



A Koopman–Takens Theorem: Linear Least Squares Prediction of Nonlinear Time Series

Péter Koltai^{1,2} , Philipp Kunde³

¹ Department of Mathematics, University of Bayreuth, Nürnberger Straße 38, 95448 Bayreuth, Germany.

E-mail: peter.koltai@uni-bayreuth.de

² Institute of Mathematics, Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany

³ Faculty of Mathematics and Computer Science, Jagiellonian University in Kraków, ul. Prof. S. Łojasiewicza 6, 30-348 Kraków, Poland. E-mail: pkunde.math@gmail.com

Received: 31 August 2023 / Accepted: 30 March 2024

© The Author(s) 2024

Abstract: The least squares linear filter, also called the Wiener filter, is a popular tool to predict the next element(s) of time series by linear combination of time-delayed observations. We consider observation sequences of deterministic dynamics, and ask: Which pairs of observation function and dynamics are predictable? If one allows for nonlinear mappings of time-delayed observations, then Takens' well-known theorem implies that a set of pairs, large in a specific topological sense, exists for which an exact prediction is possible. We show that a similar statement applies for the linear least squares filter in the infinite-delay limit, by considering the forecast problem for invertible measure-preserving maps and the Koopman operator on square-integrable functions.

1. Introduction

The prediction (or forecast or extrapolation) of time series is important in diverse applications [1–5] and attracts lively research activity [6–15]. It can be framed as follows. Given a discrete-time dynamics $T : \mathbb{X} \rightarrow \mathbb{X}$ on state space \mathbb{X} and observation function $f : \mathbb{X} \rightarrow \mathbb{C}$, are we able to assess from the time series $f(x), f(Tx), \dots, f(T^n x)$ its next element $f(T^{n+1}x)$?

We are going to consider for this problem the linear least (mean) squares filter [16, section 9.7]. Going back to Wiener [17] and Kolmogorov [18, 19], it is also called *Wiener filter*. It computes linear combination coefficients for d consecutive observations $f(T^i x), \dots, f(T^{i+d-1} x)$ such that the square error of the linear combination to the next observations in the sequence, $f(T^{i+d} x)$, averaged over the sequence is minimal. The positive integer d is called *delay depth*.

If we were to allow nonlinear transformations to map the time-delayed observations, Takens' celebrated theorem [13] implies that an exact prediction is possible for a “large” set of pairs of map T and observable f :

Takens’ theorem. *Let \mathbb{M} be a compact manifold of dimension m . The set of pairs of mappings and observables $(T, f) \in \text{Diff}^2(\mathbb{M}) \times C^2(\mathbb{M}, \mathbb{R})$, for which the delay-coordinate map*

$$\Phi_{T,f} : \mathbb{M} \rightarrow \mathbb{R}^{2m+1}, \quad \Phi_{T,f}(x) = \left(f(x), f(Tx), \dots, f(T^{2m}x) \right)$$

is an embedding, is open and dense with respect to the C^1 topology.

To be more precise, exact predictability is implied by the consequence of Takens’ theorem that the mapping $(f(T^i x), \dots, f(T^{i+d-1} x)) \mapsto (f(T^{i+1} x), \dots, f(T^{i+d} x))$ with $d = 2m + 1$ is one-to-one as the composition $\Phi_{T,f}(x) \mapsto x \mapsto Tx \mapsto \Phi_{T,f}(Tx)$ of (smooth) one-to-one mappings. Note that a part of this pipeline is the *reconstruction* step $\Phi_{T,f}(x) \mapsto x$, possible by the injectivity of the map $\Phi_{T,f}$. Closer to our way of looking at the problem, [8] considers the (less restrictive) problem of predicting from d consecutive terms of the time series its future values. Notice that such *prediction* takes place not in the original phase space \mathbb{M} , but in $\Phi_{T,f}(\mathbb{M}) \subset \mathbb{R}^d$, considered as a model space for the system. See also [20].

The descriptor “large” from above refers to the adjectives ‘open and dense’ from Takens’ theorem—often called genericity in this context. It is customary to call a property generic if it holds for all points of a set that is large in a topological sense. This largeness, however, has no unique description throughout the literature. Indeed, in most cases the word “generic” refers to a property that holds for a set of points containing a dense G_δ set. Here, a set is called G_δ if it is a countable intersection of open sets, cf. Definition 15 below. Note that this is slightly weaker than the genericity used in Takens’ theorem itself.

Takens’ theorem has been extended over the years in different directions, such as less smoothness of the setting [6, 7, 11, 21, 22] (see also [8, 23] for a comprehensive overview) or stronger concepts for the largeness of the set for which it holds. One such concept is *prevalence* [11], which is a generalization of “Lebesgue-almost every” for infinite-dimensional normed spaces; cf. Definition 16.

Takens-type theorems serve as a justification of the validity of time-delay based approaches used in applications, and have been the basis of further methodological developments, see e.g. [2–5, 9, 10, 24–28].

At this point the prediction problem seems to reduce to learning the scalar function $(f(x), \dots, f(T^{d-1} x)) \mapsto f(T^d x)$ for a sufficiently large delay depth. However, this function is nonlinear and the curse of dimensionality becomes a burden, as no additional exploitable structural property of it is known. Although, for instance, deep-learning based approaches have been deployed to this learning problem with success [15, 29–31], training such artificial neuronal networks remains a nontrivial optimization problem and performance guarantees are not available. Similar difficulties arise for “composite” methods, where for a fixed set of nonlinear ansatz functions a linear combination optimizing the forecast error is sought [28, 32, 33] and for the related family of regression methods based on the Mori–Zwanzig formalism [34, 35], utilizing a projection-based quantification of how the unobserved past propagates to the present and future [36–39]. Due to its simplicity and (also theoretic) accessibility, the linear filter remains a viable tool for prediction. In addition, as we will see below, the forecast problem is all about approximating a linear operator.

Indeed, we interpret the prediction problem by the linear least squares filter in the setting of ergodic theory, with an essentially invertible measure-preserving transformation $T : (\mathbb{X}, \mathfrak{B}, \mu) \rightarrow (\mathbb{X}, \mathfrak{B}, \mu)$ and its associated Koopman operator $\mathcal{U} : L^2_\mu \rightarrow L^2_\mu$

with $\mathcal{U}f = f \circ T$ [40,41], often also called composition operator. We will see in Proposition 3 below that the least squares filter in the limit of infinite data corresponds to an orthogonal projection of \mathcal{U} to the Krylov subspace $\mathcal{K}_d = \text{span}\{f \circ T^{-d+1}, \dots, f \circ T^{-1}, f\} \subset L^2_\mu$ spanned by time-delayed observables.

Without an explicit mentioning of it, a first connection of the linear least squares filter and the Koopman operator in this projective sense seems to have appeared in [42]. Therein, the focus was on the connection to a more recent linear method, Dynamic Mode Decomposition (DMD), which arose in the context of fluid dynamics [43]. The relationship of DMD and the Koopman operator was observed in [44], after which generalizations such as Extended DMD [45] appeared. This led to the reintroduction of the Wiener filter with the name “Hankel DMD” in [42].

The marriage of Koopman operator and delay embedding has since then witnessed a lively activity. Intermittency in chaotic systems was investigated by the “Hankel alternative view of Koopman analysis” [26,27], spectral approximation results of Hankel DMD have been derived in [46], while others have also sought for spectral approximation beyond the linear filter scenario [46–51].

Note that prediction requires the inference of $\mathcal{U}f(x)$ from past observations $\dots, f \circ T^{-2}(x), f \circ T^{-1}(x), f(x)$. The usual trade-off in Koopman-operator based methods is that the original nonlinear but (potentially) finite-dimensional system is considered through a linear, but infinite-dimensional operator. It is thus not surprising that in general perfect prediction by linear manipulations of time-delayed observations cannot be achieved, merely in the limit of infinite delay depth, i.e. “asymptotically”. We will consider two notions of *asymptotic linear predictiveness*. The first is (strong) predictiveness, requiring that the closure $\overline{\bigcup_{d \in \mathbb{N}} \mathcal{K}_d} = L^2_\mu$; i.e., that the Krylov spaces “fill” L^2_μ , cf. Definition 7. The second, weak predictiveness, we define by $\mathcal{U}f \in \overline{\bigcup_{d \in \mathbb{N}} \mathcal{K}_d}$, cf. Definition 10. With regard to the above discussion after Takens’ theorem, we note that the stronger version of predictiveness here is closer to the notion of reconstruction (of the Koopman operator), while the weaker one is in line with what one considered above as prediction in model space.

In the following we consider the Wiener filter and our results will mainly concern (strong) linear asymptotic predictiveness. We provide some Takens-type predictiveness theorems where the prediction map is linear. Using the notions of standard probability space without atoms (that is, a measure space isomorphic to $[0, 1]$ with Lebesgue measure) and weak topology on measure-preserving essential bijections from Sect. 3.1.2, our main results are summarized in a simplified fashion as follows:

Koopman–Takens theorem. *Let $(\mathbb{X}, \mathcal{B}, \mu)$ be a standard probability space without atoms and let MPT be the space of μ -preserving essential bijections on \mathbb{X} with the weak topology.*

- (a) *Theorem 23: For a dense set of transformations $T \in MPT$, the collection of asymptotically linearly predictive observables f is prevalent and generic in L^2_μ .*
- (b) *Proposition 26: For a generic set of transformations $T \in MPT$, the collection of asymptotically linearly predictive observables f is generic in L^2_μ .*
- (c) *Proposition 14: For a dense set of transformations $T \in MPT$, every L^2_μ -observable f is weakly linearly asymptotically predictive. Specifically, the conclusion holds for all transformations T with discrete spectrum.*

Practically this means that for a large set of systems almost any measurement procedure (i.e., observable) results in a predictive scenario. However, it also means, that on finite

delay depths the Wiener filters of predictive and of unpredictable systems can be arbitrary close (Proposition 38). A result of our analysis is also that there are mixing systems that are predictive (Sect. 3.5). Since mixing systems “lose their memory” over time, it was not clear a priori whether observations from arbitrary far in the past can contribute enough for this to be the case.

We can compare the Takens and Koopman–Takens theorems along different lines:

- (i) While for Takens’ theorem the (often unknown) dimension of the manifold \mathbb{M} provides a bound on the minimal delay depth for predictability to hold, in the Koopman–Takens setting statements hold in the asymptotic limit of infinite delay depth. In fact, in this latter ergodic setting the structure of a manifold is not needed; we will only require Lebesgue–Rokhlin spaces below (Sect. 3.1.2). However, in Sect. 3.6 we can extend our results to the setting of measure-preserving homeomorphisms on compact connected manifolds.
- (ii) Takens-type results require the mapping and the observable to be at least (sometimes Lipschitz) continuous, while our framework works with measurable maps and observables. It should also be observed that the prediction problem can be considered for both invertible and non-invertible transformations T , while the possibility of dynamical reconstruction (embedding) by delay coordinates maps is generally restricted to invertible systems in the Takens framework. It is important to note that from a theoretical point of view, reconstruction implies prediction but not vice versa. In the Koopman–Takens framework we only consider essentially invertible maps.
- (iii) For Takens’ theorem, one can identify a class of $T \in \text{Diff}^2(\mathbb{M})$ such that the conclusion of the theorem holds: If T has finitely many periodic points of period at most $2m$ and for any periodic point x of period $p \leq 2m$ the eigenvalues of $DT^p(x)$ are all distinct, then for generic $f \in C^2(\mathbb{M}, \mathbb{R})$ the map $\Phi_{T,f}$ is an embedding [52, Theorem 2]. In our setting, one can also give a class of systems for which weak predictiveness holds: Mappings with discrete spectrum are weakly predictive for every L^2_μ -observable (see Proposition 14).

The paper is structured as follows. Section 2 sets up the prediction problem in an ergodic-theoretic context and shows that an important role is played by so-called cyclic vectors of the Koopman operator. The size of the set of predictive pairs of dynamics and observables is considered in Sect. 3, where we give different genericity results for predictiveness—both of constructive and nonconstructive nature. Further practical properties of the least squares linear filter for time series analysis are collected in Sect. 4. We illustrate our findings on four numerical examples in Sect. 5, before we conclude with Sect. 6.

2. The Least-Squares Filter

2.1. Preliminaries on operator-based considerations in ergodic theory. The measure-theoretic consideration of dynamical systems starts with a measure space $(\mathbb{X}, \mathfrak{B}, \mu)$ and a measure-preserving transformation (mpt) $T : \mathbb{X} \rightarrow \mathbb{X}$, i.e., $\mu = \mu \circ T^{-1}$. If the σ -algebra \mathfrak{B} is clear from the context, we suppress it in the notation, and write (\mathbb{X}, μ) . We will work with probability spaces, thus $\mu(\mathbb{X}) = 1$.

The system (or, if no ambiguity arises, the *dynamics* or the measure) is called ergodic if there are only trivial invariant sets, that is, for every set $E \in \mathfrak{B}$ with $\mu(T^{-1}(E)\Delta E) = 0$ we have $\mu(E) = 0$ or $\mu(E) = 1$. Every system can be restricted to subsets $\mathbb{X}' \subset \mathbb{X}$, such that $T|_{\mathbb{X}'}$ is ergodic; so-called ergodic components. For ergodic systems, the Birkhoff

Ergodic Theorem states that for $f \in L^p_\mu := L^p_\mu(\mathbb{X})$, $1 \leq p < \infty$, the orbital averages converge, more precisely

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} f(T^i x) = \int_{\mathbb{X}} f d\mu$$

for μ -a.e. $x \in \mathbb{X}$ and also as a function in L^p_μ . Note that this is a functional version of the indecomposability statement defining ergodicity: The orbital averages of any observables converge to a constant. Again, if T is not ergodic, this statement holds restricted to its ergodic components.

This functional view on ergodic theory (also) motivates to consider observables and their evolution under the dynamics, giving rise to the Koopman operator $\mathcal{U} : L^2_\mu \rightarrow L^2_\mu$, $\mathcal{U}f = f \circ T$. The analogously defined operator could be considered on L^p_μ , $1 \leq p \leq \infty$, as well, but the additional structure of a Hilbert space and the possibility to consider autocorrelations leads to our choice of L^2_μ . Sometimes, to omit ambiguity, we write \mathcal{U}_T to denote the Koopman operator associated to T . Since T preserves the measure, we have that $\|\mathcal{U}f\|_{L^2_\mu} = \|f\|_{L^2_\mu}$, and if T is (essentially) invertible, then \mathcal{U} is *unitary*. Then, in particular, $\mathcal{U}_{T^{-1}} = \mathcal{U}_T^{-1}$ and the spectrum of \mathcal{U} is contained in $\{z \in \mathbb{C} \mid |z| = 1\}$, since both \mathcal{U} and \mathcal{U}^{-1} are non-expansive. This spectrum can be further classified as follows. The mpt T (or the operator \mathcal{U}) is said to have (purely) *discrete spectrum*, if \mathcal{U} has a countable set of eigenvalues with associated eigenfunctions yielding an orthonormal basis $\phi_i, i \in \mathbb{N}$, of L^2_μ . Note that $\lambda_1 = 1$ is always an eigenvalue of \mathcal{U} with eigenfunction $\phi_1 = \mathbb{1}$ (the constant function with value one). If 1 is the only eigenvalue of \mathcal{U} , then T (and \mathcal{U}) is said to have (purely) *continuous spectrum*. This is for instance the case for mixing systems, while rotations of abelian groups have purely discrete spectrum [53, §3.3]. One speaks of T having mixed spectrum if neither of the above pure cases hold. We note that the terminology is consistent with the functional-analytic partitioning of spectra into the point, continuous, and residual parts. Indeed, unitary operators have no residual spectrum: one way to see this known fact is to combine [54, Theorem 9.2-4] and [54, section 10.6].

