397025-1.00276-1>.

- 1644 Anna C. Schapiro, Lauren V. Kustner, and Nicholas B. Turk-Browne. Shaping of object representations  
1645 in the human medial temporal lobe based on temporal regularities. *Current Biology*, 22(17):1622–  
1646 1627, Sep 2012. ISSN 0960-9822. doi:[10.1016/j.cub.2012.06.056](https://doi.org/10.1016/j.cub.2012.06.056). URL <http://dx.doi.org/10.1016/j.cub.2012.06.056>.
- 1648 Anna C Schapiro, Timothy T Rogers, Natalia I Cordova, Nicholas B Turk-Browne, and Matthew M  
1649 Botvinick. Neural representations of events arise from temporal community structure. *Nature*  
1650 *Neuroscience*, 16(4):486–492, Feb 2013. ISSN 1546-1726. doi:[10.1038/nn.3331](https://doi.org/10.1038/nn.3331). URL <http://dx.doi.org/10.1038/nn.3331>.
- 1652 Nicolas W Schuck and Yael Niv. Sequential replay of nonspatial task states in the human hippocampus.  
1653 *Science*, 364(6447):eaaw5181, 2019. doi:[10.1126/science.aaw5181](https://doi.org/10.1126/science.aaw5181).
- 1654 Nicolas W Schuck, Robert Gaschler, and Peter A Frensch. Implicit learning of what comes  
1655 when and where within a sequence: The time-course of acquiring serial position-item and item-  
1656 item associations to represent serial order. *Advances in cognitive psychology*, 8(2):83–97, 2012a.  
1657 doi:[10.2478/v10053-008-0106-0](https://doi.org/10.2478/v10053-008-0106-0). URL <https://pubmed.ncbi.nlm.nih.gov/22679464>.
- 1658 Nicolas W. Schuck, Robert Gaschler, Aysha Keisler, and Peter A. Frensch. Position–item associ-  
1659 ations play a role in the acquisition of order knowledge in an implicit serial reaction time task.  
1660 *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(2):440–456, 2012b.  
1661 doi:[10.1037/a0025816](https://doi.org/10.1037/a0025816). URL <https://doi.org/10.1037/a0025816>.
- 1662 Nicolas W. Schuck, Ming Bo Cai, Robert C. Wilson, and Yael Niv. Human orbitofrontal cortex  
1663 represents a cognitive map of state space. *Neuron*, 91(6):1402 – 1412, 2016. ISSN 0896-6273.  
1664 doi:[10.1016/j.neuron.2016.08.019](https://doi.org/10.1016/j.neuron.2016.08.019). URL <http://www.sciencedirect.com/science/article/pii/S0896627316305116>.
- 1666 Carol Augart Seger. Implicit learning. *Psychological Bulletin*, 115(2):163–196, 1994. ISSN 0033-2909.  
1667 doi:[10.1037/0033-2909.115.2.163](https://doi.org/10.1037/0033-2909.115.2.163). URL <http://dx.doi.org/10.1037/0033-2909.115.2.163>.
- 1668 C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):  
1669 379–423, 1948. doi:[10.1002/j.1538-7305.1948.tb01338.x](https://doi.org/10.1002/j.1538-7305.1948.tb01338.x).
- 1670 Brynn E Sherman, Kathryn N Graves, and Nicholas B Turk-Browne. The prevalence and importance  
1671 of statistical learning in human cognition and behavior. *Current Opinion in Behavioral Sciences*,  
1672 32:15–20, Apr 2020. ISSN 2352-1546. doi:[10.1016/j.cobeha.2020.01.015](https://doi.org/10.1016/j.cobeha.2020.01.015). URL <http://dx.doi.org/10.1016/j.cobeha.2020.01.015>.
- 1674 Stephen M. Smith and J. Michael Brady. SUSAN - a new approach to low level image pro-  
1675 cessing. *International Journal of Computer Vision*, 23(1):45–78, May 1997. ISSN 0920-5691.  
1676 doi:[10.1023/a:1007963824710](https://doi.org/10.1023/a:1007963824710). URL <http://dx.doi.org/10.1023/A:1007963824710>.
- 1677 Joan G Snodgrass and Mary Vanderwart. A standardized set of 260 pictures: norms for name  
1678 agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psy-*  
1679 *chology: Human learning and memory*, 6(2):174–215, 1980. doi:[10.1037/0278-7393.6.2.174](https://doi.org/10.1037/0278-7393.6.2.174). URL  
1680 <https://doi.org/10.1037/0278-7393.6.2.174>.
- 1681 Vanessa V. Sochat, Cameron J. Prybol, and Gregory M. Kurtzer. Enhancing reproducibility in sci-  
1682 entific computing: Metrics and registry for singularity containers. *PLoS ONE*, 12(11):e0188511,

- 1683 Nov 2017. doi:[10.1371/journal.pone.0188511](https://doi.org/10.1371/journal.pone.0188511). URL [http://dx.doi.org/10.1371/journal.pone.](http://dx.doi.org/10.1371/journal.pone.0188511)  
1684 [0188511](http://dx.doi.org/10.1371/journal.pone.0188511).
- 1685 Alec Solway, Carlos Diuk, Natalia Córdova, Debbie Yee, Andrew G. Barto, Yael Niv, and Matthew M.  
1686 Botvinick. Optimal behavioral hierarchy. *PLoS Computational Biology*, 10(8):1–10, 08 2014.  
1687 doi:[10.1371/journal.pcbi.1003779](https://doi.org/10.1371/journal.pcbi.1003779). URL <https://doi.org/10.1371/journal.pcbi.1003779>.
- 1688 Kimberly L Stachenfeld, Matthew M Botvinick, and Samuel J Gershman. The hippocampus  
1689 as a predictive map. *Nature Neuroscience*, 20(11):1643–1653, Oct 2017. ISSN 1546-1726.  
1690 doi:[10.1038/nn.4650](https://doi.org/10.1038/nn.4650). URL <http://dx.doi.org/10.1038/nn.4650>.
- 1691 Bryan A. Strange, Andrew Duggins, William Penny, Raymond J. Dolan, and Karl J. Friston. Infor-  
1692 mation theory, novelty and hippocampal responses: unpredicted or unpredictable? *Neural Net-*  
1693 *works*, 18(3):225–230, 2005. ISSN 0893-6080. doi:<https://doi.org/10.1016/j.neunet.2004.12.004>.  
1694 URL <https://www.sciencedirect.com/science/article/pii/S0893608005000067>.
- 1695 Richard S. Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM*  
1696 *SIGART Bulletin*, 2(4):160–163, Jul 1991. ISSN 0163-5719. doi:[10.1145/122344.122377](https://doi.org/10.1145/122344.122377). URL  
1697 <http://dx.doi.org/10.1145/122344.122377>.
- 1698 Arielle Tambini and Lila Davachi. Persistence of hippocampal multivoxel patterns into postencoding  
1699 rest is related to memory. *Proceedings of the National Academy of Sciences*, 110(48):19591–19596,  
1700 Nov 2013. ISSN 1091-6490. doi:[10.1073/pnas.1308499110](https://doi.org/10.1073/pnas.1308499110). URL <http://dx.doi.org/10.1073/pnas.1308499110>.
- 1702 Arielle Tambini and Lila Davachi. Awake reactivation of prior experiences consolidates memories  
1703 and biases cognition. *Trends in Cognitive Sciences*, 23(10):876–890, Oct 2019. ISSN 1364-6613.  
1704 doi:[10.1016/j.tics.2019.07.008](https://doi.org/10.1016/j.tics.2019.07.008). URL <http://dx.doi.org/10.1016/j.tics.2019.07.008>.
- 1705 Edward C. Tolman. Cognitive maps in rats and men. *Psychological Review*, 55(4):189–208, 1948.  
1706 ISSN 0033-295X. doi:[10.1037/h0061626](https://doi.org/10.1037/h0061626). URL <http://dx.doi.org/10.1037/h0061626>.
- 1707 John W. Tukey. Comparing individual means in the analysis of variance. *Biometrics*, 5(2):99–114,  
1708 Jun 1949. ISSN 0006-341X. doi:[10.2307/3001913](https://doi.org/10.2307/3001913). URL <http://dx.doi.org/10.2307/3001913>.
- 1709 Nicholas B. Turk-Browne, Justin A. Jungé, and Brian J. Scholl. The automaticity of visual statistical  
1710 learning. *Journal of Experimental Psychology: General*, 134(4):552–564, 2005. ISSN 0096-3445.  
1711 doi:[10.1037/0096-3445.134.4.552](https://doi.org/10.1037/0096-3445.134.4.552). URL <http://dx.doi.org/10.1037/0096-3445.134.4.552>.
- 1712 Nicholas J Tustison, Brian B Avants, Philip A Cook, Yuanjie Zheng, Alexander Egan, Paul A  
1713 Yushkevich, and James C Gee. N4itk: Improved n3 bias correction. *IEEE Transactions on Med-*  
1714 *ical Imaging*, 29(6):1310–1320, Jun 2010. ISSN 1558-254X. doi:[10.1109/tmi.2010.2046908](https://doi.org/10.1109/tmi.2010.2046908). URL  
1715 <http://dx.doi.org/10.1109/TMI.2010.2046908>.
- 1716 Matthijs A A. van der Meer and Aaron David Redish. Covert expectation-of-reward in rat ven-  
1717 tral striatum at decision points. *Frontiers in Integrative Neuroscience*, 3, 2009. ISSN 1662-5145.  
1718 doi:[10.3389/neuro.07.001.2009](https://doi.org/10.3389/neuro.07.001.2009). URL <http://dx.doi.org/10.3389/neuro.07.001.2009>.
- 1719 Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA,  
1720 2009. ISBN 1441412697.



