Freie Universität Berlin

Master Thesis at the Institute of Computer Science

Human-Centered Computing (HCC)

Ideator Types in Electronic Brainstorming

Kim Kern

Supervisor: Maximilian Mackeprang

First Reviewer: Prof. Dr. Claudia Müller-Birn

Second Reviewer: Prof. Dr. Christoph Benzmüller

Berlin, 2019-09-13

Eidesstattliche Erklärung

Ich versichere hiermit an Eides Statt, dass diese Arbeit von niemand anderem als meiner Person verfasst worden ist. Alle verwendeten Hilfsmittel wie Berichte, Bücher, Internetseiten oder ähnliches sind im Literaturverzeichnis angegeben, Zitate aus fremden Arbeiten sind als solche kenntlich gemacht. Die Arbeit wurde bisher in gleicher oder ähnlicher Form keiner anderen Prüfungskommission vorgelegt und auch nicht veröffentlicht.

Berlin, den 13. September 2019

Kim Kern

Acknowledgments

First of all, I want to thank Maximilian Mackeprang, Claudia Müller-Birn and the whole HCC research group for their constructive and honest feedback on my presentations, proposal and study design as well as the positive and warm atmosphere in the office.

Special thanks are due to Maximilian, especially, the long coding, post-it and whiteboard sessions with him were not only very helpful and productive but actually enjoyable. I really appreciate the many things I learned from him during this year.

I am very grateful for the statistical advice and emotional support I received from Daniel.

And last but not least, I want to thank my friends, especially Jeroen, for proofreading my thesis.

Abstract

Large scale ideation continues to gain importance in today's economy, driving research and development. One of its best-known techniques is electronic brainstorming (EBS), which has proven to be successful at producing a large number of ideas but struggles with a high quantity of low-quality ideas. This weakness motivated the question: How can the quality of these ideas be improved?

While much research on the improvement of idea quality has been done regarding what kind of inspiration is given during ideation and in what way, little has been looked into individual differences of ideators who may have different needs in terms of inspiration and therefore show different behavior when exposed to inspiration. This thesis examines whether such types of ideators with individual differences can be identified.

As individual differences may have various dimensions, first, an in-situ exploratory study was conducted to identify individual differences in the context of inspirations, that can be tested for its impact on the ideation outcome in a subsequent, quantitative study.

The exploratory study induced the idea of the existence of the ideator types *inspiration seeker* (benefiting from inspiration) and *inspiration avoider* (feeling distracted by inspiration). The analysis of data from previous EBS studies showed that this idea applies not only to classical group ideation but to a large-scale ideation setting as well.

In order to understand the impact of the identified ideator types, a quantitative study was conducted. It aimed at replicating a recent study on the influence of inspirational stimuli (Siangliulue et al., 2015), while additionally examining different effects of these stimuli on *inspiration seekers* and *inspiration avoiders*.

The analysis of the study showed that the ideator type did not seem to have an impact on the number of submitted ideas or their value. However, *avoiders* produced ideas with a higher maximum novelty per session than *seekers* across all inspiration conditions with the greatest difference between the types when no inspiration at all was provided.

The results show that individual differences regarding inspirational stimuli exist and do impact the fluency and quality of ideas. Paying attention to these differences is a promising approach to improve the quality of the ideas produced in electronic innovation systems. This classification could potentially be used to create personalized inspiration systems catering to the needs of different ideator types.

Contents

List of Figures

List of Tables

1. Introduction

In a globalized world, the pressure for innovation is rapidly increasing for international companies [AI02]. However, generating innovative thinking remains difficult in complex R&D process frameworks. Hence, companies are continuously looking for "leaner and more effective (...) innovation techniques" [GR99]. One of these techniques is the well-known *brainstorming* method introduced by Alex F. Osborn in 1948, in which groups of people freely propose new ideas for a specific problem or domain while the judgment of these ideas is deferred [Osb48].

Although this is a widely used approach, there are several problems connected to it, especially for groups with growing sizes: Most importantly, since only one person can talk at the same time while leaving the other participants idle, unnecessary *resource and production blocking* can be observed [DS87, Osb53, Nij00]. Additionally, several studies have proven the adverse effect of *social loafing*, meaning that people work less effectively in group tasks when their efforts cannot be directly attributed [HJ85, PN03]. *Evaluation apprehension* can act as a further hindrance for productivity as people fear to be judged by their peers [CJV90]. These effects can become so potent that people working individually outperform groups of the same size in terms of quality and quantity [DS87, NSL02, PBGH99]. Especially for big groups, the adverse effects predominate since the procedure does not scale well.

As a solution, *electronic brainstorming (EBS)* has been used to do large-scale ideation with the help of computers and the internet. This allows for participants to work simultaneously and therefore reduces production blocking. In a business context, it also facilitates access to the external view from complete outsiders or customers, which offers an "enormous potential for generating innovation" [BBLK10, MKGP06].

For instance, IBM has conducted several ideation sessions, which they called innovation jams [BW08]. Although it seemed easy for them to generate a high number of ideas, they struggled with the quality of the ideas. Edward Bevan, IBM vice president for technology/innovation, said:

> "Idea generation is in some ways the 'easy' part (...) of innovation, whereas advancing, refining and building support for those ideas is the really tough part." [BW08]

Further research could confirm this finding and described most ideas as "simple, mundane and repetitive" [SCDG16] or even found only 10-30% of the ideas from ideation engagements to be considered as high-quality [BBLK10, KG15].

In response to this problem, the question arose at IBM, "whether a systematic process could be developed" to increase the idea quality and reduce post-processing efforts.