If $f, g \in L^2_\mu$, then $f\mathcal{U}^k g \in L^1_\mu$ for $k \geq 0$, and the Birkhoff Ergodic Theorem yields

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} f(T^i x) g(T^{i+k} x) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} (f g \circ T^k)(T^i x) = \int_{\mathbb{X}} f\mathcal{U}^k g d\mu \quad (1)$$

for μ -a.e. $x \in \mathbb{X}$.

Finally, the form of the observation sequence suggests to consider sequences of the form $(f, \mathcal{U}f, \mathcal{U}^2 f, \dots)$. If the so-called *cyclic subspace* of f , defined by

$$C_f := C_{f, \mathcal{U}} := \overline{\text{span}\{f, \mathcal{U}f, \mathcal{U}^2 f, \dots\}},$$

is equal to L^2_μ , then we call f a *cyclic vector* of \mathcal{U} . As we will see, such cyclic observables play an important role for the prediction problem.

2.2. The filter as L^2 -orthogonal projection. In all what follows, we assume T to be a μ -ergodic essentially invertible mpt on \mathbb{X} . Let a trajectory

$$x, Tx, \dots, T^m x$$

of length $(m + 1)$ be given. For a given delay depth $d \in \mathbb{N}$, let us consider the linear least squares filtering problem for forecasting the next observation from the previous d ones:

$$\tilde{J}_m(\tilde{c}) := \frac{1}{m - d + 1} \sum_{t=0}^{m-d} \left| f(T^{t+d}x) - \sum_{i=0}^{d-1} \tilde{c}_i f(T^{t+i}x) \right|^2 \rightarrow \min_{\tilde{c} \in \mathbb{C}^d}!, \tag{2}$$

where $\rightarrow \min_{\tilde{c} \in \mathbb{C}^d}!$ expresses the goal of minimizing $J(\tilde{c})$ over $\tilde{c} = (\tilde{c}_0, \dots, \tilde{c}_{d-1}) \in \mathbb{C}^d$. For μ -a.e. $x \in \mathbb{X}$ the value function converges as $m \rightarrow \infty$ by (1) and we obtain the limiting filtering problem

$$\tilde{J}(\tilde{c}) := \int \left| f(T^d x) - \sum_{i=0}^{d-1} \tilde{c}_i f(T^i x) \right|^2 d\mu(x) \rightarrow \min_{\tilde{c} \in \mathbb{C}^d}! \tag{3}$$

Replacing \tilde{J}_m by \tilde{J} is reasonable in the following sense.

Lemma 1. *Let $d \in \mathbb{N}$ be fixed and let $\{f, f \circ T, \dots, f \circ T^{d-1}\}$ be a linearly independent set in L^2_μ . Then the minimizer of \tilde{J}_m converges for almost every $x \in \mathbb{X}$ to the unique minimizer of \tilde{J} for $m \rightarrow \infty$.*

Proof. The linear independence assumption guarantees that $\tilde{J} : \mathbb{C}^d \rightarrow \mathbb{R}$ is a quadratic loss function with symmetric positive definite $d \times d$ Hessian matrix H . Thus, \tilde{J} has a unique minimizer, that can be characterized as the unique solution of a linear system $H\tilde{c} = b$ with a suitable right-hand side b . Also, $\tilde{J}_m : \mathbb{C}^d \rightarrow \mathbb{R}$ is a quadratic loss function, whose finitely many parameters converge by the Birkhoff Ergodic Theorem for a.e. $x \in \mathbb{X}$ to those of \tilde{J} . Thus, \tilde{J}_m has almost surely a unique minimizer for sufficiently large m , which converges essentially because $(H, b) \mapsto H^{-1}b$ is continuous for invertible matrices H . □

From now on, we will thus consider the infinite-data case. More precisely, we will work with the following equivalent formulation of problem (3). Note that by invariance of μ with respect to T and hence with respect to T^{d-1} as well, we have with $c_i := \tilde{c}_{d-1-i}$ that

$$\begin{aligned} \tilde{J}(\tilde{c}) &= \int \left| f(Tx) - \sum_{i=0}^{d-1} \tilde{c}_i f(T^{i-d+1}x) \right|^2 d\mu(x) \\ &\stackrel{j=d-1-i}{=} \int \left| f(Tx) - \sum_{j=0}^{d-1} c_j f(T^{-j}x) \right|^2 d\mu(x) =: J(c). \end{aligned} \tag{4}$$

We note that at this point it is convenient to write the (finite) trajectory, by redefining its starting point, as

$$T^{-m}x, \dots, T^{-1}x, x,$$

whenever we interpret the ergodic limit in the finite-data case.

To interpret the minimizer of J in (4), we define the Krylov ‘‘basis’’ $\mathcal{B}_d := (f, f \circ T^{-1}, \dots, f \circ T^{-d+1})$ and the associated Krylov space $\mathcal{K}_d := \text{span } \mathcal{B}_d$. The Krylov basis is a basis of \mathcal{K}_d if and only if it forms a linearly independent set.

Remark 2. Note that if f is a cyclic vector of U^{-1} , then \mathcal{B}_d is a basis for any $d \geq 1$, assuming that L^2_μ is infinite-dimensional.

Since $J(c) \rightarrow \min!$ is the best approximation problem in L^2_μ , the minimizer $c = (c_0, \dots, c_{d-1})$ of J induces a function

$$\mathcal{U}_d f := \sum_{j=0}^{d-1} c_j f \circ T^{-j} \in \mathcal{K}_d$$

that is the unique L^2_μ -orthogonal projection of $\mathcal{U}f$ onto \mathcal{K}_d . There will be no confusion of notation between \mathcal{U}_T and \mathcal{U}_d , as T will always be a map and d an integer. If \mathcal{B}_d is a basis, the associated coefficient vector c is unique. The orthogonal projection property is equivalent to the variational equation (by setting the gradient of J to zero)

$$\langle \mathcal{U}_d f - \mathcal{U}f, f \circ T^{-i} \rangle_{L^2_\mu} = 0 \quad \forall i = 0, \dots, d - 1. \tag{5}$$

We can extend the operator \mathcal{U}_d to \mathcal{K}_d by defining it to have the matrix representation

$$U_d = \begin{bmatrix} c_0 & 1 & & & \\ c_1 & 0 & 1 & & \\ \vdots & & \ddots & \ddots & \\ c_{d-2} & & & 0 & 1 \\ c_{d-1} & & & & 0 \end{bmatrix} \tag{6}$$

with respect to \mathcal{B}_d , where the first column is the vector c , the first upper diagonal contains ones, and all other entries are zero. By setting up the variational equation for the L^2_μ -orthogonal projection of the Koopman operator \mathcal{U} on the space \mathcal{K}_d , it is immediate to see the following

Proposition 3. *The operator $\mathcal{U}_d : \mathcal{K}_d \rightarrow \mathcal{K}_d$ with matrix representation U_d with respect to \mathcal{B}_d is the L^2_μ -orthogonal projection of \mathcal{U} onto \mathcal{K}_d .*

Remark 4. (Hankel DMD) A statement similar to Proposition 3 has been made in [42, section 3] and was used there in the case $\mathcal{K}_{d+1} = \mathcal{K}_d$ for some $d \in \mathbb{N}$. Therein, the filter was coined *Hankel Dynamic Mode Decomposition (Hankel DMD)*, because it arises as the least squares solution of the (usually overdetermined) linear system

$$\begin{bmatrix} f(T^{-1}x) & \dots & f(T^{-d}x) \\ f(T^{-2}x) & \dots & f(T^{-d-1}x) \\ \vdots & & \vdots \\ f(T^{-m+d-1}x) & \dots & f(T^{-m}x) \end{bmatrix} \begin{bmatrix} c_0 \\ \vdots \\ c_{d-1} \end{bmatrix} = \begin{bmatrix} f(x) \\ f(T^{-1}x) \\ \vdots \\ f(T^{-m+d}x) \end{bmatrix} \tag{7}$$

involving a Hankel-type matrix (more precisely, a rectangular Toeplitz matrix) in the construction of (Extended) Dynamic Mode Decomposition [43,45].

Remark 5. (Computational aspects) There are several issues that need to be taken into account when computing the filter numerically from finite data.

The problem of obtaining the filter vector c via (4) inherits the numerical condition of the Krylov basis \mathcal{B}_d . If the basis elements are highly non-orthogonal, the computation of c is ill-conditioned. This need not effect the performance of the filter, though. This is

because the filter’s performance is measured by the objective function J in L^2_μ and not in the space of coefficients c .

A more severe source of error is the finite-sample approximation (2) of J . From an ergodic-theoretic perspective, the convergence in Lemma 1 as $m \rightarrow \infty$ can be arbitrarily slow. From the perspective of quadrature, integrals of the form (5) involve functions $f \circ T^{-d+1}$, which might become increasingly irregular as d grows and thus slowing down convergence. Such integrals enter implicitly while computing inner products in the solution of the least squares problem (7). A remedy (simple, but without performance guarantee) is to compute the filter for larger sample sizes m (or start with reducing the set, if obtaining new samples is not an option) and thus checking whether convergence can be safely assumed.

We now return to our original question, whether the filter is able to continue an observation sequence $\dots, f(T^{-1}x), f(x)$. Hence, we would like to infer $\mathcal{U}f(x)$. While there are also other mathematical frameworks for quantifying predictability and information content in ensemble predictions (see [55–57] and references therein), we focus on the *mean squared forecast error* and can give the following result.

Proposition 6. *Let f be a cyclic vector for \mathcal{U}^{-1} . Then, as $d \rightarrow \infty$, we have*

$$\|\mathcal{U}_d f - \mathcal{U}f\|_{L^2_\mu} \rightarrow 0.$$

Proof. Since f is a cyclic vector, the space $\bigcup_{d \geq 0} \mathcal{K}_d$ is dense in L^2_μ . Since \mathcal{U}_d is the orthogonal projection of \mathcal{U} to \mathcal{K}_d , and the spaces \mathcal{K}_d are nested, i.e., $\mathcal{K}_{d_1} \subset \mathcal{K}_{d_2}$ for $d_1 \leq d_2$, this proves that $\mathcal{U}_d f$ converges to $\mathcal{U}f$. \square

Proposition 6 motivates the following notion of predictiveness.

Definition 7 (Predictiveness). We call the pair (T, f) of a mpt and an L^2_μ -observable *asymptotically linearly predictive in the L^2 -sense*, if f is a cyclic vector of $\mathcal{U}^{-1} = \mathcal{U}_{T^{-1}}$. If T and μ are clear from the context, we simply say that f is *predictive*.

Remark 8. The term “predictive” is motivated by the observation that if $f \in L^2_\mu$ is a cyclic vector of \mathcal{U}^{-1} , then for any $g \in L^2_\mu$ we can approximate $\mathcal{U}g$ by a linear combination of $\mathcal{U}^{-j} f$, $j \in \mathbb{N}$, to any desired accuracy in L^2_μ norm by minimizing a modified version of $\tilde{J}(\tilde{c})$ in (3).

An interesting question is, whether L^2_μ convergence in Proposition 6 can be replaced by almost sure pointwise convergence.

Remark 9. (Almost sure asymptotic predictiveness) Convergence of a sequence in L^2_μ implies convergence almost everywhere for a subsequence. Thus, if f is a cyclic vector of \mathcal{U}^{-1} , there is a sequence $(d_k)_{k \in \mathbb{N}}$ of delay depths for which “almost sure asymptotic linear predictiveness” holds. However, we do not know in general which subsequence this is. In practice, an a posteriori test for convergence of a sequence $(\mathcal{U}_{d_k} f(x))_{k \in \mathbb{N}}$ could be our only option.

We also note that almost-everywhere pointwise convergence of orthogonal projections on L^2 spaces is considered in [58, Theorems 2 and 3], but it seems unclear whether the conditions therein (e.g., positivity) on the projection can be met in our situation (i.e., L^2_μ -orthogonal projection on \mathcal{K}_d). It seems rather unlikely for an arbitrary observation function f .

Proposition 6 shows that cyclic vectors of \mathcal{U} are of particular interest and importance to the forecast properties of the filter. In Sect. 3 we shall discuss the size of the set of cyclic vectors and the size of the set of transformations that have “many” of them. Before doing so, we consider a weaker but “sufficient” notion of predictiveness.

2.3. On weak predictiveness. Observe that in order to approximate the next element of the time series by a linear filter, it is only required.¹ that $\mathcal{U}f$ can be arbitrarily well approximated in L^2_μ by polynomials in \mathcal{U}^{-1} applied to f .

Definition 10 (Weak predictiveness). We call the pair (T, f) , or simply the observable f , *weakly predictive* if $f \circ T \in \bigcup_{d \in \mathbb{N}} \mathcal{K}_d$. Equivalently, weak predictiveness means $\lim_{d \rightarrow \infty} d_{L^2_\mu}(\mathcal{U}f, \mathcal{K}_d) = 0$, where $d_{L^2_\mu}$ denotes the distance in L^2_μ of a function to a subspace.

In the following, we are almost exclusively going to consider the stronger notion of predictiveness implied by the existence of cyclic vectors, as in Definition 7 and Proposition 6. It is nevertheless worthwhile discussing the weaker notion presented here, since observables that are eigenfunctions of \mathcal{U} (such as constant functions) will always be weakly but never strongly predictive (except for pathological cases where L^2_μ is one-dimensional). In particular, we give a characterization of weak predictiveness in Theorem 12 and show in Proposition 14 that if the system has discrete spectrum, then every observable f is weakly predictive.

To analyze weak predictiveness of a system, we will use the Spectral Theorem for unitary operators. Let $\mathbb{T} \subset \mathbb{C}$ denote the unit circle. In a simplified form (namely by restricting the operator in question to a cyclic subspace), the Spectral Theorem of unitary operators states that \mathcal{U} is isomorphic to the multiplication operator on $L^2_\nu(\mathbb{T})$ for a suitable measure ν . The general form shows isomorphy to a sum of such spaces. See, for instance, [59, Chapter 8] or [60, Chapter 18] for the general statement and also for the form discussed next in Sect. 2.3.1.

2.3.1. Spectral Theorem for unitary operators As a warm-up, consider a unitary (with respect to the Euclidean inner product $\langle \cdot, \cdot \rangle$) matrix $A \in \mathbb{C}^{k \times k}$ with eigenvalues $\lambda_i \in \mathbb{T}$, $i = 1, \dots, k$, with an orthonormal set of eigenvectors v_i . We can write the image of a vector $v = \sum_i c_i v_i$ under $p(A)$, where p is a polynomial, as $p(A)v = \sum_i p(\lambda_i) c_i v_i$. Then we have that

$$\langle p(A)v, v \rangle = \sum_i p(\lambda_i) |c_i|^2 = \int_{\mathbb{T}} p(\lambda) \left(\sum_i |c_i|^2 \delta_{\lambda_i} \right) (d\lambda),$$

where δ_x denotes the Dirac measure centered at x . The measure $\nu = \sum_i |c_i|^2 \delta_{\lambda_i}$ thus encodes a weighted version of the spectrum of A associated with the vector v .

The above finite-dimensional construction can be generalized to unitary operators on Hilbert spaces. We will not consider the most general case, it will be sufficient to restrict a unitary operator $\mathcal{A} : H \rightarrow H$ on the Hilbert space H to the cyclic subspace $C_{v, \mathcal{A}}$ of \mathcal{A} and $v \in H$:

$$C_{v, \mathcal{A}} = \overline{\text{span} \{v, \mathcal{A}v, \mathcal{A}^2v, \dots\}}.$$

¹ For a possible nonlinear analogue of this in the context of Takens’ theorem, see [8, Definition 1.8].