- 1721 Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau,  
1722 Evgeni Burovski, Pearu Peterson, Warren Weckesser, and et al. Scipy 1.0: fundamental algorithms  
1723 for scientific computing in python. *Nature Methods*, Feb 2020. ISSN 1548-7105. doi:[10.1038/s41592-  
1724 019-0686-2](https://doi.org/10.1038/s41592-019-0686-2). URL <http://dx.doi.org/10.1038/s41592-019-0686-2>.
- 1725 Matteo Visconti di Oleggio Castello, James E. Dobson, Terry Sackett, Chandana Kodiweera, James V.  
1726 Haxby, Mathias Goncalves, Satrajit Ghosh, and Yaroslav O. Halchenko. Repronim/reproim 0.6.0,  
1727 2020. URL <https://zenodo.org/record/3625000>.
- 1728 Adina S. Wagner, Laura K. Waite, Kyle Meyer, Marisa K. Heckner, Tobias Kadelka, Niels Reuter,  
1729 Alexander Q. Waite, Benjamin Poldrack, Christopher J. Markiewicz, Yaroslav O. Halchenko, Peter  
1730 Vavra, Pattarawat Chormai, Jean-Baptiste Poline, Lya K. Paas, Peer Herholz, Lisa N. Mochalski,  
1731 Nevena Kraljevic, Lisa Wiersch, Alexandre Hutton, Dorian Pustina, Hamzah Hamid Baagil, Tristan  
1732 Glatard, Sarah Oliveira, Giulia Ippoliti, Christian Mönch, Dorien Huijser, and Michael Hanke. *The  
1733 DataLad Handbook*. Zenodo, 2020. doi:[10.5281/ZENODO.3905791](https://doi.org/10.5281/ZENODO.3905791). URL [https://zenodo.org/  
1734 record/3905791](https://zenodo.org/record/3905791).
- 1735 Mengni Wang, David J. Foster, and Brad E. Pfeiffer. Alternating sequences of future and past  
1736 behavior encoded within hippocampal theta oscillations. *Science*, 370(6513):247–250, 2020.  
1737 doi:[10.1126/science.abb4151](https://doi.org/10.1126/science.abb4151). URL <https://science.sciencemag.org/content/370/6513/247>.
- 1738 Nikolaus Weiskopf, Chloe Hutton, Oliver Josephs, and Ralf Deichmann. Optimal EPI parameters for  
1739 reduction of susceptibility-induced BOLD sensitivity losses: A whole-brain analysis at 3 T and 1.5  
1740 T. *NeuroImage*, 33(2):493–504, Nov 2006. ISSN 1053-8119. doi:[10.1016/j.neuroimage.2006.07.029](https://doi.org/10.1016/j.neuroimage.2006.07.029).  
1741 URL <http://dx.doi.org/10.1016/j.neuroimage.2006.07.029>.
- 1742 Andrew M Wikenheiser and Aaron David Redish. Decoding the cognitive map: ensemble hippocampal  
1743 sequences and decision making. *Current Opinion in Neurobiology*, 32:8–15, Jun 2015a. ISSN 0959-  
1744 4388. doi:[10.1016/j.conb.2014.10.002](https://doi.org/10.1016/j.conb.2014.10.002). URL <http://dx.doi.org/10.1016/j.conb.2014.10.002>.
- 1745 Andrew M. Wikenheiser and Aaron David Redish. Hippocampal theta sequences reflect current goals.  
1746 *Nature Neuroscience*, 18(2):289–294, 2015b. doi:[10.1038/nn.3909](https://doi.org/10.1038/nn.3909). URL [https://doi.org/10.  
1747 1038/nn.3909](https://doi.org/10.1038/nn.3909).
- 1748 Robert C. Wilson, Yuji K. Takahashi, Geoffrey Schoenbaum, and Yael Niv. Orbitofrontal  
1749 cortex as a cognitive map of task space. *Neuron*, 81(2):267–279, 2014. ISSN 0896-6273.  
1750 doi:[10.1016/j.neuron.2013.11.005](https://doi.org/10.1016/j.neuron.2013.11.005). URL [http://www.sciencedirect.com/science/article/pii/  
1751 S0896627313010398](http://www.sciencedirect.com/science/article/pii/S0896627313010398).
- 1752 Lennart Wittkuhn and Nicolas W. Schuck. Dynamics of fMRI patterns reflect sub-second activation  
1753 sequences and reveal replay in human visual cortex. *Nature Communications*, 12(1795), 2021.  
1754 doi:[10.1038/s41467-021-21970-2](https://doi.org/10.1038/s41467-021-21970-2). URL <https://doi.org/10.1038/s41467-021-21970-2>.
- 1755 Lennart Wittkuhn, Samson Chien, Sam Hall-McMaster, and Nicolas W. Schuck. Replay in minds  
1756 and machines. *Neuroscience & Biobehavioral Reviews*, 129:367–388, 2021. ISSN 0149-7634.  
1757 doi:[10.1016/j.neubiorev.2021.08.002](https://doi.org/10.1016/j.neubiorev.2021.08.002). URL [https://www.sciencedirect.com/science/article/  
1758 pii/S0149763421003444](https://www.sciencedirect.com/science/article/pii/S0149763421003444).

- 1759 Shengjin Xu, Wanchen Jiang, Mu-ming Poo, and Yang Dan. Activity recall in a visual cortical  
1760 ensemble. *Nature Neuroscience*, 15(3):449–455, Jan 2012. ISSN 1546-1726. doi:[10.1038/nm.3036](https://doi.org/10.1038/nm.3036).  
1761 URL <http://dx.doi.org/10.1038/nm.3036>.
- 1762 Tal Yarkoni, Christopher Markiewicz, Alejandro de la Vega, Krzysztof Gorgolewski, Taylor Salo,  
1763 Yaroslav Halchenko, Quinten McNamara, Krista DeStasio, Jean-Baptiste Poline, Dmitry Petrov,  
1764 Valérie Hayot-Sasson, Dylan Nielson, Johan Carlin, Gregory Kiar, Kirstie Whitaker, Elizabeth  
1765 DuPre, Adina Wagner, Lee Tirrell, Mainak Jas, Michael Hanke, Russell Poldrack, Oscar Esteban,  
1766 Stefan Appelhoff, Chris Holdgraf, Isla Staden, Bertrand Thirion, Dave Kleinschmidt, John Lee,  
1767 Matteo di Castello, Michael Notter, and Ross Blair. PyBIDS: Python tools for BIDS datasets.  
1768 *Journal of Open Source Software*, 4(40):1294, aug 2019a. doi:[10.21105/joss.01294](https://doi.org/10.21105/joss.01294). URL <https://doi.org/10.21105/joss.01294>.  
1769 <https://doi.org/10.21105/joss.01294>.
- 1770 Tal Yarkoni, Christopher J. Markiewicz, Alejandro de la Vega, Krzysztof J. Gorgolewski, Yaroslav O.  
1771 Halchenko, Taylor Salo, Quinten McNamara, Krista DeStasio, Jean-Baptiste Poline, Dmitry Petrov,  
1772 Valérie Hayot-Sasson, Dylan M. Nielson, Johan Carlin, Gregory Kiar, Kirstie Whitaker, Adina  
1773 Wagner, Elizabeth DuPre, Stefan Appelhoff, Alexander Ivanov, Johannes Wennberg, Lee S. Tirrell,  
1774 Oscar Esteban, Mainak Jas, Michael Hanke, Russell Poldrack, Chris Holdgraf, Isla Staden, Ariel  
1775 Rokem, Bertrand Thirion, Chadwick Boulay, Dave F. Kleinschmidt, Erin W Dickie, John A. Lee,  
1776 Matteo Visconti di Oleggio Castello, Michael Philipp Notter, Pauline Roca, and Ross Blair. bids-  
1777 standard/pybids: 0.9.3, 2019b. URL <https://zenodo.org/record/3363985>.
- 1778 Linda Q. Yu, Robert C. Wilson, and Matthew R. Nassar. Adaptive learning is structure  
1779 learning in time. *Neuroscience & Biobehavioral Reviews*, 128:270–281, 2021. ISSN 0149-  
1780 7634. doi:<https://doi.org/10.1016/j.neubiorev.2021.06.024>. URL <https://www.sciencedirect.com/science/article/pii/S0149763421002657>.  
1781 <https://www.sciencedirect.com/science/article/pii/S0149763421002657>.
- 1782 Y. Zhang, M. Brady, and S. Smith. Segmentation of brain MR images through a hidden markov  
1783 random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical  
1784 Imaging*, 20(1):45–57, 2001. ISSN 0278-0062. doi:[10.1109/42.906424](https://doi.org/10.1109/42.906424). URL [http://dx.doi.org/  
1785 10.1109/42.906424](http://dx.doi.org/10.1109/42.906424).