Note that $C_{v, \mathcal{A}}$ is invariant under \mathcal{A} . By Riesz duality, one finds a unique (v -dependent) measure ν (which we will call the “trace of the spectral measure for v ”, or simply “trace measure”) with

$$\langle h(\mathcal{A})v, v \rangle = \int_{\mathbb{T}} h d\nu \quad \forall h \in C(\mathbb{T}).$$

In fact, this equality can be extended to hold for every $h \in L^2_\nu(\mathbb{T})$, and one can define a unitary transformation

$$W : L^2_\nu(\mathbb{T}) \rightarrow H, \quad Wh := h(\mathcal{A})v.$$

Note that if $\mathbb{1}$ denotes the constant function with value 1 on \mathbb{T} , then $W\mathbb{1} = v$. Ultimately, by applying W to the function $\zeta \mapsto \zeta h(\zeta)$, one has that

$$W^{-1} \mathcal{A} W h(\zeta) = \zeta h(\zeta) \quad \forall h \in L^2_\nu(\mathbb{T}). \tag{8}$$

Thus, on its spectrum, \mathcal{A} can be characterized by the multiplication operator M , with $(Mh)(\zeta) := \zeta h(\zeta)$, just as for matrices. The Spectral Theorem comprises the statements that W is a unitary transformation and that (8) holds.

2.3.2. Application of the Spectral Theorem For weak predictiveness, we would need $\mathcal{U}^{-1} f \in C_{f, \mathcal{U}}$ (for simplicity of notation, we swapped the roles of T and T^{-1}). In other words, for any $\varepsilon > 0$ we would like to have a polynomial p_ε such that

$$\|\mathcal{U}^{-1} f - p_\varepsilon(\mathcal{U})f\|_{L^2_\mu} < \varepsilon.$$

Note that the Spectral Theorem can be used in its above simplified form on the invariant cyclic subspace $C_{f, \mathcal{U}}$ for $v = f$ and $\mathcal{A} = \mathcal{U}$. With this and $W\mathbb{1} = f$, one has that

$$\begin{aligned} \|\mathcal{U}^{-1} f - p_\varepsilon(\mathcal{U})f\|_{L^2_\mu} < \varepsilon &\Leftrightarrow \|WM^{-1}W^{-1}f - Wp_\varepsilon(M)W^{-1}f\|_{L^2_\mu} < \varepsilon \\ &\Leftrightarrow \|M^{-1}\mathbb{1} - p_\varepsilon(M)\mathbb{1}\|_{L^2_\nu} < \varepsilon. \end{aligned}$$

Spelling out the last expression, we would like to approximate the function $\zeta \mapsto \zeta^{-1}$ arbitrarily well by polynomials in the L^2_ν norm. Note that if $\nu = \text{Leb}$, then this is impossible, because on the complex circle \mathbb{T} the monomials $\zeta \mapsto \zeta^k$, $k \in \mathbb{Z}$, are orthogonal trigonometric polynomials. If, however, ν is a finite sum of weighted Dirac measures (for finitely many distinct eigenvalues λ_i), then there is a polynomial p with $p(\lambda_i) = \frac{1}{\lambda_i}$ for all i . This follows from the invertibility of a Vandermonde matrix. A much broader answer, containing these previous two examples, can be given by invoking an implication of Szegő’s theorem:

Theorem 11 (Kolmogorov’s Density Theorem [61, Theorem 2.11.5]). *Let ν be a probability measure on the complex circle \mathbb{T} of the form $d\nu(\theta) = w(\theta) d\theta + d\nu_s(\theta)$, where ν_s is singular and $d\theta$ measures arc length. Then the polynomials are dense in $L^2_\nu(\mathbb{T})$ if and only if*

$$\int \log(w(\theta)) d\theta = -\infty.$$

Theorem 12. *Let $dv(\theta) = w(\theta) d\theta + dv_s(\theta)$ denote the decomposition of the trace measure ν associated with T and f as in Theorem 11. Then f is weakly predictive if and only if the Szegő condition $\int \log(w(\theta)) d\theta = -\infty$ holds.*

Proof. We recall that—with abbreviating notation such that ζ^n denotes the rational function $\zeta \mapsto \zeta^n$ for $n \in \mathbb{Z}$ —weak predictiveness is equivalent to that ζ^{-1} can be arbitrary well approximated by monomials $1, \zeta, \zeta^2, \dots$ in L^2_ν . Equivalently, we write $\zeta^{-1} \in \overline{\text{span}\{1, \zeta, \zeta^2, \dots\}}$.

“ \Leftarrow ”. Assume that the Szegő condition holds. Then, by Theorem 11 we have that $L^2_\nu = \overline{\text{span}\{1, \zeta, \zeta^2, \dots\}}$, and in particular $\zeta^{-1} \in \overline{\text{span}\{1, \zeta, \zeta^2, \dots\}}$.

“ \Rightarrow ”. The proof of this direction is due to Friedrich Philipp and is inspired by [62, Lemma 7.4]. We include it with his permission.

Assume weak predictiveness, i.e., $\zeta^{-1} \in \overline{\text{span}\{1, \zeta, \zeta^2, \dots\}}$. By [61, Lemma 2.11.3] we have that $\zeta^{-n} \in \overline{\text{span}\{1, \zeta, \zeta^2, \dots\}}$ for every $n \in \mathbb{N}$. In particular, for any two polynomials p, q we have that $p(\zeta)+q(\bar{\zeta}) \in \overline{\text{span}\{1, \zeta, \zeta^2, \dots\}}$, showing that $\overline{\text{span}\{1, \zeta, \zeta^2, \dots\}}$ is closed under complex conjugation (here, $\bar{\zeta}$ stands for the complex conjugate of ζ). The Stone–Weierstraß theorem now implies that $\overline{\text{span}\{1, \zeta, \zeta^2, \dots\}}$ is dense in $C(\mathbb{T})$, the usual space of continuous function on the unit circle. Since ν is a Borel probability measure on \mathbb{T} , it is regular, and [63, Theorem 3.14] shows that $C(\mathbb{T})$ is dense in L^2_ν . Theorem 11 now implies that the Szegő condition holds. \square

The above proof also shows that weak predictiveness implies that $U^n f$ can be approximated from past observations for every $n \in \mathbb{N}$, not merely for $n = 1$.

Remark 13. (Time flip)

Note that swapping the roles of T and T^{-1} is merely a reflection of the trace measure ν on the real axis. Thus, all statements on weak predictiveness that we consider here are valid both for T and T^{-1} .

Furthermore, the proof of Theorem 12 shows that weak predictiveness is equivalent to the cyclic subspace of f under U being equal to its cyclic subspace under U^{-1} , i.e., $C_{f,U} = C_{f,U^{-1}}$. To see this, first note that the proof shows the equivalence of weak predictiveness and $\overline{\text{span}\{\dots, \zeta^{-2}, \zeta^{-1}, 1, \zeta, \zeta^2, \dots\}} = \overline{\text{span}\{1, \zeta, \zeta^2, \dots\}} = L^2_\nu$. By complex conjugation one obtains that also $\overline{\text{span}\{1, \zeta^{-1}, \zeta^{-2}, \dots\}} = L^2_\nu$, which implies our claim.

If U has discrete spectrum, then for any observable f the decomposition of the associated trace measure ν in Theorem 11 has no absolutely continuous part, i.e., $w \equiv 0$. Theorem 12 thus implies:

Proposition 14. *If the invertible mpt T has discrete spectrum, then every observable $f \in L^2_\mu$ is weakly predictive.*

We note that Theorem 12 allows for mixing systems (that have purely continuous spectrum) also to be predictive; for instance by having a subinterval of \mathbb{T} (i.e., an arc) of positive (arc-length) measure on which $w = 0$. Note that this is compatible with our observation in Sect. 3.5 below, that there are mixing systems which even possess cyclic vectors.

3. Genericity of Asymptotically Linearly Predictive Pairs (T, f)

Predictiveness in Proposition 6 requires cyclic vectors of \mathcal{U}^{-1} . For simplicity of notation, in the remainder of this section we are going to speak about cyclic vectors of \mathcal{U} . We ask the reader to keep in mind that the statements hold true for \mathcal{U}^{-1} just as well, because the properties discussed herein are shared by T and T^{-1} .

In Sect. 3.1.2 below we define MPT as the set of all invertible mpts and endow it with the weak topology. We are going to investigate the size of the set of pairs (T, f) such that f is a cyclic vector of $\mathcal{U} = \mathcal{U}_T$. For instance, we show in Sect. 3.3 that a dense set of measure-preserving transformations admits a prevalent set of cyclic vectors. The collection of mpts with a dense G_δ set of cyclic vectors is even a dense G_δ set in MPT (see Sect. 3.4). This result is extended to the setting of measure-preserving homeomorphisms in Sect. 3.6.

In Sect. 3.5 we discuss so-called rank one systems. We will see that the rank one property is generic in MPT and guarantees existence of cyclic vectors.

3.1. Preliminaries.

3.1.1. Some terminology from topology

Definition 15 (G_δ and generic). In a topological space a G_δ set is a countable intersection of open sets. We call a property *generic* if the set with this property contains a dense G_δ set.

Generic properties represent sets which are large in a topological sense, and in particular, nonempty. Since a unitary transformation is an isometric bijection, it maps G_δ sets to G_δ sets. It is also easy to see that a unitary transformation maps dense sets to dense ones; we will spell this out in Lemma 21 for a specific setting.

For comparison, we also mention another related concept of genericity: Some authors define a generic property as one that holds on a *residual set* (that is, a countable intersection of dense open sets), with the dual concept being a *meagre set* (i.e., a countable union of nowhere dense closed sets). We observe as an important practical aspect for applications that, if a property holds on a residual set, it may not hold for every point, but perturbing the point slightly will generally land one inside the residual set (by nowhere density of the components of the meagre set). Thus, generic sets constitute the most important case to address in theorems and algorithms.

In this paper we will actually study generic sets in completely metrizable topological spaces. In those spaces a residual set is dense by Baire category theorem and, hence, the two concepts of genericity coincide.

The next “largeness” concept, prevalence, arose from the desire to generalize the notion “Lebesgue-almost surely” to infinite-dimensional spaces.

Definition 16 (Prevalence [11]). Let \mathbb{V} be a vector space. A set $\mathbb{S} \subset \mathbb{V}$ is called *prevalent* if there is a finite-dimensional *probe space* $\mathbb{E} \subset \mathbb{V}$ such that for all $v \in \mathbb{V}$ (Lebesgue-)almost every point in $v + \mathbb{E}$ belongs to \mathbb{S} .

We also note that unitary transformations map prevalent sets to prevalent sets. To see this, we pick an orthonormal basis $\{e_1, \dots, e_{\dim \mathbb{E}}\}$ of the finite-dimensional probe space \mathbb{E} and define the Lebesgue measure $\mu_{v+\mathbb{E}}$ on $v + \mathbb{E}$ via the coefficients with respect to this basis, that is, for any subset A of $v + \mathbb{E}$ we define $\mu_{v+\mathbb{E}}(A) = \lambda(A)$, where

$A = \left\{ c \in \mathbb{C}^{\dim \mathbb{E}} \mid v + \sum_{i=1}^{\dim \mathbb{E}} c_i e_i \in A \right\}$ and λ is the Lebesgue measure on $\mathbb{C}^{\dim \mathbb{E}}$. Since a unitary transformation U is isometric, it is measure-preserving as a map from $(v + \mathbb{E}, \mu_{v+\mathbb{E}})$ to $(Uv + U\mathbb{E}, \mu_{Uv+U\mathbb{E}})$.

3.1.2. Some terminology from measure theory Two measure spaces $(\mathbb{X}, \mathcal{B}, \mu)$ and $(\mathbb{X}', \mathcal{B}', \mu')$ are said to be *isomorphic* if there are null sets $N \subset \mathbb{X}$, $N' \subset \mathbb{X}'$ and an invertible map $\phi : \mathbb{X} \setminus N \rightarrow \mathbb{X}' \setminus N'$ such that ϕ and ϕ^{-1} both are measurable and measure-preserving maps. Such a map ϕ is called a *measure-theoretic isomorphism*. In the case when $(\mathbb{X}, \mathcal{B}, \mu) = (\mathbb{X}', \mathcal{B}', \mu')$ one sometimes also calls such an isomorphism ϕ an *automorphism*. We denote the collection of automorphisms of a given measure space by *MPT*. As usual, two measure-preserving transformations are identified if they differ on a set of measure zero only.

If $\mu(\mathbb{X}) = 1$, then the measure space $(\mathbb{X}, \mathcal{B}, \mu)$ is a probability space. An important class of probability spaces are the *Lebesgue–Rokhlin spaces*, that is, they are complete and separable probability spaces. We refer to [64, section 9.4] for details. Every Lebesgue–Rokhlin spaces is isomorphic to a disjoint union of the unit interval with Lebesgue measure and at most countably many atoms [64, Theorem 9.4.7].

In the following, we restrict ourselves to Lebesgue–Rokhlin spaces without any atoms and call them *standard measure spaces*. Since every standard measure space is isomorphic to the unit interval with Lebesgue measure on the Borel sets, every invertible mpt of a standard measure space is isomorphic² to an invertible Lebesgue-measure-preserving transformation on $[0, 1]$. Hence, it often suffices to prove statements for mpts on $[0, 1]$.

We endow *MPT* with the *weak topology*.

Definition 17 (Weak topology on *MPT* [65, p. 62]). A subbasis of the weak topology is given by all sets of the form

$$N(T, E, \varepsilon) := \{S \in \text{MPT} \mid \mu(T E \Delta S E) < \varepsilon\},$$

where E is measurable and $\varepsilon > 0$. Convergence in the weak topology of the sequence of mpts $(T_n)_n$ to the mpt T , denoted by $T_n \xrightarrow{w} T$, is thus defined such that $\mu(T_n E \Delta T E) \rightarrow 0$ as $n \rightarrow \infty$ for every $E \in \mathfrak{B}$.

This weak topology is even completely metrizable, see [65, p. 64]. To quote Halmos: “Since, however, this rather artificial construction [of the metric] does not seem to throw any light on the structure of *MPT*, I see no point in studying it further.” Weak convergence of mpts can be characterized as follows. The proof can be found in “Appendix A”.

Lemma 18. For $T, T_n \in \text{MPT}$, $n \in \mathbb{N}$, the following are equivalent:

1. $T_n \xrightarrow{w} T$ as $n \rightarrow \infty$.
2. $T_n^{-1} \xrightarrow{w} T^{-1}$ as $n \rightarrow \infty$.
3. $\mathcal{U}_{T_n} \rightarrow \mathcal{U}_T$ as $n \rightarrow \infty$ in the strong operator topology, i.e., $\|\mathcal{U}_{T_n} g - \mathcal{U}_T g\|_{L^2_\mu} \rightarrow 0$ for all $g \in L^2_\mu$.
4. $\mathcal{U}_{T_n} \rightarrow \mathcal{U}_T$ as $n \rightarrow \infty$ in the weak operator topology, i.e., $\langle \mathcal{U}_{T_n} g - \mathcal{U}_T g, h \rangle_{L^2_\mu} \rightarrow 0$ for all $g, h \in L^2_\mu$.
5. For every $g \in L^2_\mu$, one has that $\|g \circ T_n^i - g \circ T^i\|_{L^2_\mu} \rightarrow 0$ as $n \rightarrow \infty$ for every $i \in \mathbb{Z}$.

² Two mpts $T : \mathbb{X} \rightarrow \mathbb{X}$ and $S : \mathbb{X}' \rightarrow \mathbb{X}'$ are called *isomorphic* if there is an isomorphism $\phi : \mathbb{X} \rightarrow \mathbb{X}'$ such that $S = \phi \circ T \circ \phi^{-1}$.

3.2. *Ergodic circle rotations have a prevalent set of cyclic vectors.* Recall that predictiveness requires cyclic vectors of \mathcal{U}^{-1} , while, for simplicity of notation, we consider cyclic vectors of \mathcal{U} in the entire Sect. 3. Note that the statement of the next result holds just as well for T^{-1} as for T . Let $S^1 = \mathbb{R}/\mathbb{Z}$ denote the unit circle (the onedimensional unit torus) closed under addition.

Proposition 19. *Let $T : S^1 \rightarrow S^1$, $Tx = x + \alpha \pmod 1$, be an ergodic circle rotation and $\mathcal{U} : L^2(S^1) \rightarrow L^2(S^1)$ the associated Koopman operator, where the measure underlying $L^2(S^1)$ is the Lebesgue measure. Then \mathcal{U} has a prevalent set of cyclic vectors. Furthermore, the set of cyclic vectors contains a dense G_δ set.*

Proof. **1. Prevalence.** We recall that the translation semigroup $T^t f = f(\cdot - t)$, $t \in \mathbb{R}$, is strongly continuous on $L^2(S^1)$, which means that $\lim_{t \rightarrow s} T^t f = T^s f$ for $f \in L^2(S^1)$, cf. [66, I.4.18]. This implies, since the sequence $(n\alpha)_n$ is dense in S^1 by ergodicity of T , that

$$\liminf_{n \rightarrow \infty} \|\mathcal{U}^n f - f(\cdot - t)\|_{L^2} = 0 \quad \forall t \in S^1. \tag{9}$$

Assume that $f \in L^2(S^1)$ is not cyclic. Then there is a $0 \neq g \in L^2(S^1)$ such that

$$g \perp \overline{\text{span}\{f, \mathcal{U}f, \mathcal{U}^2 f, \dots\}},$$

which is by (9) equivalent to

$$0 = \int_{S^1} g f(\cdot - t) = (g * \mathcal{R}f)(t) \quad \forall t \in S^1, \tag{10}$$

where $\mathcal{R}f$, defined by $(\mathcal{R}f)(x) := f(-x)$ for $x \in S^1$, is the reflection of f and $g * h$ denotes the convolution of g with h . Denoting by $\mathcal{F} : L^2(S^1) \rightarrow \ell^2(\mathbb{Z})$ the Fourier transform, we have that (10) is equivalent to

$$\mathcal{F}g(k) \cdot \mathcal{F}\mathcal{R}f(k) = 0 \quad \forall k \in \mathbb{Z}. \tag{11}$$

Consider the probe vector $p \in L^2(S^1)$ such that the Fourier transform of its reflection is given by

$$\mathcal{F}\mathcal{R}p(k) := \begin{cases} 1, & k = 0, \\ \frac{1}{k^2}, & k \neq 0. \end{cases}$$

Alternatively, one could consider the real-valued probe vector p by replacing $\frac{1}{k^2}$ by $\frac{1}{k^2}(1 + i \text{sign}(k))$, where i denotes the imaginary unit. Let $\lambda \in \mathbb{C}$. By (11), the requirement for $f + \lambda p$ not to be cyclic is that

$$\mathcal{F}g(k) \cdot (\mathcal{F}\mathcal{R}f(k) + \lambda \mathcal{F}\mathcal{R}p(k)) = 0 \quad \forall k \in \mathbb{Z}$$

for some (possibly λ -dependent) $g \neq 0$. However, $g \neq 0$ implies that $\mathcal{F}g(k) \neq 0$ for at least one $k \in \mathbb{Z}$. Thus, there is a $k \in \mathbb{Z}$ with $\mathcal{F}\mathcal{R}f(k) + \lambda \mathcal{F}\mathcal{R}p(k) = 0$, implying $\lambda = -k^2 \mathcal{F}\mathcal{R}f(k)$. Hence, cyclicity of $f + \lambda p$ can fail at most at countably many λ . In other words, $f + \lambda p$ is cyclic for Lebesgue-almost every λ , rendering the set of cyclic vectors of \mathcal{U} prevalent.

2. Dense G_δ . By the above we have the implication

$$f \text{ not cyclic} \implies \exists k \in \mathbb{Z} \text{ such that } \mathcal{F}\mathcal{R}f(k) = 0. \tag{12}$$

For $k \in \mathbb{Z}$ let us define

$$A_k := \left\{ f \in L^2(S^1) \mid \mathcal{FR}f(k) \neq 0 \right\} = (\mathcal{FR})^{-1} \left(\left\{ s \in \ell^2(\mathbb{Z}) \mid s_k \neq 0 \right\} \right),$$

where we denote by s_k the k -th element of a sequence s . With (12) we have that $\bigcap_{k \in \mathbb{Z}} A_k$ is contained in the set of cyclic vectors. Note that both \mathcal{F} and \mathcal{R} are unitary transformations, hence $(\mathcal{FR})^{-1}$ is a bounded operator preserving openness and denseness with respect to the norm topologies. As the sets $\{s \in \ell^2(\mathbb{Z}) \mid s_k \neq 0\}$ are open and dense for every $k \in \mathbb{Z}$, the set $\bigcap_{k \in \mathbb{Z}} A_k$ is G_δ . Since the set $\{s \in \ell^2(\mathbb{Z}) \mid s_k \neq 0 \forall k \in \mathbb{Z}\}$ is also dense, so is its preimage $\bigcap_{k \in \mathbb{Z}} A_k$ under \mathcal{FR} . \square

Remark 20. Note that the set $\{s \in \ell^2(\mathbb{Z}) \mid s_k \neq 0 \forall k \in \mathbb{Z}\}$ is not open in $\ell^2(\mathbb{Z})$. Thus, “dense G_δ ” cannot be replaced in the statement of Proposition 19 by “open and dense”.

3.3. One map with a prevalent set of cyclic vectors implies densely many. We just showed that ergodic circle rotations possess a prevalent set of cyclic vectors. Through the isomorphy of the nonatomic and separable probability space (\mathbb{X}, μ) and $([0, 1], \text{Leb})$ (see Sect. 3.1.2 and also [65, p. 61]) we can build an ergodic “rotation” on (\mathbb{X}, μ) , hence a transformation with a prevalent set of cyclic vectors. First, one transfers a cyclic vector via isomorphy:

Lemma 21. *Let $T : \mathbb{X} \rightarrow \mathbb{X}$ be a μ -preserving transformation with a cyclic vector c . For every μ -preserving essential bijection $S : \mathbb{X} \rightarrow \mathbb{X}$, the mpt $\tilde{T} = S^{-1}TS$ has the cyclic vector $c \circ S$.*

Proof. Note that $\mathcal{U}_T^n f = \mathcal{U}_S \circ \mathcal{U}_T^n \circ \mathcal{U}_S^{-1}$ for $n \in \mathbb{N}$. In particular,

$$\mathcal{U}_{\tilde{T}}^n (c \circ S) = \mathcal{U}_S (\mathcal{U}_T^n c).$$

Since $\text{span}\{\mathcal{U}_T^n c \mid n \in \mathbb{N}\}$ is dense, it is sufficient to show that for any dense subspace $\mathbb{U} \subset L^2_\mu$ also the space $\mathcal{U}_S \mathbb{U}$ is dense. For any fixed $\varepsilon > 0$ and $f \in L^2_\mu$ let $u \in \mathbb{U}$ be such that $\|u - f \circ S^{-1}\|_{L^2_\mu} < \varepsilon$. Since S is measure-preserving and an essential bijection, we have that $\|u \circ S - f\|_{L^2_\mu} < \varepsilon$. Thus, $\mathcal{U}_S \mathbb{U}$ is dense and the claim follows. \square

To show denseness of transformations with specific ergodic properties, the following tool is extremely useful.

Lemma 22 (Conjugacy Lemma, [65, p. 77]). *Let T_0 be an aperiodic μ -preserving transformation (i.e., the set of periodic points has zero measure). In the weak topology the set of all transformations of the form $S^{-1}T_0S$, where S is a μ -preserving essential bijection, is dense in MPT.*

It follows that the set of transformations that has prevalently many cyclic vectors is dense in MPT, because we can match every cyclic vector c of T with a cyclic vector $c \circ S$ of $S^{-1}TS$. The mapping $c \mapsto c \circ S$ is a unitary transformation, since S is a measure-preserving essential bijection, and thus it preserves the properties discussed in Sect. 3.1.1.

Altogether we conclude the following denseness result.

Theorem 23. *The collection of μ -preserving transformations with a prevalent and dense G_δ set of cyclic vectors in L^2_μ lies dense in MPT with respect to the weak topology.*

Proof. We take any fixed irrational rotation R_α on S^1 (that has a dense G_δ and prevalent set of cyclic vectors by Proposition 19) and conjugate it with an isomorphism ϕ from \mathbb{X} to S^1 to get an aperiodic mpt $T_0 = \phi^{-1} \circ R_\alpha \circ \phi$ on \mathbb{X} . By Lemma 21, T_0 has a dense G_δ and prevalent set of cyclic vectors. Then we apply the Conjugacy Lemma (Lemma 22) to get a dense set of mpts, each of which have a dense G_δ and prevalent set of cyclic vectors, again by Lemma 21. \square

Remark 24. (Weakly predictive mpts are dense) In Theorem 23 we used an ergodic circle rotation to generate densely many ergodic transformations on an arbitrary Lebesgue–Rokhlin space $(\mathbb{X}, \mathcal{B}, \mu)$ without atoms, such that all these transformations yield a prevalent set of cyclic vectors. Note that for such spaces $L^2_\mu(\mathbb{X})$ is separable. From [53, §3.3] we know that ergodic group rotations have discrete spectrum. Thus, also the dense set of transformations constructed in Theorem 23 have discrete spectrum, and by Proposition 14 all observables for these transformations are weakly predictive.

It is natural to ask whether for an ergodic transformation the possession of discrete spectrum already implies that the set of cyclic vectors is prevalent. The Representation Theorem [53, Theorem 3.6] states that an ergodic measure-preserving transformation with discrete spectrum on a probability space is isomorphic to an ergodic rotation on a compact abelian group. The strong analogy between (ergodic) circle rotations and group rotations, together with the prevalence statement of Proposition 19 could lead one to formulate the following conjecture:

Conjecture 25. *Let the ergodic T in MPT possess discrete spectrum. Then $\mathcal{U} = \mathcal{U}_T$ has a prevalent set of cyclic vectors in L^2_μ .*

In Sect. 3.5 we will discuss another type of systems that possesses cyclic vectors.

3.4. A generic MPT has a dense G_δ set of cyclic vectors. By Theorem 23 the collection of measure-preserving transformations with a dense G_δ set of cyclic vectors lies dense in MPT with respect to the weak topology. In fact, we can even prove that it is a dense G_δ set in MPT. We use similar methods as in [67, Theorem 8.26], where it is shown that the collection of measure-preserving transformations admitting a cyclic vector is a dense G_δ set in MPT with respect to the weak topology.

Proposition 26. *The collection of μ -preserving transformations with a dense G_δ set of cyclic vectors in L^2_μ is a dense G_δ set in MPT with respect to the weak topology.*

Proof. By Theorem 23 we already know that the collection of measure-preserving transformations with a dense G_δ set of cyclic vectors is dense in MPT.

To show that it is a G_δ set, we let $\{f_j\}_{j \in \mathbb{N}}$ be a dense set in L^2_μ . For $g \in L^2_\mu$, $m, N \in \mathbb{N}$ we introduce the set

$$\begin{aligned} \tilde{V}_i(m, N, g) &= \left\{ \mathcal{U} \in \mathcal{B}(L^2_\mu) \mid \exists \lambda_k^{(i)} \in \mathbb{C}, k = 0, \dots, N : \left\| f_i - \sum_{k=0}^N \lambda_k^{(i)} \mathcal{U}^k g \right\|_{L^2_\mu} < \frac{1}{m} \right\} \\ &= \bigcup_{p \in P_N} \left\{ \mathcal{U} \in \mathcal{B}(L^2_\mu) \mid \|f_i - p(\mathcal{U})g\|_{L^2_\mu} < \frac{1}{m} \right\}, \end{aligned}$$

where P_N denotes the set of polynomials of degree not larger than N . Such a set $\tilde{V}_i(m, N, g)$ is open in the space $\mathcal{B}(L^2_\mu)$ of bounded linear operators with respect to

the strong operator topology because $\mathcal{U}_n \rightarrow \mathcal{U}$ implies $p(\mathcal{U}_n) \rightarrow p(\mathcal{U})$ for every polynomial p by Lemma 18(5). Since the weak topology on MPT coincides with the topology induced from the strong operator topology on Koopman operators (cf. Lemma 18), we see that the sets

$$V_i(m, N, g) = \left\{ T \in \text{MPT} \mid \exists \lambda_k^{(i)} \in \mathbb{C}, k = 0, \dots, N : \left\| f_i - \sum_{k=0}^N \lambda_k^{(i)} \mathcal{U}_T^k g \right\|_{L^2_\mu} < \frac{1}{m} \right\}.$$

are open with respect to the weak topology. Thus, the finite intersections

$$\begin{aligned} V(m, n, N, g) &:= \bigcap_{i=1}^n V_i(m, N, g) \\ &= \left\{ T \in \text{MPT} \mid \forall i = 1, \dots, n \exists \lambda_k^{(i)} \in \mathbb{C}, k = 0, \dots, N : \left\| f_i - \sum_{k=0}^N \lambda_k^{(i)} \mathcal{U}_T^k g \right\|_{L^2_\mu} < \frac{1}{m} \right\} \end{aligned}$$

are also open. With this terminology, the collection of mpts with a dense G_δ set of cyclic vectors is given by the G_δ set

$$\mathcal{T} := \bigcap_{m \in \mathbb{N}} \bigcap_{n \in \mathbb{N}} \bigcap_{t \in \mathbb{N}} \bigcap_{s \in \mathbb{N}} \bigcup_{N \in \mathbb{N}} \bigcup_{g \in L^2_\mu, \|g - f_s\|_{L^2_\mu} < \frac{1}{t}} V(m, n, N, g),$$

where it remains to show that $T \in \mathcal{T}$ has a dense G_δ set of cyclic vectors in L^2_μ . Heuristically, this dense G_δ set of cyclic vectors is obtained by the intersections over t and s which guarantee for any function f_s from our dense family and any $t \in \mathbb{N}$ the existence of a cyclic vector that is $\frac{1}{t}$ -close to f_s . To show this in detail, let $B_r(f) \subset L^2_\mu$ denote the open ball with radius $r > 0$ and center $f \in L^2_\mu$ and consider

$$\begin{aligned} \mathcal{T}(m, n) &= \bigcap_{t \in \mathbb{N}} \bigcap_{s \in \mathbb{N}} \bigcup_{N \in \mathbb{N}} \bigcup_{g \in L^2_\mu, \|g - f_s\|_{L^2_\mu} < \frac{1}{t}} V(m, n, N, g) \\ &= \left\{ T \in \text{MPT} \mid \forall s, t \in \mathbb{N} \exists g \in B_{1/t}(f_s) \text{ and } \exists N \in \mathbb{N} \text{ s.t. } \forall i = 1, \dots, n \right. \\ &\quad \left. \exists p \in P_N \text{ with } \|p(\mathcal{U}_T)g - f_i\|_{L^2_\mu} < \frac{1}{m} \right\}. \end{aligned}$$

The functions $g \in L^2_\mu$ that appear in the characterization of $\mathcal{T}(m, n)$ are candidates for cyclic vectors upon intersection over all $m, n \in \mathbb{N}$. We denote this set by $C(T, m, n) \subset L^2_\mu$, more precisely, for a fixed $T \in \mathcal{T}(m, n)$ let

$$C(T, m, n) := \bigcap_{t \in \mathbb{N}} \bigcup_{s \in \mathbb{N}} \left[B_{1/t}(f_s) \cap \left\{ g \in L^2_\mu \mid \begin{array}{l} \exists N \in \mathbb{N} \text{ s.t. } \forall i = 1, \dots, n \exists p \in P_N \\ \text{s.t. } \|p(\mathcal{U}_T)g - f_i\|_{L^2_\mu} < \frac{1}{m} \end{array} \right\} \right].$$

For $T \in \mathcal{T} = \bigcap_{m, n \in \mathbb{N}} \mathcal{T}(m, n)$, the set of cyclic vectors is given by $\bigcap_{m, n \in \mathbb{N}} C(T, m, n)$. To see that this is a G_δ set, it is sufficient to show that the $C(T, m, n)$ are themselves G_δ for every $m, n \in \mathbb{N}$. This is straightforward by noting that $p(\mathcal{U}_T)$ is a continuous mapping on L^2_μ for every $p \in P_N$ and that

$$C(T, m, n) = \bigcap_{t \in \mathbb{N}} \bigcup_{s \in \mathbb{N}} \left[B_{1/t}(f_s) \cap \bigcap_{i=1}^n \bigcup_{N \in \mathbb{N}} \bigcup_{p \in P_N} p(\mathcal{U}_T)^{-1} B_{1/m}(f_i) \right].$$

Arbitrary unions and finite intersections of open sets are open, hence $C(T, m, n)$ and with it the set of cyclic vectors for $T \in \mathcal{T}$ is G_δ , concluding the proof. \square

While this result gives us the genericity of transformations with cyclic vectors, it does not provide us with a criterion to check if a given transformation has cyclic vectors. In the following section we meet such a sufficient condition to guarantee the existence of cyclic vectors.