## 1786 Glossary

1787 **AC-PC** anterior commissure - posterior commissure.

1788 **AIC** Akaike information criterion.

1789 **ANOVA** analysis of variance.

1790 **ANTs** Advanced Normalization Tools.

1791 **A-P** anterior-to-posterior.

1792 **BIDS** brain imaging data structure.

1793 **BOBYQA** Bound Optimization BY Quadratic Approximation.

1794 **BOLD** blood-oxygen-level dependent.

1795 **CSF** cerebrospinal fluid.

1796 **DGPs** German Psychological Society.

1797 **DICOM** Digital Imaging and Communications in Medicine.

1798 **EPI** echo-planar imaging.

1799 **FA** flip angle.

1800 **FD** framewise displacement.

1801 **FDR** false discovery rate.

1802 **fMRI** functional magnetic resonance imaging.

1803 **FOV** field of view.

1804 **FSL** FMRIB Software Library.

1805 **FWHM** Full Width at Half Maximum.

1806 **GM** gray-matter.

1807 **GR** gradient recalled.

1808 **HMM** hidden markov model.

1809 **HPC** high performance computing.

1810 **HRF** The hemodynamic response function (HRF) characterizes an fMRI response that results from  
1811 a brief, spatially localized pulse of neuronal activity.

1812 **HSD** honest significant difference.

1813 **INU** intensity non-uniformity.

1814 **IQR** interquartile range.

1815 **ITI** inter-trial interval.

1816 **LME** linear mixed effects.

1817 **LTS** long-term support.

1818 **MB** multi-band.

1819 **MEG** magnetoencephalography.

1820 **min** minute.

1821 **MPRAGE** Magnetization Prepared Rapid Gradient Echo.

1822 **MRI** magnetic resonance imaging.

1823 **ms** millisecond.

1824 **MTL** medial temporal lobe.

1825 **PFC** prefrontal cortex.

1826 **PMU** Physiological Measurement Unit.

1827 **ROI** region of interest.

1828 **s** second.

1829 **SEM** standard error of the mean.

1830 **SI** supplementary information.

1831 **SR** successor representation.

1832 **S-R** stimulus-response.

1833 **SRI** stimulus-response interval.

1834 **SUSAN** Smallest Univalve Segment Assimilating Nucleus.

1835 **T1w** T1-weighted.

1836 **TD** temporal difference.

1837 **TE** echo time.

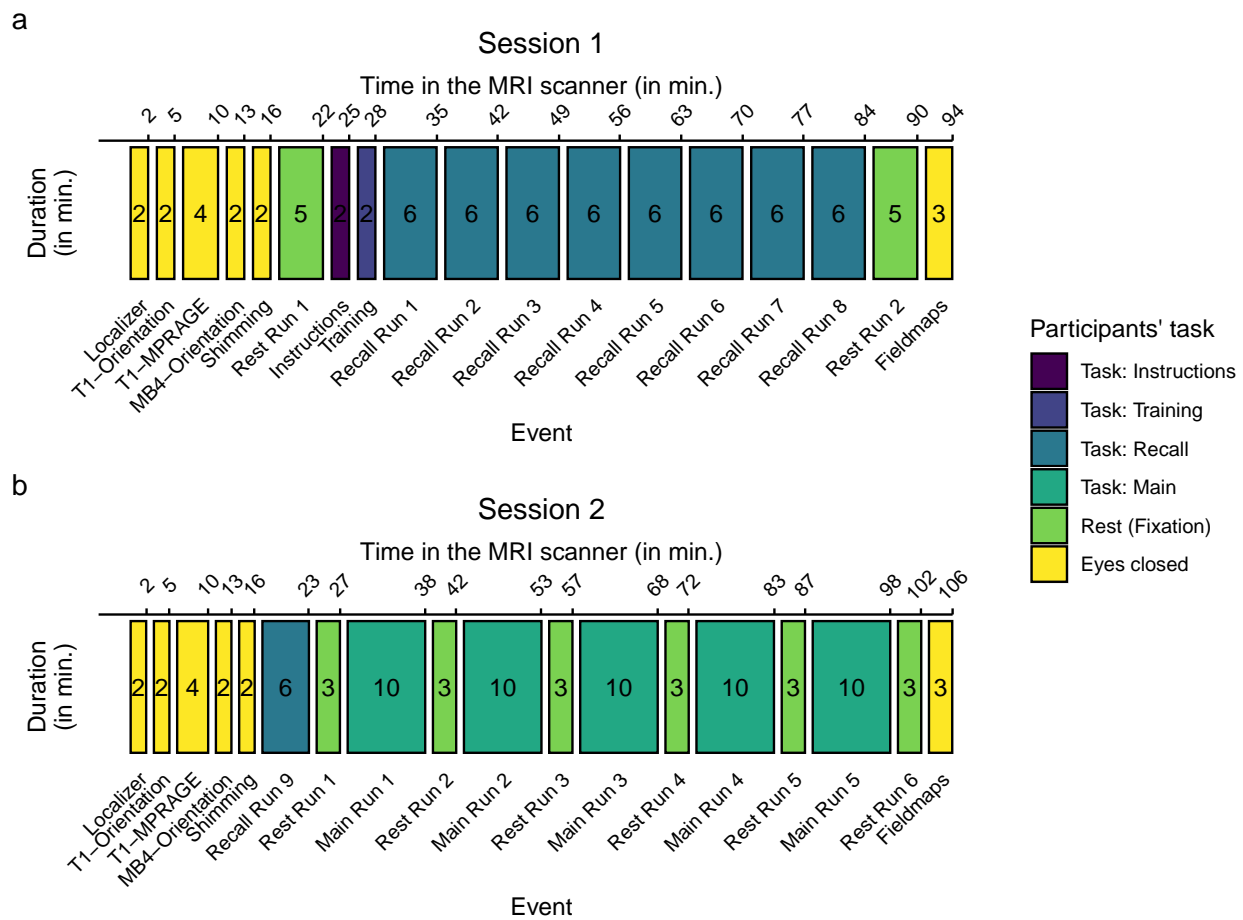
1838 **TI** inversion time.

1839 **TR** repetition time.