3.5. Rank one systems. In this section we discuss *rank one systems* which constitute a generic class of measure-preserving transformations possessing cyclic vectors. We start with the admittedly technical definition of rank one systems.³

Definition 27 (*Rank one system*). A system $T : (\mathbb{X}, \mu) \rightarrow (\mathbb{X}, \mu)$ is of *rank one* if for any measurable set $A \subset \mathbb{X}$ and any $\varepsilon > 0$ there exist $F \subset \mathbb{X}$, $h \in \mathbb{Z}^+$ and a measurable subset $A' \subset \mathbb{X}$ such that

- the sets $T^k F$, $k = 0, 1, \dots, h - 1$ are disjoint;
- $\mu(A \Delta A') < \varepsilon$;
- $\mu(\bigcup_{k=0}^{h-1} T^k F) > 1 - \varepsilon$;
- A' is measurable with respect to the partition $\xi(T)$ formed by the sets $F, TF, \dots, T^{h-1}F$ and $\mathbb{X} \setminus \bigcup_{k=0}^{h-1} T^k F$.

One calls $\mathcal{T} = \{F, TF, \dots, T^{h-1}F\}$ a *Rokhlin tower* (or *column*) of height h and with base F .

Thus, a system is rank one if for any fixed set and a tolerance one can choose some set that has disjoint iterates which achieve filling the space and approximating the first set both within the tolerance. Furthermore, one says that a sequence of towers \mathcal{T}_n is *exhaustive* if $\xi(\mathcal{T}_n)$ converges to the decomposition into points as $n \rightarrow \infty$, that is, for every measurable set $A \subseteq \mathbb{X}$ and for every $n \in \mathbb{N}$ there exists a set A_n , which is a union of elements of $\xi(\mathcal{T}_n)$, such that $\lim_{n \rightarrow \infty} \mu(A \Delta A_n) = 0$. By definition, rank one systems have exhaustive sequences of towers.

If $T \in \text{MPT}$ is rank one, so is T^{-1} . To see this, note that for every $\varepsilon > 0$ and associated tower \mathcal{T} with base F for T we get the same tower with base $T^{h-1}F$, now for T^{-1} , satisfying the requirements of Definition 27.

Examples of rank one systems include rotations and odometer transformations; see our example in Sect. 5.4. In general, it is true that discrete spectrum together with ergodicity implies rank one [69]. Rank one systems were introduced and actively studied in the framework of so-called *cutting-and-stacking constructions* in ergodic theory (see e.g. [70, section 5.2]). Here, one can view a rank one system as a transformation obtained from a cutting-and-stacking construction with a single tower but there is a scarcity of explicit examples.

Our interest in the rank one property is based on the fact that it is a sufficient condition to guarantee the existence of cyclic vectors.⁴

³ According to Ferenczi it is the lecturer's nightmare to define rank one systems [68, p. 40] We refer to his survey article for several equivalent definitions of the rank one property.

⁴ We also mention that there is a general notion of *rank* of a measure-preserving transformation. This rank is always greater than or equal to the spectral multiplicity of the associated Koopman operator. Since we are mainly interested in systems with simple spectrum, we focus on the case of rank one. We refer the interested reader to [67, chapter 5].

Proposition 28. *A rank one system has cyclic vectors.*

Proof. We combine arguments from [67, chapter 4] and [70, Proposition 5.8].

Let T be a rank one transformation. Recall that $C_g = \text{span}\{g, \mathcal{U}_T g, \mathcal{U}_T^2 g, \dots\}$ denotes the cyclic subspace of $g \in L^2_\mu$. Furthermore, $d_{L^2_\mu}(f, C_g)$ denotes the L^2_μ distance of $f \in L^2$ from C_g .

Suppose that T does not admit cyclic vectors. Then there exist orthogonal unit vectors $f_1, f_2 \in L^2_\mu$ such that

$$d_{L^2_\mu}^2(f_1, C_g) + d_{L^2_\mu}^2(f_2, C_g) \geq 1 \tag{13}$$

for every $g \in L^2_\mu$ (see [71, Lemma 3.1] or [67, Theorem 4.3] for a proof of a more general result). As in the definition of the rank one property, we consider an exhaustive sequence of towers \mathcal{T}_n with base F_n and height h_n approximating T . Then for each n the images of the characteristic function χ_{F_n} under \mathcal{U}_T^i , $i = 0, 1, \dots, h_n - 1$, are characteristic functions of the disjoint levels of the tower \mathcal{T}_n . Then for each n there is a cyclic subspace $C_{\chi_{F_n}}$ which contains all characteristic functions of the levels of the tower and their linear combinations. Since our sequence of towers is exhaustive, we can choose n sufficiently large such that $d_{L^2_\mu}(f_1, C_{\chi_{F_n}}) < \frac{1}{2}$ and $d_{L^2_\mu}(f_2, C_{\chi_{F_n}}) < \frac{1}{2}$. We obtained a contradiction to (13). Hence, T has cyclic vectors. \square

In the converse direction to Footnote 4, the simplicity of the spectrum does not imply the rank one property; see [68, section 2.4.2] for various types of counterexamples. Interestingly, there are also mixing systems of rank one; see e.g. [68, section 1.4.4] for a celebrated construction by Ornstein. We would like to emphasize the following:

Remark 29. (Mixing can be predictive) On the one hand, measure-theoretic (strong) mixing implies gaining asymptotic independence on the past. Intuitively, this could lead to the impossibility of asymptotic predictiveness, since the observation $f \circ T^{-k}$ gives “increasingly less information” about $f \circ T$ as k grows. On the other hand, the existence of mixing systems that are rank one—and thus possess cyclic vectors—tells us that this intuition is not necessarily correct (recall also the discussion after Proposition 14).

Given the significance of cyclic vectors for predictiveness, it is an interesting question if a rank one system has a prevalent set of cyclic vectors. We do not provide an answer to this question, but note that if it were positive, then also Conjecture 25 would be true, as ergodic systems with discrete spectrum are rank one.

We provide a proof of the following folklore result by adapting proofs of similar results in [71, Theorem 1.1] and [72, Theorem 3.1].

Proposition 30. *The collection of rank one systems is a dense G_δ set in MPT with respect to the weak topology.*

Proof. Let $\{E_j\}_{j \in \mathbb{N}}$ be a dense collection of measurable subsets in $[0, 1]$. We define \mathcal{A}_n as the set of transformations $T \in \text{MPT}$ that admit a tower approximating E_1, \dots, E_n with accuracy $1/n$. We claim that $\mathcal{R} := \bigcap_{n \in \mathbb{N}} \mathcal{A}_n$ is a dense G_δ set and coincides with the collection of rank one systems.

To see that every $T \in \mathcal{R}$ has rank one, we let A be a measurable set and $\varepsilon > 0$. Since $\{E_j\}_{j \in \mathbb{N}}$ is dense, there is E_k with $\mu(A \Delta E_k) < \varepsilon/2$. Let $n > \max(k, 2/\varepsilon)$. Since $T \in \mathcal{A}_n$, there is a tower approximating E_k with accuracy $1/n < \varepsilon/2$. Hence, the tower approximates A with accuracy ε .

In the next step we show that each \mathcal{A}_n is dense. For this purpose, we choose $K \in \mathbb{N}$ sufficiently large such that for all $k \geq K$ each E_1, \dots, E_n can be approximated by unions of dyadic intervals

$$I_{k,i} := \left[\frac{i}{2^k}, \frac{i+1}{2^k} \right), \quad i = 0, 1, \dots, 2^k - 1,$$

with accuracy $1/n$. On the partition $\{I_{k,i} \mid i = 0, 1, \dots, 2^k - 1\}$ we consider cyclic permutations of order k , that is, maps cyclically permuting the dyadic intervals of rank k by translations. We denote the collection of cyclic permutations of order k by Π_k . In particular, each such permutation admits a tower of exactly the dyadic intervals of rank k . Hence, $\Pi_k \subset \mathcal{A}_n$. By the so-called *Weak Approximation Theorem* from [65, p. 65] the collection $\bigcup_k \Pi_k$ of cyclic permutations is dense in MPT with respect to the weak topology. Since all but finitely many of the cyclic permutations are in \mathcal{A}_n , \mathcal{A}_n still contains a dense set.

One can also verify that each \mathcal{A}_n is open. Then the Baire category theorem yields that \mathcal{R} is a dense G_δ set. □

3.6. A generic measure-preserving homeomorphism has a dense G_δ set of cyclic vectors.

In this section we want to extend our results to the setting of measure-preserving homeomorphisms. Let \mathbb{M} be a compact connected manifold and μ be a so-called *OU measure*.⁵, that is, μ is nonatomic, locally positive (i.e., it is positive on every nonempty open set), and zero on the manifold boundary. We consider the collection $\text{Homeo}(\mathbb{M}, \mu)$ of μ -preserving homeomorphisms of \mathbb{M} with the uniform topology defined by the complete metric

$$d_{\text{unif}}(f, g) = \|f - g\|_{\text{unif}} = \max_{x \in \mathbb{M}} d(f(x), g(x)) + \max_{x \in \mathbb{M}} d(f^{-1}(x), g^{-1}(x)),$$

where d is a given metric on \mathbb{M} compatible with its topology.⁶

In this setting, any measure-theoretic property which is generic for abstract measure-preserving transformations is also generic for measure-preserving homeomorphisms of compact manifolds. This can be stated as follows.

Lemma 31 ([73, Corollary 10.4]). *Let μ be an OU measure on a compact connected manifold \mathbb{M} . Let \mathcal{V} be a conjugate invariant dense G_δ subset of MPT with the weak topology. Then $\mathcal{V} \cap \text{Homeo}(\mathbb{M}, \mu)$ is a dense G_δ subset of $\text{Homeo}(\mathbb{M}, \mu)$ with the uniform topology.*

Since the collection of mpts with a dense G_δ set of cyclic vectors is conjugate invariant by Lemma 21, we can apply Lemma 31 on Proposition 26. This gives the following genericity result for homeomorphisms with cyclic vectors.

Proposition 32. *Let \mathbb{M} be a compact connected manifold and μ be an OU measure. The collection of μ -preserving homeomorphisms on \mathbb{M} with a dense G_δ set of cyclic vectors in L^2_μ is a dense G_δ set in $\text{Homeo}(\mathbb{M}, \mu)$ with the uniform topology.*

Thus, a generic homeomorphism admits a generic set of predictive L^2 observables. Note that Takens-type theorems require observables to be continuous.

⁵ The name OU measure was coined in [73] in honor of Oxtoby and Ulam. Sometimes these measures are also called *good measures*.

⁶ Note that Urysohn’s metrization theorem implies that every manifold is metrizable. On compact spaces, all such metrics inducing the topology are uniformly equivalent. Hence, they also induce the same uniform topology on $\text{Homeo}(\mathbb{M}, \mu)$.

4. Properties of the Least Squares Filter

4.1. *Forecast properties.* Beyond the prediction error considered in Proposition 6, a further dynamic object of interest is the autocorrelation function; that is

$$A(n) := \langle \mathcal{U}^n f, f \rangle_{L^2_\mu} = \int f \circ T^n f \, d\mu.$$

It turns out that autocorrelations for $n \leq d$ are always reproduced correctly, irrespective of whether the observation function is a cyclic vector of \mathcal{U} or not.

To show this, let us define the shorthand notations $\Pi_{\mathcal{V}}$ (or Π_h) for the orthogonal projection in L^2_μ to a subspace $\mathcal{V} \subset L^2_\mu$ (or to the onedimensional subspace spanned by $h \in L^2_\mu$) and $\mathcal{K}_d^- := \text{span}\{f \circ T^{-d+1}, \dots, f \circ T^{-1}\}$. Note, in particular, that $\mathcal{K}_d^- = \mathcal{U}^{-1}\mathcal{K}_{d-1}$ and $\mathcal{K}_d = \mathcal{K}_d^- + \text{span}\{f\}$. Let $\Pi_{\mathcal{V}}^\perp := \text{id} - \Pi_{\mathcal{V}}$ denote the orthogonal projection on the orthogonal complement of \mathcal{V} .

Proposition 33. *Let $A_d(n) := \langle \mathcal{U}_d^n f, f \rangle_{L^2_\mu}$ denote the filter’s autocorrelation. We have for every map T and observable f that*

$$|A(n) - A_d(n)| \leq \begin{cases} 0, & n \leq d, \\ \|\mathcal{U}_d f - \mathcal{U} f\|_{L^2_\mu} \|f\|_{L^2_\mu}, & n = d + 1, \\ (n - d) \frac{\|\Pi_{\mathcal{K}_d^-}^\perp \mathcal{U} f\|_{L^2_\mu}}{\|\Pi_{\mathcal{K}_d^-}^\perp f\|_{L^2_\mu}} \|f\|_{L^2_\mu}^2, & n \geq d + 2. \end{cases}$$

Proof. We can express the statement of Proposition 3 as

$$\langle \mathcal{U}_d g - \mathcal{U} g, f \circ T^{-i} \rangle_{L^2_\mu} = 0 \quad \forall g \in \mathcal{K}_d, \quad i = 0, \dots, d - 1. \tag{14}$$

For $n = 1$, the claim $A(n) = A_d(n)$ translates to

$$\langle \mathcal{U}_d f - \mathcal{U} f, f \rangle_{L^2_\mu} = 0.$$

Note that (14) implies this with $g = f$ and $i = 0$. By expressing the difference $\mathcal{U}_d^n - \mathcal{U}^n$ as a telescopic sum, we can extend this result to higher powers n :

$$\begin{aligned} A_d(n) - A(n) &= \sum_{i=0}^{n-1} \left\langle \mathcal{U}^i \mathcal{U}_d^{n-i} f - \mathcal{U}^{i+1} \mathcal{U}_d^{n-i-1} f, f \right\rangle_{L^2_\mu} \\ &= \sum_{i=0}^{n-1} \langle \mathcal{U}_d^{n-i} f - \mathcal{U} \mathcal{U}_d^{n-1-i} f, f \circ T^{-i} \rangle_{L^2_\mu} \\ &= \sum_{i=0}^{n-1} \langle (\mathcal{U}_d - \mathcal{U}) \mathcal{U}_d^{n-1-i} f, f \circ T^{-i} \rangle_{L^2_\mu} \\ &= \sum_{i=0}^{d-1} \langle (\mathcal{U}_d - \mathcal{U}) \mathcal{U}_d^{n-1-i} f, f \circ T^{-i} \rangle_{L^2_\mu} \\ &\quad + \sum_{i=d}^{n-1} \left\langle \mathcal{U}^i (\mathcal{U}_d - \mathcal{U}) \mathcal{U}_d^{n-1-i} f, f \right\rangle_{L^2_\mu}, \end{aligned} \tag{15}$$

where we used $\langle \mathcal{U}^i h, f \rangle_{L_\mu^2} = \langle h, f \circ T^{-i} \rangle_{L_\mu^2}$ for any $h \in L_\mu^2$, implied by the invariance of μ . Since $i \leq n - 1$, note that $g := \mathcal{U}_d^{n-i-1} f \in \mathcal{K}_d$ and (14) implies that for $i = 0, \dots, d - 1$,

$$\langle (\mathcal{U}_d - \mathcal{U})\mathcal{U}_d^{n-i-1} f, f \circ T^{-i} \rangle_{L_\mu^2} = 0.$$

The first term on the right-hand side of (15) is hence zero. Thus, $A_d(n) = A(n)$ for $n \leq d$, as the second sum on the right-hand side of (15) is empty.