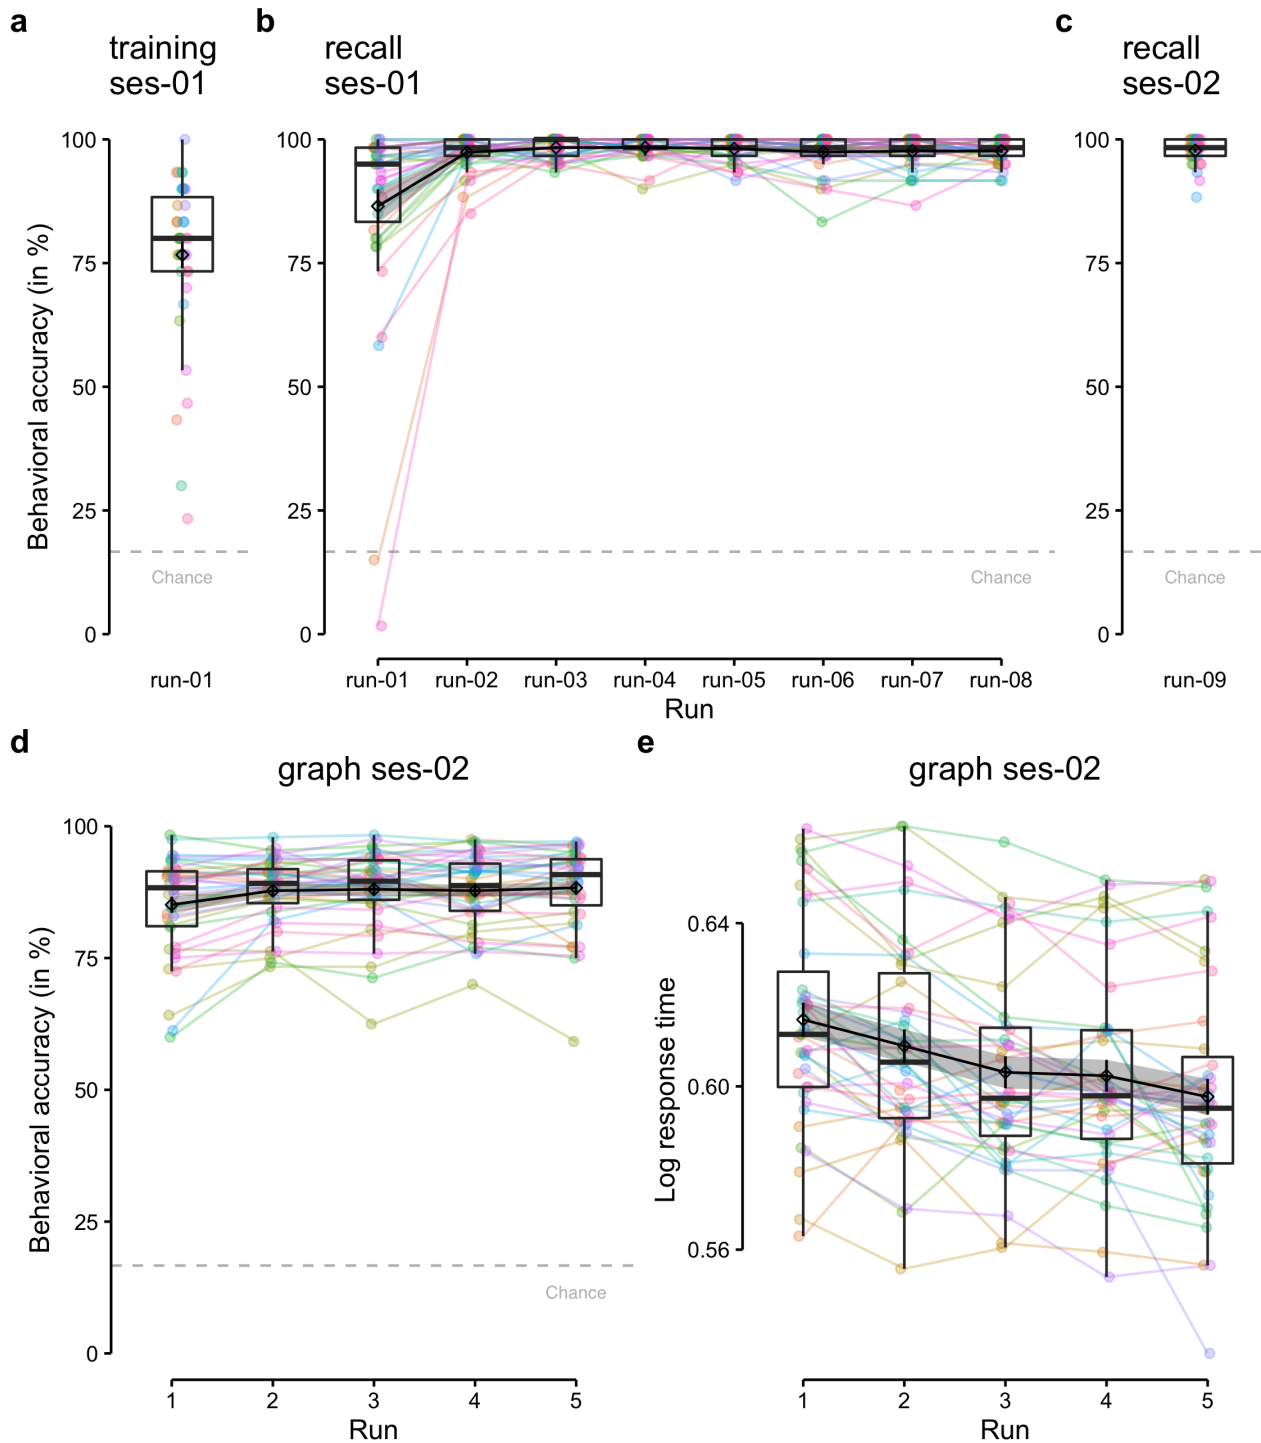
1840 **WM** white-matter.

## 1 Supplementary Information

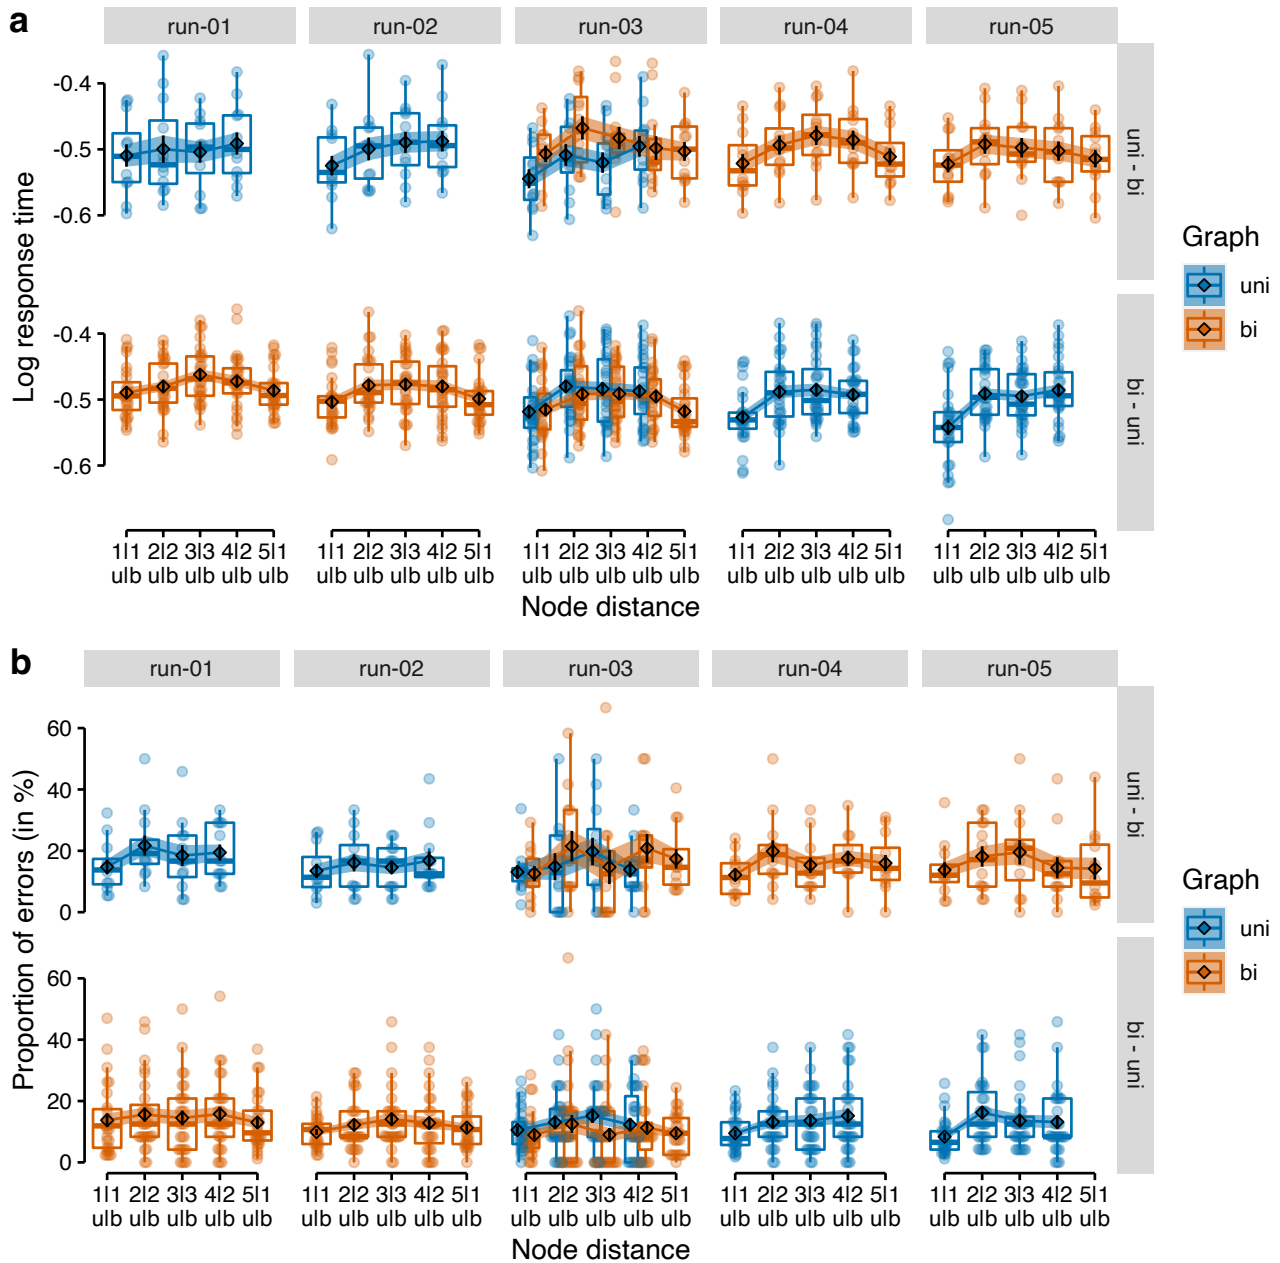
### 2 Supplementary Figures



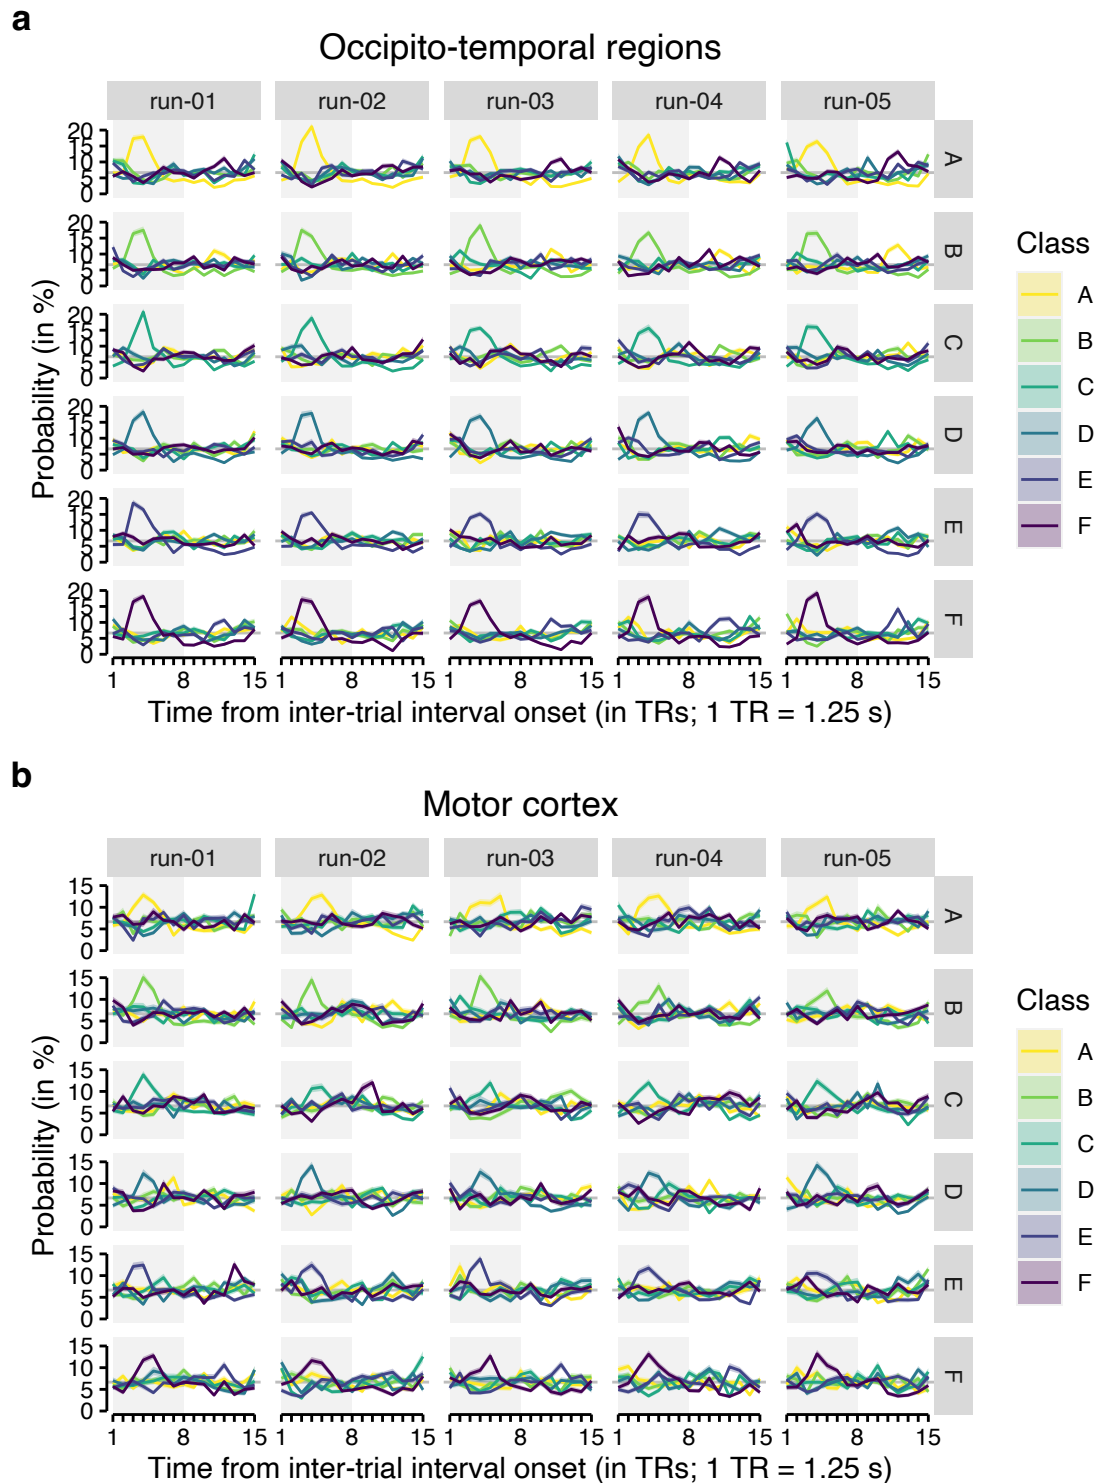
**Supplementary Figure S1: Study procedure.** (a) Session 1 started with a 5 min resting-state scan before participants read the task instructions and completed the training condition of the task. Participants then completed eight runs of the recall condition of ca. 6 min each before another 5 min resting-state scan was recorded. (b) Session 2 started with another run of the recall condition of ca. 6 min. Participants then completed all five runs of the graph learning task of about 10 min each which were interleaved with six resting-state scans of 3 min each. Both experimental sessions started with a short localizer scan and a T1w anatomical scan and ended with the acquisition of fieldmaps. During these scans and additional preparations by the study staff (e.g., orientation of the FOV) participants were asked to keep their eyes closed. Numbers inside the rectangles indicate approximate duration of each step in minutes (mins). Colors indicate participants' task (see legend).



**Supplementary Figure S2: Behavioral accuracy and response times per task run in training, recall, and graph trials.** Mean behavioral accuracy (in %; y-axis) per task run of the study (x-axis) in (a) training trials, (b) recall trials in session 1, (c) recall trials in session 2, and (d) graph trials in session 2. (e) Mean log response time (y-axis) per task run of the study (x-axis) in graph trials. The chance-level (gray dashed line) is at 16.67%. Each dot corresponds to averaged data from one participant. Colored lines connect data across runs for each participant. Boxplots indicate the median and IQR. The lower and upper hinges correspond to the first and third quartiles (the 25<sup>th</sup> and 75<sup>th</sup> percentiles). The upper whisker extends from the hinge to the largest value no further than 1.5\* IQR from the hinge (where IQR is the interquartile range (IQR), or distance between the first and third quartiles). The lower whisker extends from the hinge to the smallest value at most 1.5\* IQR of the hinge. The diamond shapes show the sample mean. Error bars and shaded areas indicate  $\pm 1$  SEM. All statistics have been derived from data of  $n = 39$  human participants who participated in one experiment.