For $n = d + 1$ the second sum on the right-hand side of (15) reduces to

$$\left\langle \mathcal{U}^d (\mathcal{U}_d - \mathcal{U})\mathcal{U}_d^0 f, f \right\rangle_{L_\mu^2}.$$

The inequality for $n = d + 1$ thus follows with $\|\mathcal{U}\| = 1$. Note that since \mathcal{U} is unitary and \mathcal{U}_d is the orthogonal projection of \mathcal{U} , by the non-expansiveness of orthogonal projections we have that $\|\mathcal{U}_d\| \leq 1$.

For $n \geq d + 2$ we can estimate the terms in the second sum on the right-hand side of (15) by

$$\left| \left\langle \mathcal{U}^i (\mathcal{U}_d - \mathcal{U})\mathcal{U}_d^{n-1-i} f, f \right\rangle_{L_\mu^2} \right| \leq \|\mathcal{U}\|^i \|(\mathcal{U}_d - \mathcal{U})|_{\mathcal{K}_d}\| \|\mathcal{U}_d\|^{n-1-i} \|f\|_{L_\mu^2}^2. \quad (16)$$

It remains to bound $\|(\mathcal{U}_d - \mathcal{U})|_{\mathcal{K}_d}\|$. We note that for every $f^- \in \mathcal{K}_d^-$ we have that $\mathcal{U}_d f^- = \mathcal{U} f^-$ because $\mathcal{U} f^- \in \mathcal{K}_d$. Thus, we can write

$$\|(\mathcal{U}_d - \mathcal{U})|_{\mathcal{K}_d}\| = \sup_{\substack{g \in \mathcal{K}_d \\ g \neq 0}} \frac{\|\mathcal{U}_d g - \mathcal{U} g\|_{L_\mu^2}}{\|g\|_{L_\mu^2}} = \sup_{\substack{g=f+f^- \\ f^- \in \mathcal{K}_d^-}} \frac{\|\mathcal{U}_d g - \mathcal{U} g\|_{L_\mu^2}}{\|g\|_{L_\mu^2}}, \quad (17)$$

where the second equation holds because for any g without a contribution from f (i.e., $g \in \mathcal{K}_d^-$) the fraction on the right is zero, clearly dominated by the supremum as written. We obtain

$$\|(\mathcal{U}_d - \mathcal{U})|_{\mathcal{K}_d}\| = \sup_{f^- \in \mathcal{K}_d^-} \frac{\|\mathcal{U}_d f - \mathcal{U} f\|_{L_\mu^2}}{\|f + f^-\|_{L_\mu^2}} = \sup_{f^- \in \mathcal{K}_d^-} \frac{\|\Pi_{\mathcal{K}_d}^\perp \mathcal{U} f\|_{L_\mu^2}}{\|f + f^-\|_{L_\mu^2}} = \frac{\|\Pi_{\mathcal{K}_d}^\perp \mathcal{U} f\|_{L_\mu^2}}{\|\Pi_{\mathcal{K}_d}^\perp f\|_{L_\mu^2}}.$$

The claim follows by applying the bound (16) to every term in the sum. □

Thus, if f is predictive (in particular, if it is a cyclic vector of \mathcal{U}), then even

$$\lim_{d \rightarrow \infty} A_d(d + 1) = A(d + 1).$$

By the Wiener–Khinchin theorem, autocorrelation is related to the notion of power spectrum through Fourier transformation [74]. Since Proposition 33 does not indicate anything about autocorrelation beyond a finite horizon, statements about the reconstruction of the power spectrum seem to be elusive. Methods more powerful than the least squares filter can approximate the power spectrum, e.g. [75, §4.5]. In the next section we consider spectral approximation of \mathcal{U} . We note that the rightmost term in (17) will reappear in the pseudospectral bound given there.

4.2. *Some spectral properties.* We start with the following trivial observation:

Proposition 34. *The spectrum of U_d is contained in the unit complex disk.*

Proof. As in the proof of Proposition 33: Since \mathcal{U} is unitary and \mathcal{U}_d is the orthogonal projection of \mathcal{U} , by the non-expansiveness of orthogonal projections we have that $\|\mathcal{U}_d\| \leq 1$, implying the claim. \square

Under additional assumptions, stronger statements about the position and the distribution of the spectrum of U_d can be found in [46, Corollaries 2 and 3].

How close is the spectrum of \mathcal{U}_d to that of \mathcal{U} ? The theory of pseudospectra [76] for normal operators offers a way to approach this question. For simplicity of notation, we use $\|\cdot\| = \|\cdot\|_{L^2_\mu}$ for the remainder of this section. A scalar $\lambda \in \mathbb{C}$ is said to be in the ε -pseudospectrum of \mathcal{U} if there is a $h \in L^2_\mu$ with $\|h\| = 1$ such that $\|\mathcal{U}h - \lambda h\| < \varepsilon$. Since \mathcal{U} is unitary, hence normal, every value in the ε -pseudospectrum is at most distance ε away from a true spectral value [76].

One can now take an eigenpair (λ_d, ϕ_d) of \mathcal{U}_d with $\|\phi_d\| = 1$ and estimate $\|\mathcal{U}\phi_d - \lambda_d\phi_d\|$ to assess how far is λ_d from the spectrum of \mathcal{U} . Recall the notation introduced prior to Proposition 33. With it, we have:

Proposition 35. *Let λ_d be an eigenvalue of \mathcal{U}_d . Then λ_d is contained in the ε -pseudospectrum of \mathcal{U} , where*

$$\varepsilon = \frac{\|\Pi_{\mathcal{K}_d}^\perp \mathcal{U}f\|}{\|\Pi_{\mathcal{K}_d}^\perp f\|}. \tag{18}$$

Proof. Let (λ_d, ϕ_d) be an eigenpair of \mathcal{U}_d with $\|\phi_d\| = 1$. Note that we seek $\varepsilon = \|\mathcal{U}\phi_d - \lambda_d\phi_d\|$. Let us write $\phi_d = af^- + bf$ with $f^- \in \mathcal{K}_d^-$ and $a, b \in \mathbb{C}$. Then, since $\lambda_d\phi_d = \mathcal{U}_d\phi_d$, we have that

$$\varepsilon = \|(\mathcal{U} - \Pi_{\mathcal{K}_d} \mathcal{U})\phi_d\| = \|\Pi_{\mathcal{K}_d}^\perp \mathcal{U}(bf)\|,$$

because $\Pi_{\mathcal{K}_d}^\perp \mathcal{U}f^- = 0$. We thus need to bound $|b|$ from above, given $\|af^- + bf\| = 1$. This can be done by basic trigonometry. Let us consider the triangle with corners 0, af^- , and ϕ_d , which then have the opposite “sides” of lengths $\|bf\|$, $\|\phi_d\| = 1$, and $\|af^-\|$, respectively. Denote the angle at the corner 0 by β and the one at af^- by θ . It is then immediate that

$$\sin \theta = \frac{\|\Pi_{f^-}^\perp \phi_d\|}{\|bf\|} \geq \frac{\|\Pi_{\mathcal{K}_d}^\perp \phi_d\|}{\|bf\|} = \frac{\|\Pi_{\mathcal{K}_d}^\perp f\|}{\|f\|},$$

where we used $\text{span}\{f^-\} \subset \mathcal{K}_d^-$ for the inequality and $\Pi_{\mathcal{K}_d}^\perp f^- = 0$ for the equality. The sine law for this triangle yields

$$\frac{\sin \theta}{\|\phi_d\|} = \frac{\sin \beta}{\|bf\|},$$

giving $|b| = \frac{\sin \beta}{\|f\| \sin \theta}$. Since $\sin \beta \leq 1$, combining this with the above proves the claim. \square

Remark 36.

- (a) Note that in the expression (18) for ε , the denominator represents the error to approximate f by the $(d - 1)$ observables $f \circ T^{1-d}, \dots, f \circ T^{-1}$ preceding it, and the numerator represents the error to approximate $\mathcal{U}f$ by the d observables $f \circ T^{1-d}, \dots, f$ preceding it.
- (b) In [77, Lemma 2], there was an attempt to make a similar pseudospectral statement. The result, however, is questionable, since its proof seems to be incorrect: Their pseudo-eigenfunction $\tilde{\mathbf{e}} \cdot \tilde{\mathbf{f}}$ does not have L^2_μ -norm equal to 1, in general. This also impairs [77, Theorem 5].

Despite the appealing form of (18), we do not seem to be able to draw practical conclusions from it for $d \rightarrow \infty$. With going to a subsequence, however, we can make connections to the point spectrum. The following is a consequence of [78, Theorem 4], by noting that \mathcal{U}_d is the L^2_μ -orthogonal projection on \mathcal{K}_d .

Proposition 37. *Let (λ_d, ϕ_d) be eigenpairs of \mathcal{U}_d for $d \in \mathbb{N}$ with $\|\phi_d\| = 1$. If f is a cyclic vector of \mathcal{U} , then there is a subsequence $(d_k)_k \subset \mathbb{N}$ such that $\phi_{d_k} \xrightarrow{w} \phi$ and $\lambda_{d_k} \rightarrow \lambda$, and $\mathcal{U}\phi = \lambda\phi$. In particular, if $\phi \neq 0$, then (λ, ϕ) is an eigenpair of \mathcal{U} .*

At this point it is important to note that other methods than the least squares filter allow to make stronger claims about faithful approximation of (parts of) the spectrum of \mathcal{U} , see [46–51].

4.3. Approximate filter. Both Theorem 23 and Proposition 26 imply that the set of systems that are predictive for a large set of observables is weakly dense in MPT. As the weak topology is metrizable [65, p. 64], there is some predictive map “arbitrary close” to our original system T . Recall, however, that a practical interpretation of this metric seems hardly possible according to Halmos. To remedy this fact somewhat, we will consider this closeness statement on the level of filters. In particular, we show in Proposition 38 below that there is a predictive system that produces a filter arbitrary close to that produced by T .

Note that, by ergodicity of T , for almost every $x \in \mathbb{X}$ we obtain the same filter c (cf. Lemma 1). A quick inspection of the structure of the objective function J in (4) reveals that the actual observation sequence is of little relevance for the filter c ; what matters are the correlation integrals $\int (f \circ T^{-i})(f \circ T^{-j}) d\mu$, $i, j = -1, 0, \dots, d - 1$. Note that Lemma 18(5) gives for a sequence $(T_n)_n \subset \text{MPT}$ with $T_n \xrightarrow{w} T$ that

$$\lim_{n \rightarrow \infty} \int (f \circ T_n^{-i})(f \circ T_n^{-j}) d\mu = \int (f \circ T^{-i})(f \circ T^{-j}) d\mu \tag{19}$$

for all $i, j = -1, 0, \dots, d - 1$. Since for a fixed $d \in \mathbb{N}$ there are finitely many pairs (i, j) to consider, this convergence is uniform across these pairs. Thus, Theorem 23, Equation (19), together with Lemma 1 gives

Proposition 38. *Let $T \in \text{MPT}$ and $d \in \mathbb{N}$ a fixed delay depth. Let $f \in L^2_\mu$ be an observable such that \mathcal{B}_d is a basis. Then there are predictive $S \in \text{MPT}$ that produce filters arbitrary close to that of T .*

Remark 39. In particular, Proposition 38 implies that based on the filter alone, which is computed from (finite or infinite) observation data, one is not able to infer whether the

system at hand is predictive. This also implies that the spectrum of \mathcal{U}_d (which is for a fixed d a continuous function of the filter vector c) is not able to distinguish whether the observed time series comes from a predictive system.

We stress that Proposition 38 is a statement for a fixed delay depth, and in particular does not imply anything about the forecast error $\|\mathcal{U}_d f - \mathcal{U} f\|_{L^2_\mu}$ in the limit $d \rightarrow \infty$.

5. Numerical Examples

We illustrate our results with a variety of examples properties of the linear least squares filter. In particular, we will consider the prediction error and the autocorrelations.

5.1. Numerical setup. The filter is learned on (observations of) a long “training” trajectory of length $(m + 1)$, cf. (2). The mean squared forecast error is computed on a different “testing” trajectory, of the same length $(m + 1)$, by sliding the filter along the trajectory and comparing the prediction with the next (the true) element on the trajectory. By ergodicity, the mean square error thus approximates $\|\mathcal{U}_d f - \mathcal{U} f\|_{L^2_\mu}^2$ in the limit of $m \rightarrow \infty$.

The autocorrelations are compared as follows. We generate a large sample of μ -distributed initial points $x_1^{(i)}$, $i = 1, \dots, N$. This is done either by random uniform sampling, as in most cases it is the unique ergodic distribution, or by a long trajectory of a rationally independent stepsize compared with the one used for learning and testing, in the case of the one time-continuous system. From each of these initial conditions we generate trajectory “snippets” of length d , denoted by $x_1^{(i)}, x_2^{(i)} = T x_1^{(i)}, \dots, x_d^{(i)} = T^{d-1} x_1^{(i)}$, and continue each snippet in two ways:

- (i) By the true map, and apply the observation map to the elements at the end, giving

$$f(x_1^{(i)}), f(x_2^{(i)}), \dots, f(x_d^{(i)}), f(x_{d+1}^{(i)} = T x_d^{(i)}), f(x_{d+2}^{(i)} = T^2 x_d^{(i)}), \dots$$

- (ii) By applying the filter iteratively to the observed trajectory “snippet” to generate an observation sequence of the same length as in (i), giving

$$z_1^{(i)} = f(x_1^{(i)}), z_2^{(i)} = f(x_2^{(i)}), \dots, z_d^{(i)} = f(x_d^{(i)}), z_{d+1}^{(i)}, z_{d+2}^{(i)}, \dots,$$

where $z_n^{(i)}$, $n > d$, is obtained by the filter applied to $z_{n-d}^{(i)}, \dots, z_{n-1}^{(i)}$.

We then approximate the autocorrelations of the system and the filter by

$$A(n) = \frac{1}{N} \sum_{i=1}^N f(x_d^{(i)}) f(x_{d+n}^{(i)}),$$

$$A_d(n) = \frac{1}{N} \sum_{i=1}^N z_d^{(i)} z_{d+n}^{(i)},$$

respectively. Unless stated differently, we use $m = N = 10^4$ in the examples.