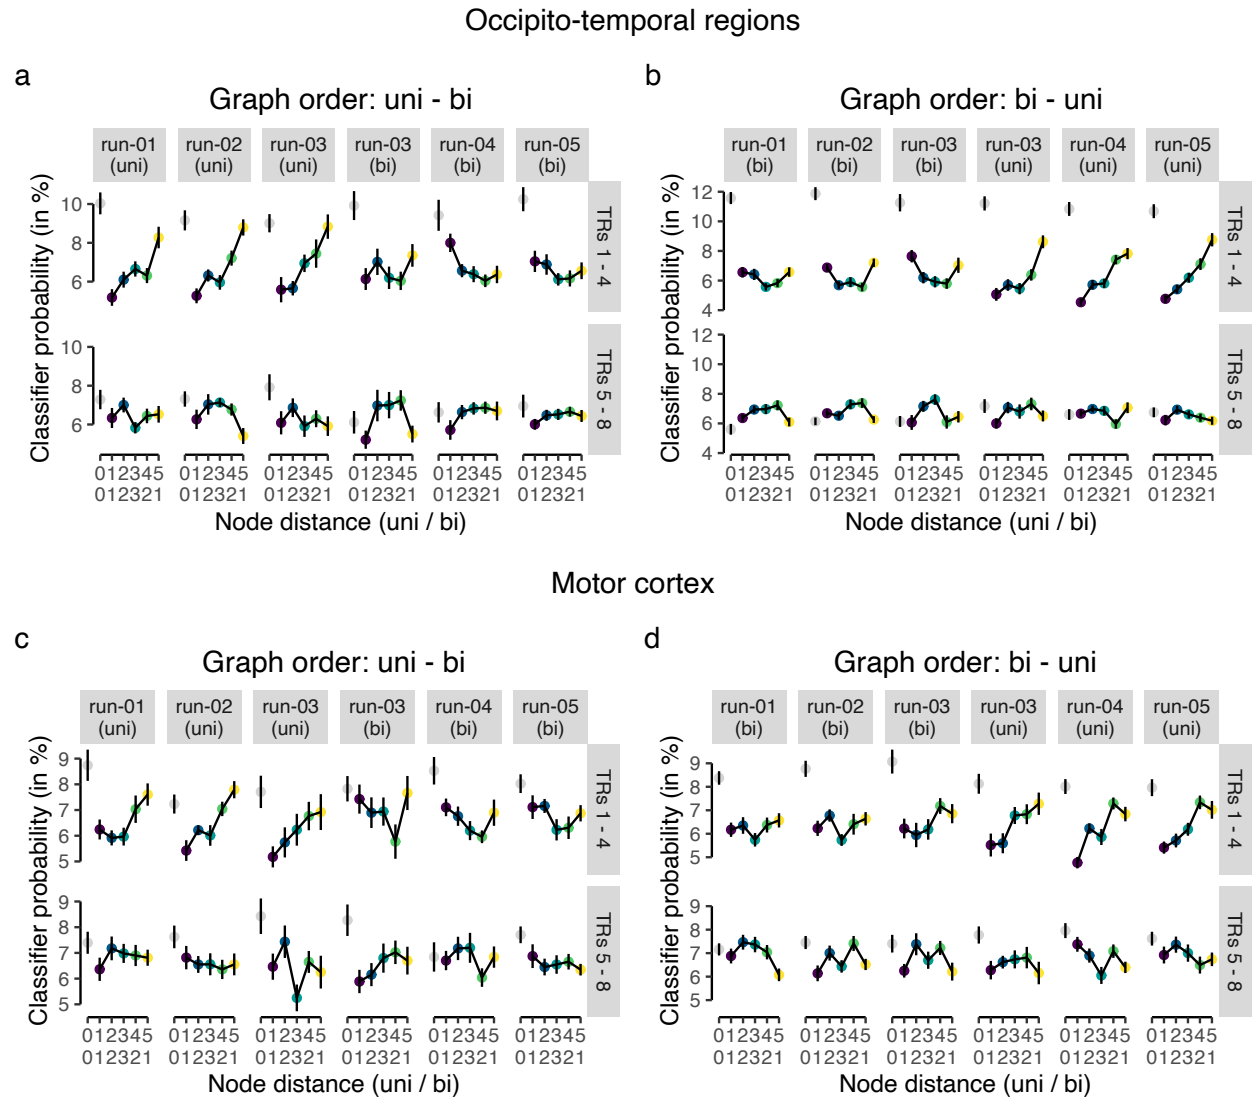


**Supplementary Figure S3: Behavioral responses across task runs.** (a) Log response times (y-axis) as a function of node distance (x-axis) in the graph structure (colors) for each task run (vertical panels) and graph order (uni – bi vs. bi – uni; horizontal panels). (b) Proportion of errors (in %; y-axis; relative to the total number of trials per node distance and run) as a function of node distance (x-axis) in the graph structure (colors) for each task run (vertical panels) and graph order (uni – bi vs. bi – uni; horizontal panels). Boxplots indicate the median and IQR. The lower and upper hinges correspond to the first and third quartiles (the 25<sup>th</sup> and 75<sup>th</sup> percentiles). The upper whisker extends from the hinge to the largest value no further than 1.5\* IQR from the hinge (where IQR is the interquartile range (IQR), or distance between the first and third quartiles). The lower whisker extends from the hinge to the smallest value at most 1.5\* IQR of the hinge. The diamond shapes show the sample mean. Each dot corresponds to averaged data from one participant. Error bars and shaded areas represent  $\pm 1$  SEM. All statistics have been derived from data of  $n = 39$  human participants who participated in one experiment.

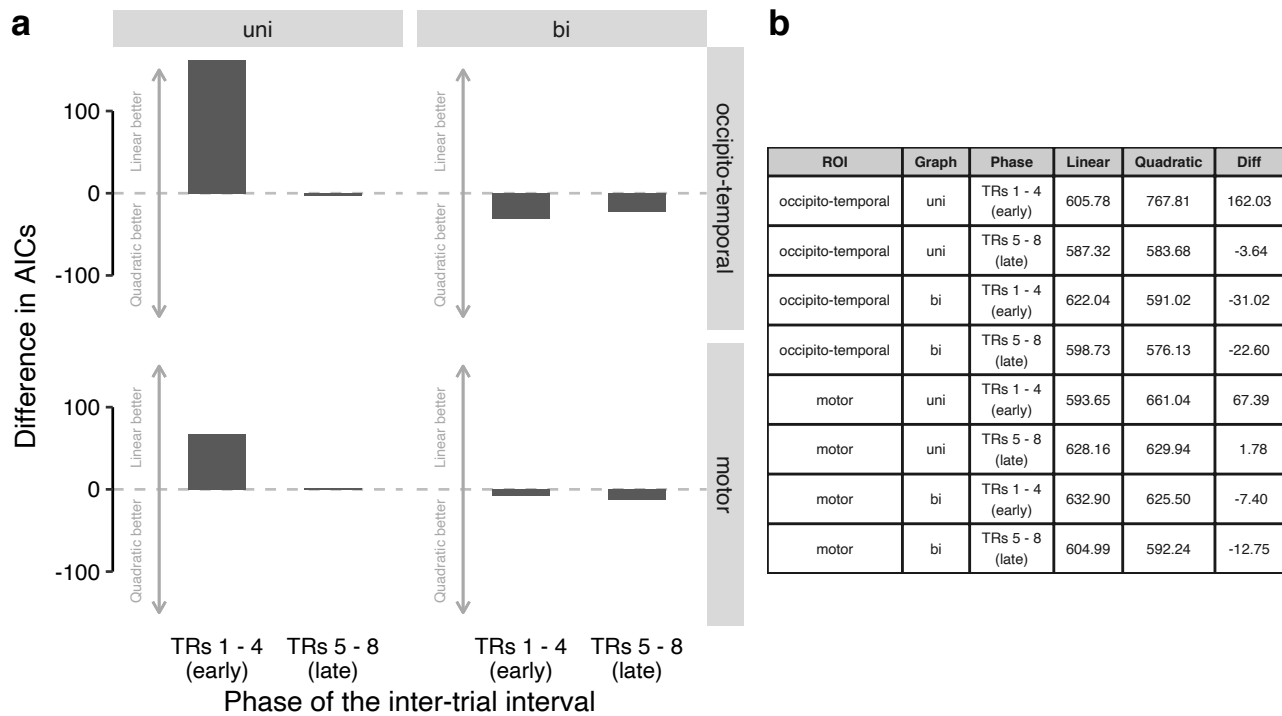


**Supplementary Figure S4: Classifier probabilities in long ITIs of graph trials.** Time courses (in TRs from the onset of the ITIs; x-axis) of classifier probabilities (in %; y-axis) per class (colors; see legend) and run (vertical panels). Substantial delayed and extended increases in classifier probability were found for the class that occurred on a given trial (horizontal panels) in both occipito-temporal brain regions (a) and motor and somatosensory cortex (b), peaking around the fourth TR following ITI onset, as expected given that classifier were trained on the fourth TR from event onset in fMRI data from recall trials. Each line represented averaged data across all trials of all participants. All shaded areas represent  $\pm 1$  SEM. Gray rectangles indicate the long ITI (TRs 1–8). All statistics have been derived from data of  $n = 39$  human participants who participated in one experiment.