Fig. 1. Ergodic torus rotation, results for observable f_1 . Left: L^2 prediction error. Right: autocorrelations




Fig. 2. Ergodic torus rotation. L^2 prediction error for observable f_2

5.2. Ergodic torus rotation. First we consider the ergodic rotation of the unit 2-torus $S^1 \times S^1$:

$$T(x) = \begin{pmatrix} x_1 + \sqrt{2} \\ x_2 + \sqrt{3} \end{pmatrix} \pmod 1.$$

We use the observables $f_1(x) = \exp(\sin(4\pi x_1) + \cos(6\pi x_2))$ and $f_2 = \chi_{[0, 1/2] \times [1/2, 1]}$, where χ_A denotes the characteristic function of the set A .

In the mean squared forecast error we observe exponential convergence in the number of delays, see Figs. 1 (left) and 2. Note the scale on the vertical axes: The convergence is slower the rougher the observable is, and in particular for non-continuous functions. Why in Fig. 2 around $d = 35$ the sudden slow-down occurs is yet to be understood. In Fig. 1 (right) we see that the filter’s autocorrelations $A_d(n)$ are exact up to the chosen delay depth ($d = 21$ there), and qualitatively keep the correct behavior even up to $n = 3d$.

5.3. The affine twist map. We consider

$$T(x) = \begin{pmatrix} x_1 + \sqrt{2} \\ x_1 + x_2 \end{pmatrix} \pmod 1,$$




Fig. 3. L^2 prediction error (left) and autocorrelations for the affine twist map (right)

again on the unit 2-torus. This map has mixed spectrum: Its Kronecker factor⁷ is the irrational circle rotation $R_{\sqrt{2}}$, the spectrum in the orthogonal complement to this factor is countable Lebesgue [70, Example 3.17]: Roughly speaking this part of the spectrum can be represented by the Lebesgue measure on \mathbb{T} and has everywhere infinite multiplicity. In summary, the affine twist map has so-called *quasi-discrete spectrum*.⁸

We use the observable $f(x) = \exp(2 \sin(4\pi x_1) + \cos(6\pi x_2))$.

In Fig. 3 (left), we observe that at some point the prediction error ceases to improve with increasing delay. This indicates that (T, f) is not predictive (not even weakly).

Unsurprisingly, reducing the dependence of the observable f on the “mixing coordinate” x_2 improves the prediction error. For $f(x) = \exp(2 \sin(4\pi x_1) + 0.01 \cos(6\pi x_2))$ the prediction error curve in Fig. 3 (left) looks similar in shape, but converges to a nonzero value less than 10^{-3} (not shown). For observables satisfying $f(x) = \tilde{f}(x_1)$ for some \tilde{f} the prediction error behaves as in the previous example.

The autocorrelations behave as predicted by Proposition 33 (we chose $d = 21$), they seem to maintain the right “frequency” even for $n > d$, however, their “amplitudes” decay faster.

5.4. Von Neumann–Kakutani’s transformation. The von Neumann–Kakutani’s transformation [79], $T : [0, 1] \rightarrow [0, 1]$, is also known under the names dyadic odometer or Coput’s transformation. In [68, section 1.3.1], it is given in the form

$$T\left(1 - \frac{1}{2^n} + y\right) = \frac{1}{2^{n+1}} + y \quad \forall 0 \leq y < \frac{1}{2^{n+1}}, \quad n \in \mathbb{N}_0.$$

After some manipulations, the following form can be obtained:

$$T(x) = x - 1 + \frac{3}{2^{n(x)}}, \quad n(x) := \lfloor -\log_2(1 - x) \rfloor.$$

⁷ The Kronecker factor of a transformation is the maximal factor with pure point spectrum.

⁸ A quasi-eigenfunction of order one is an ordinary eigenfunction of U_T . A function $f \in L^2$, $f \neq 0$, is called a *quasi-eigenfunction* for T of order $n + 1$ if $U_T f = \varphi f$, where φ is a quasi-eigenfunction of order n . One says that T has a *quasi-discrete spectrum* if the span of quasi-eigenfunctions of all possible orders is dense in L^2 .




Fig. 4. The von Neumann–Kakutani’s transformation. L^2 prediction error (left) and autocorrelations (right)




Fig. 5. Spectra for the von Neumann–Kakutani’s transformation. Blue dots indicate the spectrum of U_d and the black line is the complex unit circle. Left: $d = 64$, red circles show “dyadic rationals” $\exp(2\pi i \frac{k}{64})$, $k = 0, \dots, 63$, that are known to be eigenvalues of \mathcal{U} . Right: $d = 250$

This transformation is rank one and has discrete spectrum, with the eigenvalues being the dyadic rationals [68]. The main difference to the torus rotation (which also has discrete spectrum) is that while the spectrum of the latter in Sect. 5.2 is generated by the two irrational factors $\sqrt{2}$ and $\sqrt{3}$, here there are countably many numbers generating the spectrum. We use the observable $f(x) = \exp(\sin(6\pi x))$, and $N = 10^5$ for increased testing accuracy.

Although we observe monotonic decrease in the prediction error in Fig. 4 (left), the error decrease is slowing down for larger delay depths. We attribute this to the complexity of the system’s spectrum, of which the method is picking up more and more as d increases (see also Fig. 5). The autocorrelations also show a more complex picture as compared with the previous examples, thus we depict it for a larger delay depth $d = 51$ in Fig. 4 (right). Despite the complexity, we note that the autocorrelations are very well reproduced on the considered time scale $n = 3d$.

We show the spectra in Fig. 5, which are in this case also reproduced by the least squares filter to a high accuracy.




Fig. 6. Observable for the L63 system, contours shown at $z = 27$. Grey dots indicate the projections of the training points onto the (x_1, x_2) plane

5.5. *The Lorenz '63 model.* We consider the Lorenz '63 (L63) system

$$x'_1 = \sigma(x_2 - x_1), \quad x'_2 = x_1(\rho - x_3) - x_2, \quad x'_3 = x_1x_2 - \beta x_3,$$

with the usual choice of parameters $\sigma = 10, \beta = \frac{8}{3}, \rho = 28$. The discrete time system T that generates the time series is taken as the time- t -flow map ϕ^t of this system for $t = 0.05$. It is realized by the “classical” 4th order Runge–Kutta method with step size $5 \cdot 10^{-4}$. This flow time is considerably smaller than the Lyapunov timescale—the inverse of the largest Lyapunov exponent, used to measure the timescale on which chaotic motion can be observed—of the L63 system, known to be approximately $t_{\text{Lyap}} \approx 1.1$ [80].

Two equilibria of the L63 system are at $x_{\pm} = (\pm\sqrt{\beta(\rho - 1)}, \pm\sqrt{\beta(\rho - 1)}, \rho - 1)$. We set $\delta = \sqrt{\beta(\rho - 1)}$ and consider the two observables

$$f_1(x) = x_1 \quad \text{and} \quad f_2(x) = \exp\left(-\frac{\|x - x_+\|_2^2}{(\delta/3)^2}\right),$$

with f_2 depicted in Fig. 6.

The initial points $x_1^{(i)}$ for the autocorrelation computation are obtained from a trajectory of length $N = 10^6$ generated by the time- t -flow map of L63 with $t = 40h$, using the step size $h = 5.3 \cdot 10^{-4}$ and a random initial point close to $(4, 7, 16)^\top$. We choose $m = 10^5$, as we have validated separately that increasing m does not increase the fidelity of our results.

The Lorenz system is mixing [81] and hence has purely continuous spectrum. Theorem 12 tells us to expect predictiveness only if the trace measure satisfies the Szegő condition. For the observable f_1 the trace measure, and in particular the associated density w , has been computed in [46, Figure 11 (top left)], showing that w is bounded away from zero.⁹ This implies that the Szegő condition is not satisfied and that the prediction error cannot converge to zero. Indeed, we observe that for increasing d the prediction error saturates at a nonzero value, see Fig. 7 (left). The left panel actually

⁹ In [46], the flow time $t = 0.2$ is used, whereas we took $t = 0.05$. We used their method to confirm that both spectral densities of $\phi^{0.05}$, with respect to f_1 and f_2 , are bounded away from zero (not shown).




Fig. 7. L63 system with observable f_1 . Left: L^2 prediction error for prediction time 0.4, using flow times $t \in \{0.05, 0.1, 0.2, 0.4\}$. Right: Autocorrelation for the filter with $d = 21$ using flow time $t = 0.05$

shows a comparison of multistep prediction errors of filters for several delay depths d and flow times $t \in \{0.05, 0.1, 0.2, 0.4\}$ (blue, red, orange, and purple, respectively). The prediction was made for time 0.4, i.e., for 8, 4, 2, 1 steps of the filters associated to the respective flow times. The corresponding errors are shown versus the product dt of delay depth multiplied by the flow time. We observe that the error does not improve after $dt \approx 5$, i.e., the filter gains no advantage from accessing observations from more than a handful Lyapunov times in the past. Additionally, the (in d saturated) errors $\|\mathcal{U}^{0.4} f - \mathcal{U}_d^{0.4} f\|_{L^2}^2$ increase with increasing t . This is in accordance with the intuition that “larger flow times lose more information about the initial condition”. An interesting question, that we do not pursue any further here, would be how the error behaves for $t \rightarrow 0$ as dt is kept constant. We also note that the blue and red curves ($t = 0.05$ and $t = 0.1$, respectively) are not monotonic in d . This is not in conflict with the proof of Proposition 6, because there monotony is shown only for the one-step forecast error, and here we depict multistep errors for these flow times.

Figure 7 (right) shows the autocorrelation comparison as in the previous examples, for $t = 0.05$ and $d = 21$. For steps beyond the delay depth the filter underestimates the true autocorrelation.

The observable f_2 is chosen to pick up a finer-scale dynamics around the equilibrium x_+ . Note that by Footnote 9, we expect it to be unresponsive as well. Indeed, this is confirmed by the error saturating at a nonzero value in Fig. 8 (left). The autocorrelations up to the delay depth $d = 21$ are reproduced exactly, and their exponential decay and qualitative behavior persist beyond that lag; shown in Fig. 8 (right) up to $n = 80 \approx 4d$.

6. Discussion and Outlook

We considered the linear least squares filter (Wiener filter) for prediction of time series and showed Takens-type theorems about the size of the set of dynamics-observables pairs that allow asymptotically exact prediction in the limit of infinite delay depth. Since our setup was ergodic-theoretic and a central role was played by the Koopman operator, we call such a result a “Koopman–Takens” theorem. The main results, to be found in Sects. 2.3 and 3, show among other things that a generic measure-preserving transformation together with a generic square-integrable observable is a predictive pair.




Fig. 8. L63 system with observable f_2 . L^2 prediction error (left) and autocorrelations (right)

Our results showed that systems with discrete spectrum are particularly well suited for the prediction problem (see, e.g., Proposition 14 and Theorem 23); this has also been observed by others, see, e.g. [47]. In general, we conjectured that discrete spectrum together with ergodicity (or perhaps even merely the rank one property) implies that a prevalent set of cyclic observables exist, see Conjecture 25. Interestingly, however, mixing systems (that are known to have purely continuous spectrum) can be predictive as well, see the discussion after Proposition 14 and Remark 29. Further, Proposition 38 indicated that the filter with finite delay depth alone does not allow us to draw conclusions about predictiveness. This can be done by validating the mean squared forecast error on a test set. In numerical experiments not reported here we observed that the spectrum of U_d can suggest information about the type of spectrum of the system; see also the discussion on the spectral convergence of Hankel DMD in [46]. This needs further rigorous elaboration.

On a quantitative level our numerical experiments suggest, on the one hand, that the forecast error decays with increasing delay depth faster for smooth observables than for nonsmooth (e.g., discontinuous) ones. On the other hand, less complicated spectrum has a similar effect: While the torus rotation in Sect. 5.2 has a spectrum generated by two scalars (the entries of the translation vector), the von Neumann–Kakutani’s transformation has a spectrum generated by all 2^{-n} , $n \in \mathbb{N}$.

Natural next steps would be to consider vector-valued observations, investigate the effect of noise in the observations, and to derive quantitative error bounds for finite delay depths (in contrast to the asymptotic results given herein). Further, even closer to the spirit of Takens’ theorem would be to have results on some measure-theoretic reconstruction of the state space in terms of the infinite observation sequences.

Acknowledgements We thank Friedrich Philipp for discussions on weak predictiveness and for a part of the proof of Theorem 12. We also thank the referees of the first version of this paper for their careful reading and insightful comments. Péter Koltai has been partially supported by Deutsche Forschungsgemeinschaft (DFG) through grant CRC 1114 “Scaling Cascades in Complex Systems”, Project Number 235221301, Project A08 “Characterization and prediction of quasi-stationary atmospheric states”. The research of Philipp Kunde is part of the project No. 2021/43/P/ST1/02885 co-funded by the National Science Centre and the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no. 945339. He also thanks Freie Universität Berlin for its hospitality during three research visits related to this paper.

Funding Open Access funding enabled and organized by Projekt DEAL. Péter Koltai has been partially supported by Deutsche Forschungsgemeinschaft (DFG) through Grant CRC 1114 “Scaling Cascades in Complex

Systems”, Project Number 235221301, Project A08 “Characterization and prediction of quasi-stationary atmospheric states”. The research of Philipp Kunde is part of the project No. 2021/43/P/ST1/02885 co-funded by the National Science Centre and the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant agreement No. 945339.

Data Availability Not applicable/All data generated or analysed during this study are included in this published article.

Declarations

Conflict of interest The authors have no Conflict of interest to declare that are relevant to the content of this article.

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 licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence 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 licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

A. Proof of Lemma 18

Proof of Lemma 18. The statements (1) \Leftrightarrow (3) \Leftrightarrow (4) can also be found in [65, p. 62]. (1) \Leftrightarrow (2). Let $A \in \mathfrak{B}$. Since T preserves μ , we have $\mu(T_n A \Delta T A) = \mu(T^{-1} T_n A \Delta A)$. Let $B := T_n A$. We obtain $\mu(T_n A \Delta T A) = \mu(T^{-1} B \Delta T_n^{-1} B)$. As T, T_n are measure-preserving essential bijections, $B \in \mathfrak{B}$, and in fact $T\mathfrak{B} = \mathfrak{B} = T_n\mathfrak{B}$ for every n . We thus obtain the equivalence of $T_n \xrightarrow{w} T$ and $T_n^{-1} \xrightarrow{w} T^{-1}$ as $n \rightarrow \infty$.

(1) \Rightarrow (5). We will show the claim for $i \geq 0$, while $i < 0$ then follows with (1) \Leftrightarrow (2). There is nothing to show for $i = 0$, let us thus consider $i = 1$ first. For arbitrary $\varepsilon > 0$ let $g_\varepsilon = \sum_{j=1}^{N_\varepsilon} \alpha_j \mathbb{1}_{A_j}$, $A_j \in \mathfrak{B}$, be a simple function with $\|g - g_\varepsilon\|_{L^2_\mu} < \varepsilon$. We note that

$$\|g_\varepsilon \circ T - g_\varepsilon \circ T_n\|_{L^2_\mu} \leq \sum_{j=1}^{N_\varepsilon} |\alpha_j| \|\mathbb{1}_{A_j} \circ T - \mathbb{1}_{A_j} \circ T_n\|_{L^2_\mu} = \sum_{j=1}^{N_\varepsilon} |\alpha_j| \mu(T_n A_j \Delta T A_j)^{1/2}.$$

By assumption, we thus have $\|g_\varepsilon \circ T - g_\varepsilon \circ T_n\|_{L^2_\mu} \rightarrow 0$ as $n \rightarrow \infty$. Using this and the measure-preserving property of T, T_n , we obtain for sufficiently large n that

$$\begin{aligned} \|g \circ T - g \circ T_n\|_{L^2_\mu} &\leq \|g \circ T - g_\varepsilon \circ T\|_{L^2_\mu} + \|g_\varepsilon \circ T - g_\varepsilon \circ T_n\|_{L^2_\mu} \\ &\quad + \|g_\varepsilon \circ T_n - g \circ T_n\|_{L^2_\mu} \\ &< \varepsilon + \varepsilon + \varepsilon. \end{aligned}$$

Thus, $g \circ T_n \rightarrow g \circ T$ in L^2_μ as $n \rightarrow \infty$. The general statement follows by simple induction, as

$$\begin{aligned} & \|g \circ T^{i+1} - g \circ T_n^{i+1}\|_{L_\mu^2} \leq \|g \circ T^i \circ T - g \circ T^i \circ T_n\|_{L_\mu^2} \\ & + \|g \circ T^i \circ T_n - g \circ T_n^i \circ T_n\|_{L_\mu^2}, \end{aligned}$$

and T_n is measure-preserving.

(3) \Rightarrow (3). Trivial by setting $i = 1$.

(3) \Rightarrow (1). Trivial by taking g as characteristic functions.

(3) \Leftrightarrow (4). The first implication is trivial. The second is a straightforward computation, using that $\mathcal{U}_T, \mathcal{U}_{T_n}$ are unitary operators.

References

- Hamilton, F., Berry, T., Sauer, T.: Predicting chaotic time series with a partial model. *Phys. Rev. E* **92**, 010902 (2015). <https://doi.org/10.1103/PhysRevE.92.010902>
- Hsieh, C.-H., Glaser, S., Lucas, A., Sugihara, G.: Distinguishing random environmental fluctuations from ecological catastrophes for the North Pacific Ocean. *Nature* **435**, 336–340 (2005)
- Munch, S., Giron-Nava, A., Sugihara, G.: Nonlinear dynamics and noise in fisheries recruitment: a global meta-analysis. *Fish Fish.* **19**, 964–973 (2018). <https://doi.org/10.1111/faf.12304>
- Sugihara, G., Grenfell, B., May, R., Chesson, P., Platt, H.M., Williamson, M.: Distinguishing error from chaos in ecological time series. *Philos. Trans. Biol. Sci.* **330**(1257), 235–251 (1990)
- Sugihara, G., May, R.: Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series. *Nature* **344**, 734–741 (1990)
- Barański, K., Gutman, Y., Śpiewak, A.: A probabilistic Takens theorem. *Nonlinearity* **33**(9), 4940–4966 (2020)
- Barański, K., Gutman, Y., Śpiewak, A.: On the Shroer–Sauer–Ott–Yorke predictability conjecture for time-delay embeddings. *Commun. Math. Phys.* **391**, 609–641 (2022)
- Barański, K., Gutman, Y., Śpiewak, A.: Prediction of dynamical systems from time-delayed measurements with self-intersections. (2022) arXiv preprint [arXiv:2212.13509](https://arxiv.org/abs/2212.13509)
- Farmer, J.D., Sidorowich, J.J.: Predicting chaotic time series. *Phys. Rev. Lett.* **59**, 845–848 (1987). <https://doi.org/10.1103/PhysRevLett.59.845>
- Kostelich, E., Yorke, J.: Noise reduction: finding the simplest dynamical system consistent with the data. *Phys. D* **41**(2), 183–196 (1990)
- Sauer, T., Yorke, J.A., Casdagli, M.: Embedology. *J. Stat. Phys.* **65**(3), 579–616 (1991)
- Packard, N.H., Crutchfield, J.P., Farmer, J.D., Shaw, R.S.: Geometry from a time series. *Phys. Rev. Lett.* **45**(9), 712 (1980)
- Takens, F.: Detecting strange attractors in turbulence. In: *Dynamical Systems and Turbulence*, Warwick 1980, Volume 898 of Lecture Notes in Math., pp. 366–381. Springer, Berlin–New York (1981)
- Voss, H.U.: Synchronization of reconstructed dynamical systems. *Chaos* **13**, 327–334 (2003). <https://doi.org/10.1063/1.1514059>
- Wang, Z., Guet, C.: Self-consistent learning of neural dynamical systems from noisy time series. *IEEE Trans. Emerg. Top. Comput. Intell.* **6**(5), 1103–1112 (2022). <https://doi.org/10.1109/TETCI.2022.3146332>
- Percival, D.B., Walden, A.T.: *Spectral Analysis for Physical Applications*. Cambridge University Press, Cambridge (1993)
- Wiener, N.: *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*, vol. 113. The MIT Press, Cambridge, Massachusetts (1949)
- Kolmogorov, A.: Stationary sequences in Hilbert space. *Bull. Math. Univ. Moscou* **2**(6), 1–40 (1941)
- Kolmogorov, A.: Interpolation and extrapolation of stationary random sequences. *Izvestiya Rossiiskoi Akademii Nauk. Seriya Matematicheskaya* **5**, 3 (1941)
- Takens, F.: The reconstruction theorem for endomorphisms. *Bull. Braz. Math. Soc. (N.S.)* **33**(2), 231–262 (2002)
- Gutman, Y.: Takens’ embedding theorem with a continuous observable. In: *Ergodic Theory: Advances in Dynamical Systems*, pp. 134–141. De Gruyter, Berlin–Boston (2016)
- Robinson, J.C.: A topological delay embedding theorem for infinite-dimensional dynamical systems. *Nonlinearity* **18**(5), 2135–2143 (2005)
- Robinson, J.C.: *Dimensions, Embeddings, and Attractors*. Cambridge Tracts in Mathematics. Cambridge University Press, Cambridge (2010). <https://doi.org/10.1017/CBO9780511933912>
- Broomhead, D.S., King, G.P.: Extracting qualitative dynamics from experimental data. *Phys. D* **20**(2–3), 217–236 (1986)

25. Dellnitz, M., Hessel-Von Molo, M., Ziessler, A.: On the computation of attractors for delay differential equations. *J. Comput. Dyn.* **3**(1), 93–112 (2016)
26. Brunton, S.L., Brunton, B.W., Proctor, J.L., Kaiser, E., Kutz, J.N.: Chaos as an intermittently forced linear system. *Nat. Commun.* **8**(1), 19 (2017)
27. Kamb, M., Kaiser, E., Brunton, S.L., Kutz, J.N.: Time-delay observables for Koopman: theory and applications. *SIAM J. Appl. Dyn. Syst.* **19**(2), 886–917 (2020)
28. Gottwald, G.A., Reich, S.: Combining machine learning and data assimilation to forecast dynamical systems from noisy partial observations. *Chaos Interdiscip. J. Nonlinear Sci.* **31**(10), 101103 (2021)
29. Lapedes, A., Farber, R.: Nonlinear signal processing using neural networks: Prediction and system modelling. Technical report (1987)
30. Young, C.D., Graham, M.D.: Deep learning delay coordinate dynamics for chaotic attractors from partial observable data. *Phys. Rev. E* **107**(3), 034215 (2023)
31. Bakarji, J., Champion, K., Kutz, N.J., Brunton, S.L.: Discovering governing equations from partial measurements with deep delay autoencoders. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **479**(2276), 20230422 (2023). <https://doi.org/10.1098/rspa.2023.0422>
32. Gottwald, G.A., Reich, S.: Supervised learning from noisy observations: combining machine-learning techniques with data assimilation. *Phys. D* **423**, 132911 (2021)
33. Wulkow, N., Koltai, P., Sunkara, V., Schütte, C.: Data-driven modelling of nonlinear dynamics by barycentric coordinates and memory. (2021) arXiv preprint [arXiv:2112.06742](https://arxiv.org/abs/2112.06742)
34. Zwanzig, R.: Statistical mechanics of irreversibility. In: Brittin, W. (ed.) *Lectures in Theoretical Physics*, vol. 3. Wiley-Interscience, New York, NY, USA (1961)
35. Mori, H.: Transport, collective motion, and Brownian motion. *Prog. Theoret. Phys.* **33**(3), 423–455 (1965). <https://doi.org/10.1143/PTP.33.423>
36. Gouasmi, A., Parish, E.J., Duraisamy, K.: A priori estimation of memory effects in reduced-order models of nonlinear systems using the Mori-Zwanzig formalism. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **473**(2205), 20170385 (2017). <https://doi.org/10.1098/rspa.2017.0385>
37. Lin, K.K., Lu, F.: Data-driven model reduction, Wiener projections, and the Koopman-Mori-Zwanzig formalism. *J. Comput. Phys.* **424**, 109864 (2021). <https://doi.org/10.1016/j.jcp.2020.109864>
38. Gilani, F., Giannakis, D., Harlim, J.: Kernel-based prediction of non-Markovian time series. *Phys. D Nonlinear Phenom.* **418**, 132829 (2021). <https://doi.org/10.1016/j.physd.2020.132829>
39. Lin, Y.T., Tian, Y., Perez, D., Livescu, D.: Regression-based projection for learning Mori–Zwanzig operators. *SIAM J. Appl. Dyn. Syst.* **22**(4), 2890–2926 (2023). <https://doi.org/10.1137/22M1506146>
40. Koopman, B.O.: Hamiltonian systems and transformation in Hilbert space. *Proc. Nat. Acad. Sci.* **17**(5), 315–318 (1931)
41. Koopman, B.O., Neumann, J.: Dynamical systems of continuous spectra. *Proc. Nat. Acad. Sci.* **18**(3), 255–263 (1932)
42. Arbabi, H., Mezić, I.: Ergodic theory, dynamic mode decomposition, and computation of spectral properties of the Koopman operator. *SIAM J. Appl. Dyn. Syst.* **16**(4), 2096–2126 (2017)
43. Schmid, P.J., Sesterhenn, J.: Dynamic mode decomposition of numerical and experimental data. *Bull. Am. Phys. Soc.* **656**, 5–28 (2008)
44. Rowley, C.W., Mezić, I., Bagheri, S., Schlatter, P., Henningson, D.S.: Spectral analysis of nonlinear flows. *J. Fluid Mech.* **641**, 115–127 (2009)
45. Williams, M.O., Kevrekidis, I.G., Rowley, C.W.: A data-driven approximation of the Koopman operator: extending dynamic mode decomposition. *J. Nonlinear Sci.* **25**, 1307–1346 (2015)
46. Korda, M., Putinar, M., Mezić, I.: Data-driven spectral analysis of the Koopman operator. *Appl. Comput. Harm. Anal.* **48**(2), 599–629 (2020)
47. Das, S., Giannakis, D.: Delay-coordinate maps and the spectra of Koopman operators. *J. Stat. Phys.* **175**(6), 1107–1145 (2019)
48. Giannakis, D.: Delay-coordinate maps, coherence, and approximate spectra of evolution operators. *Res. Math. Sci.* **8**(1), 8 (2021)
49. Colbrook, M.J., Townsend, A.: Rigorous data-driven computation of spectral properties of Koopman operators for dynamical systems. *Commun. Pure Appl. Math.* **77**(1), 221–283 (2024)
50. Colbrook, M.J.: The mpEDMD algorithm for data-driven computations of measure-preserving dynamical systems. *SIAM J. Numer. Anal.* **61**(3), 1585–1608 (2023)
51. Valva, C., Giannakis, D.: Consistent spectral approximation of Koopman operators using resolvent compactification. (2023) arXiv preprint [arXiv:2309.00732](https://arxiv.org/abs/2309.00732)
52. Huke, J.P.: Embedding nonlinear dynamical systems: A guide to takens’ theorem. (2006). <https://api.semanticscholar.org/CorpusID:55183186>
53. Walters, P.: *An Introduction to Ergodic Theory*, vol. 79. Springer, New York (2000)
54. Kreyszig, E.: *Introductory Functional Analysis with Applications*, vol. 17. John Wiley & Sons, New York (1991)

55. Bröcker, J.: Reliability, sufficiency, and the decomposition of proper scores. *Q. J. R. Meteorol. Soc.* **135**, 1512–1519 (2009). <https://doi.org/10.1002/qj.456>
56. Cai, D., Kleeman, R., Majda, A.: A mathematical framework for quantifying predictability through relative entropy. *Methods Appl. Anal.* **9**(3), 425–444 (2002)
57. Gneiting, T., Raftery, A.: Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* **102**(477), 359–378 (2007). <https://doi.org/10.1198/016214506000001437>
58. Duncan, R.: Some pointwise convergence results in $L^p(\mu)$, $1 < p < \infty$. *Can. Math. Bull.* **20**, 3 (1977)
59. Taylor, M.: *Partial Differential Equations II: Qualitative Studies of Linear Equations*, vol. 116. Springer, New York (2013)
60. Eisner, T., Farkas, B., Haase, M., Nagel, R.: *Operator Theoretic Aspects of Ergodic Theory*, vol. 272. Springer, Cham (2015)
61. Simon, B.: *Szegő's Theorem and Its Descendants*. Princeton University Press, Princeton (2010)
62. Philipp, F.: Bessel orbits of normal operators. *J. Math. Anal. Appl.* **448**(2), 767–785 (2017)
63. Rudin, W.: *Real and Complex Analysis*, 3rd edn. McGraw-Hill Book Co., New York City (1987)
64. Bogachev, V.I.: *Measure Theory*. Vol. I, II, pp. 500–575. Springer, Berlin (2007). <https://doi.org/10.1007/978-3-540-34514-5>
65. Halmos, P.R.: *Lectures on Ergodic Theory*. Courier Dover Publications, New York (2017)
66. Engel, K.-J., Nagel, R., Brendle, S.: *One-parameter Semigroups for Linear Evolution Equations*, vol. 194. Springer, New York (2000)
67. Nadkarni, M.: *Spectral Theory of Dynamical Systems*, 2nd edn. Springer, Singapore (2020). <https://doi.org/10.1007/978-981-15-6225-9>
68. Ferenczi, S.: Systems of finite rank. *Colloq. Math.* **73**, 35–65 (1997)
69. Del Junco, A.: Transformations with discrete spectrum are stacking transformations. *Can. J. Math.* **28**(4), 836–839 (1976). <https://doi.org/10.4153/CJM-1976-080-3>
70. Katok, A., Thouvenot, J.-P.: Spectral properties and combinatorial constructions in ergodic theory. *Handb. Dyn. Syst.* **1B**, 649–743 (2006)
71. Katok, A., Stepin, A.: Approximations in ergodic theory. *Uspehi Mat. Nauk*, 81–106 (1967)
72. Bozgan, F., Sanchez, A., Silva, C.E., Stevens, D., Wang, J.: Subsequence bounded rational ergodicity of rank-one transformations. *Dyn. Syst.* **30**, 70–84 (2015). <https://doi.org/10.1080/14689367.2014.970518>
73. Alpern, S., Prasad, V.S.: *Typical Dynamics of Volume Preserving Homeomorphisms*, Vol. 139. Cambridge University Press, Cambridge (2000)
74. Melbourne, I., Gottwald, G.A.: Power spectra for deterministic chaotic dynamical systems. *Nonlinearity* **21**(1), 179 (2007)
75. Colbrook, M.J., Ayton, L.J., Szőke, M.: Residual dynamic mode decomposition: robust and verified Koopmanism. *J. Fluid Mech.* **955**, 21 (2023)
76. Trefethen, L.N., Embree, M.: *Spectra and Pseudospectra: The Behavior of Non-normal Matrices and Operators*. Princeton University Press, Berlin (2005)
77. Mezić, I.: On numerical approximations of the Koopman operator. *Mathematics* **10**(7), 1180 (2022)
78. Korda, M., Mezić, I.: On convergence of extended dynamic mode decomposition to the Koopman operator. *J. Nonlinear Sci.* **28**, 687–710 (2018)
79. Neumann, J.: Zur Operatorenmethode in der klassischen Mechanik. *Ann. Math.* 587–642 (1932)
80. Geurts, B.J., Holm, D.D., Luesink, E.: Lyapunov exponents of two stochastic Lorenz 63 systems. *J. Stat. Phys.* **179**, 1343–1365 (2020)
81. Luzzatto, S., Melbourne, I., Paccaut, F.: The Lorenz attractor is mixing. *Commun. Math. Phys.* **260**, 393–401 (2005)