**Supplementary Figure S5: Classifier probabilities during graph trials are modulated by node distance in the graph structure.** Classifier probabilities (in %; y-axis) as a function of the distance between the nodes in the uni-directional (first line) and bi-directional (second line) graph structure averaged across TRs in the early (TRs 1–4) or late (TRs 5–8) phase (horizontal panels) of the long ITIs of the five runs (vertical panels) in graph trials for data in the occipito-temporal (a), (b) and motor cortex (c), (d) ROIs. Each dot corresponds to data averaged across participants. Error bars represent  $\pm 1$  SEM. All statistics have been derived from data of  $n = 39$  human participants who participated in one experiment.



**Supplementary Figure S6: Model comparison of LME models with linear vs. quadratic predictor of classifier probabilities in ITIs of graph trials.** (a) Difference in AIC values for LME models including a linear vs. a quadratic predictor for mean classifier probabilities for the two TR phases (early vs. later), the two graph conditions (uni vs. bi; vertical panels) and the two ROIs (occipito-temporal vs. motor; horizontal panels). Positive values indicate a better fit of the LME model with the linear predictor and negative values indicate a better fit of the LME model with the quadratic predictor. (b) Table of AIC values of LME models with linear and quadratic predictor (and their difference) for all combinations of ROI, graph condition, TR phase. All statistics have been derived from data of  $n = 39$  human participants who participated in one experiment with two sessions.

### 3 Task instructions in English

Box S1: Screen 1 of instructions for the training condition in session 1

Welcome to the study - Session 1!

Please read the following information carefully. If you have any questions, you can clarify them right away with the study instructor. Please lie as still and relaxed as possible for the entire time.

Press any key to continue.

Box S2: Screen 1 of instructions for the training condition in session 1

Your task:

You are a zookeeper in training and have to make sure that all animals are in the right cages. First you will learn in a training which animal belongs in which cage. We will now explain to you exactly how this task works.

Press any key to continue.

Box S3: Screen 3 of instructions for the training condition in session 1

Training (Part 1)

You want to become a zookeeper and start your training today. First you will learn which animal belongs in which cage. You will see six cages at the bottom of the screen. Each of the six cages belongs to one of six animals. You will select a cage with the appropriate response key. Please keep your ring, middle and index fingers on the response keys the entire time so that you can answer as quickly and accurately as possible.

Press any key to continue.

Box S4: Screen 4 of instructions for the training condition in session 1

During the training, the animals appear above their cages. Press the key for that cage as fast as you can and remember the cage where the animal belongs. Please press the correct button within 1 second. Please answer as quickly and accurately as possible. You will receive feedback if your answer was correct, incorrect or too slow. The correct cage will appear in green and the incorrect cage will appear in red.

Press any key to continue.

Box S5: Screen 5 of instructions for the training condition in session 1

It is very important that you actively remember which animal belongs in which cage. You will get a higher bonus if you remember the correct assignment. The better you remember which animal belongs in which cage, the more money you earn! You will now complete one pass of this task, which will take approximately 2 minutes.

Press any key to continue.

Box S6: Screen 1 of instructions for the recall condition in session 1

Training (part 2)

We will now check how well you have learned the assignment of the animals to their cages. The animals will now appear in the center of the screen. You are asked to remember the correct cage for each animal, and then press the correct key as quickly as possible.

Press any key to continue.

Box S7: Screen 2 of instructions for the recall condition in session 1

This time you respond only after the animal is shown. In each round, the animal will appear first in the center of the screen. Then please try to actively imagine the correct combination of animal, cage and response key. After that, a small cross will appear for a short moment. Then the cages appear and you can respond as quickly and accurately as possible. Please respond as soon as the cages appear, not earlier.

Press any key to continue.

Box S8: Screen 3 of instructions for the recall condition in session 1

You have again 1 second to respond. Please respond again as fast and accurate as possible. You will get feedback again if your response was wrong or too slow. If your response was correct, you will continue directly with the next round without feedback. You will now complete 8 passes of this task, each taking about 6 minutes. In between the rounds you will be given the opportunity to take a break.

Press any key to continue.

Box S9: Screen 1 of instructions for the recall condition in session 2

Welcome to the study - Session 2!

We will check again if you can remember the assignment of the animals to their cages. The animals will appear in the center of the screen again. You are asked to remember again the correct cage for each animal and press the correct key as quickly as possible.

Press any key to continue.

Box S10: Screen 2 of instructions for the recall condition in session 2

You answer again only after the animal has been shown. In each round, the animal appears first in the center of the screen. Then please try to actively imagine the correct combination of animal, cage and answer key. After that, a small cross will first appear for a short moment.

Then the cages appear and you can answer as quickly and accurately as possible. Please respond as soon as the cages appear, not earlier.

Press any key to continue.

Box S11: Screen 3 of instructions for the recall condition in session 2

You have again 1 second to respond. Please respond again as fast and accurate as possible.

You will get feedback again if your response was wrong or too slow. If your answer was correct, you will proceed directly to the next round without feedback. You will now complete a run-through of this task, which will again take approximately 6 minutes. After the round you will be given the opportunity to take a break. Press any key to continue.

Box S12: Screen 1 of instructions for the graph condition in session 2

You have finished the passage to memory! Well done! You are now welcome to take a short break and also close your eyes. Please continue to lie still and relaxed. When you are ready, you can continue with the instructions for the main task.

Press any key to continue.

Box S13: Screen 2 of instructions for the graph condition in session 2

Main task

Congratulations, you are now a trained zookeeper! Attention: Sometimes the animals break out of their cages! Your task is to bring the animals back to the right cages. When you see an animal on the screen, press the right button as fast as possible to bring the animal back to the right cage. This time you will not get any feedback if your answer was right or wrong. The more animals you put in the correct cages, the more bonus you get at the end of the trial!

The main task consists of 5 runs, each taking about 10 minutes to complete.

Press any key to continue.

Box S14: Screen 3 of instructions for the graph condition in session 2

You have again 1 second to respond. In the main task, you again respond immediately when you see an animal on the screen. Again, please respond as quickly and accurately as possible.

Between each round you will again see a cross for a moment. Sometimes the cross will be shown a little shorter and sometimes a little longer. It is best to stand by all the time to respond as quickly as possible to the next animal.

Press any key to continue.

Box S15: Screen 4 of instructions for the graph condition in session 2

Resting phases

After all the work as a zookeeper you also need rest. Before, between and after the main task we will take some measurements during which you should just lie still. During these rest periods, please keep your eyes open and look at a cross the entire time. Blinking briefly is perfectly fine. The background of the screen will be dark during the resting phases. Please continue to lie very still and relaxed and continue to try to move as little as possible. Please

try to stay awake the entire time.

Please wait for the study instructor.

#### 4 Task instructions in German

##### Box S16: Screen 1 of instructions for the training condition in session 1

Willkommen zur Studie - Sitzung 1!

Bitte lesen Sie sich die folgenden Informationen aufmerksam durch. Falls Sie Fragen haben, können Sie diese gleich mit der Versuchsleitung klären. Bitte liegen Sie die gesamte Zeit so ruhig und entspannt wie möglich.

Drücken Sie eine beliebige Taste, um fortzufahren.

##### Box S17: Screen 2 of instructions for the training condition in session 1

Ihre Aufgabe:

Sie sind ein\*e Zoowärter\*in in Ausbildung und sollen darauf achten, dass alle Tiere in den richtigen Käfigen sind. Zuerst werden Sie in einem Training lernen, welches Tier in welchen Käfig gehört. Wir werden Ihnen jetzt genau erklären, wie diese Aufgabe funktioniert.

Drücken Sie eine beliebige Taste, um fortzufahren.

##### Box S18: Screen 3 of instructions for the training condition in session 1

Training (Teil 1)

Sie wollen Zoowärter\*in werden und beginnen heute Ihre Ausbildung. Zuerst lernen Sie, welches Tier in welchen Käfig gehört. Sie werden gleich sechs Käfige im unteren Teil des Bildschirms sehen. Jeder der sechs Käfige gehört zu einem von sechs Tieren. Sie wählen einen Käfig mit der entsprechenden Antworttaste aus. Bitte lassen Sie Ihre Ring-, Mittel- und Zeigefinger die gesamte Zeit auf den Antworttasten, damit Sie so schnell und genau wie möglich antworten können.

Drücken Sie eine beliebige Taste, um fortzufahren.

##### Box S19: Screen 4 of instructions for the training condition in session 1

Während des Trainings erscheinen die Tiere über ihren Käfigen. Drücken Sie die Taste für diesen Käfig so schnell wie möglich und merken Sie sich den Käfig, in den das Tier gehört. Bitte drücken Sie die richtige Taste innerhalb von 1 Sekunde. Bitte antworten Sie so schnell und genau wie möglich. Sie erhalten eine Rückmeldung, wenn Ihre Antwort richtig, falsch oder zu langsam war. Dabei erscheint der richtige Käfig in Grün und der falsche Käfig in Rot.

Drücken Sie eine beliebige Taste, um fortzufahren.

Box S20: Screen 5 of instructions for the training condition in session 1

Es ist sehr wichtig, dass Sie sich aktiv merken, welches Tier in welchen Käfig gehört. Sie erhalten einen höheren Bonus, wenn Sie sich an die richtige Zuordnung erinnern. Je besser Sie sich daran erinnern, in welchen Käfig welches Tier gehört, desto mehr Geld verdienen Sie! Sie werden nun einen Durchgang dieser Aufgabe absolvieren, der circa 2 Minuten dauert.  
Drücken Sie eine beliebige Taste, um fortzufahren.

Box S21: Screen 1 of instructions for the recall condition in session 1

Training (Teil 2)

Wir werden nun überprüfen, wie gut Sie die Zuordnung der Tiere zu ihren Käfigen gelernt haben. Die Tiere werden nun in der Mitte des Bildschirms erscheinen. Sie sollen sich an den richtigen Käfig für jedes Tier erinnern und dann die richtige Taste so schnell wie möglich drücken.  
Drücken Sie eine beliebige Taste, um fortzufahren.

Box S22: Screen 2 of instructions for the recall condition in session 1

Dieses Mal antworten Sie erst nachdem das Tier gezeigt wurde. In jeder Runde erscheint zuerst das Tier in der Mitte des Bildschirms. Versuchen Sie dann bitte, sich die richtige Kombination von Tier, Käfig und Antworttaste aktiv vorzustellen. Danach erscheint zunächst ein kleines Kreuz für einen kurzen Moment. Dann erscheinen die Käfige und Sie können so schnell und genau wie möglich antworten. Bitte antworten Sie erst sobald die Käfige erscheinen, nicht früher.  
Drücken Sie eine beliebige Taste, um fortzufahren.

Box S23: Screen 3 of instructions for the recall condition in session 1

Sie haben wieder 1 Sekunde Zeit zu antworten. Bitte antworten Sie wieder so schnell und genau wie möglich. Sie erhalten wieder eine Rückmeldung, wenn Ihre Antwort falsch oder zu langsam war. Wenn Ihre Antwort richtig war, geht es ohne Rückmeldung direkt mit der nächsten Runde weiter. Sie werden nun 8 Durchgänge dieser Aufgabe absolvieren, die jeweils circa 6 Minuten dauern. Zwischen den Durchgängen werden Sie die Möglichkeit bekommen, eine Pause zu machen.  
Drücken Sie eine beliebige Taste, um fortzufahren.

Box S24: Screen 1 of instructions for the recall condition in session 2

Willkommen zur Studie - Sitzung 2!

Wir werden noch einmal überprüfen, ob Sie sich an die Zuordnung der Tiere zu ihren Käfigen erinnern können. Die Tiere werden wieder in der Mitte des Bildschirms erscheinen. Sie sollen sich wieder an den richtigen Käfig für jedes Tier erinnern und die richtige Taste so schnell wie möglich drücken.  
Drücken Sie eine beliebige Taste, um fortzufahren.



Box S25: Screen 2 of instructions for the recall condition in session 2

Sie antworten wieder erst nachdem das Tier gezeigt wurde. In jeder Runde erscheint zuerst das Tier in der Mitte des Bildschirms. Versuchen Sie dann bitte, sich die richtige Kombination von Tier, Käfig und Antworttaste aktiv vorzustellen. Danach erscheint zunächst ein kleines Kreuz für einen kurzen Moment. Dann erscheinen die Käfige und Sie können so schnell und genau wie möglich antworten. Bitte antworten Sie erst sobald die Käfige erscheinen, nicht früher.

Drücken Sie eine beliebige Taste, um fortzufahren.

Box S26: Screen 3 of instructions for the recall condition in session 2

Sie haben wieder 1 Sekunde Zeit zu antworten. Bitte antworten Sie wieder so schnell und genau wie möglich. Sie erhalten wieder eine Rückmeldung, wenn Ihre Antwort falsch oder zu langsam war. Wenn Ihre Antwort richtig war, geht es ohne Rückmeldung direkt mit der nächsten Runde weiter. Sie werden nun einen Durchgang dieser Aufgabe absolvieren, der wieder circa 6 Minuten dauert. Nach dem Durchgang werden Sie die Möglichkeit bekommen, eine Pause zu machen.

Drücken Sie eine beliebige Taste, um fortzufahren.

Box S27: Screen 1 of instructions for the graph condition in session 2

Sie haben den Durchgang zu Erinnerung beendet! Gut gemacht! Sie können jetzt gerne eine kurze Pause machen und dabei auch Ihre Augen schließen. Bitte bleiben Sie weiterhin ruhig und entspannt liegen. Wenn Sie bereit sind, können Sie mit den Instruktionen für die Hauptaufgabe fortfahren.

Drücken Sie eine beliebige Taste, um fortzufahren.

Box S28: Screen 2 of instructions for the graph condition in session 2

Hauptaufgabe

Herzlichen Glückwunsch, Sie sind nun ausgebildete\*r Zoowärter\*in! Achtung: Manchmal brechen die Tiere aus ihren Käfigen aus! Ihre Aufgabe ist es, die Tiere wieder in die richtigen Käfige zu bringen. Wenn Sie ein Tier auf dem Bildschirm sehen, drücken Sie so schnell wie möglich die richtige Taste, um das Tier zurück in den richtigen Käfig zu bringen. Dieses Mal bekommen Sie keine Rückmeldung, ob Ihre Antwort richtig oder falsch war. Je mehr Tiere Sie in die richtigen Käfige bringen, desto mehr Bonus bekommen Sie am Ende der Studie! Die Hauptaufgabe besteht aus 5 Durchgängen, die jeweils circa 10 Minuten dauern.

Drücken Sie eine beliebige Taste, um fortzufahren.

Box S29: Screen 3 of instructions for the graph condition in session 2

Sie haben wieder 1 Sekunde Zeit zu antworten. In der Hauptaufgabe antworten Sie wieder sofort, wenn Sie ein Tier auf dem Bildschirm sehen. Bitte antworten Sie wieder so schnell und genau wie möglich. Zwischen den einzelnen Runden sehen Sie wieder ein Kreuz für einen Moment. Manchmal wird das Kreuz etwas kürzer und manchmal etwas länger gezeigt. Am Besten halten Sie sich die ganze Zeit bereit, um so schnell wie möglich auf das nächste Tier zu reagieren.

Drücken Sie eine beliebige Taste, um fortzufahren.

Box S30: Screen 4 of instructions for the graph condition in session 2

Ruhephasen

Nach der ganzen Arbeit als Zoowärter\*in braucht man auch Erholung. Vor, zwischen und nach den Durchgängen der Hauptaufgabe machen wir einige Messungen bei denen Sie einfach nur ruhig liegen sollen. In diesen Ruhephasen sollen Sie bitte Ihre Augen geöffnet halten und die gesamte Zeit auf ein Kreuz schauen. Kurzes Blinzeln ist vollkommen in Ordnung. Der Hintergrund des Bildschirms wird in den Ruhephasen dunkel sein. Bitte liegen Sie weiterhin ganz ruhig und entspannt und versuchen Sie weiterhin sich so wenig wie möglich zu bewegen.

Versuchen Sie bitte die gesamte Zeit wach zu bleiben.

Bitte warten Sie auf die Versuchsleitung.