One approach in this direction is described by Siangliulue et al., who compared different ways of providing inspirations in terms of its impact on the quantity and quality of the produced ideas [SCGD15]. It was shown that the quality of ideas improved with inspirations and hence poses a good approach to tackle the problems mentioned above. However, the study did not consider whether different groups of ideators have divergent reactions, based on their individual differences.

The field of individual differences in ideation has yet to be explored in-depth, although it has recently gained interest. Girotto et al. stated that in ideation, a "one size fits all approach does not work." [GWB17] Another recent study identified opposing effects on different groups of participants by exposing proposed ideas to immediate feedback [GBD17]. Participants with high motivation and ability considered feedback to be distracting, whereas others considered it beneficial for their creativity.

Following this path of a personalized ideation experience, Girotto et al. raised the question:

> "Could individual differences between ideators affect how effective an inspiration is?" [GWB19]

As both individual differences of the participants and the utilization of inspirations have shown to be fruitful approaches, the following central research question is framed:

> Are there individual differences in how ideators react to inspirations?

As this question is too broad to be answered in this thesis, an exploratory approach is chosen to identify a specific individual difference which can be tested quantitatively in an electronic brainstorming setting.

Thesis Structure

After the general introduction and motivation in chapter 1, the thesis will begin by laying out the broad foundations for the field of large-scale ideation, while also listing relevant state-of-the-art research concerning individual differences in chapter 2. Chapter 3 presents the study design, results and conclusion for the exploratory study that was conducted to gain further insights into individual differences of ideators regarding inspirations. In chapter 4, the findings of the exploratory study are cast in a concrete definition, and with data from previous studies, it is tested whether this definition is also useful in the context of electronic brainstorming. The following quantitative MTurk study about the impact of inspirations on the two identified ideator types seekers and avoiders is presented with its study design, statistical analysis and a discussion of the results in chapter 5. Some limitations of the study design and outcome as well as potential future research are shown in chapter 6. The thesis ends with an overall conclusion of all findings in chapter 7.

This chapter introduces the historical background of brainstorming and large-scale ideation and exposes some general approaches to pave the way for the state-of-the-art related work in the next chapter.

2. Related Work

After introducing some theoretical background for ideation and creativity, this chapter outlines the development from classical brainstorming over electronic brainstorming to state-of-the-art research on large-scale ideation. The last section of this chapter focuses on recent studies on individual differences in the context of ideation.

2.1. Theoretical Background

One aspect of research in brainstorming lies in the theoretical model for the cognitive processes during brainstorming. A prominent model that is used by multiple works in the field of ideation is the Search for Ideas in Associative Memory (SIAM) theory [NS06, NSL02]. The SIAM model explains how idea generation is conducted. It distinguishes between two kinds of memory: a high-capacity long-term memory (LTM) storing highly linked images and a short-term working memory (STM) which is used for conscious, cognitive tasks but only holds one mental image at a time due to its limited capacity.

In order to use these two kinds of memory, the process of ideation can be partitioned into two loops, the knowledge activation stage and the idea production stage, see figure 5.13: In the knowledge activation stage, an ideator searches for mental images in the LTM with a search cue comprised of the ideation challenge. When a match is found, the image is loaded into STM. In the subsequent idea production stage, ideas are generated based on the concept loaded in the STM. After several failures to generate new ideas for the loaded image, again the ideator returns to the first stage, but now the search cue additionally includes all previously seen concepts and ideas.

According to SIAM, showing ideas during an ideation session can have positive and negative effects. On the one hand, it can boost the ideators' creativity as an inspiration [DFS+11]; on the other hand, it can lead to fixation and hinder the generation of novel ideas [PS94].

The spreading activation theory expands this cognitional model by illustrating how different concepts (resp. images) are interconnected in the brain [CL75]: Closely related concepts are stored close to each other. Therefore, apparent links leading to similar concepts are followed first. Only when these are exhausted, more far-fetched paths are followed, leading to unexpected and novel connections. This theory poses a possible explanation for the serial-order effect [BS12], which states that the creativity of ideas increases over the course of an ideation session.

These theoretical models can be applied in various forms of ideation. Some are outlined in the following sections.

2.2. Classical Brainstorming

The origins of brainstorming date back more than 50 years, when Alex F. Osborn first presented it in 1948 as a group ideation technique with his four brainstorming rules [Osb48]:

1. Judicial judgment is ruled out.
2. "Free-wheeling" is welcomed.
3. Quantity is wanted.
4. Combination and improvement are sought.

Its goal is, therefore, to produce a high number of diverse ideas.

There are various different forms of brainstorming: In brainwriting, for example, the ideas are not discussed in a group setting but instead written down with pen and paper. Schlicksupp described *Brainwriting Pool* as a brainwriting variant, in which the participants sit around a table [Sch75]. They write an idea down, each idea on its own note, and then put the idea in the center of the table, the idea pool. When they want inspiration, they can, at any time, take an idea from the pool and read it. This method is supposed to reduce the stress level of the ideators, since social pressure is mitigated, and there is no fixed clocking and dependence on others like in the 6-3-5 method.

Although these approaches seemed promising and are still largely used even today, there is evidence depicting it to be significantly less productive in terms of quality and quantity of the generated ideas compared to nominal brainstorming, where participants ideate individually without any interaction with other participants [DS87, MJS91]. This difference in the performance is mainly attributed to a number of process losses in the brainstorming method [PBGH99]:

Production and resource blocking Since in a verbal brainstorming session only one person can talk at a time, other participants stay idle and cannot use the time to ideate themselves [DS87, Osb53, Nij00].

Cognitive interference Listening to the ideas of others might interfere with the ideator's train of thought and can, therefore, appear as rather distracting than inspiring [DS87, NSL02, NS06].

Evaluation apprehension The fear of being judged by peers or superiors can be blocking and hinder productivity [CJV90, DS87, MJS91].

Negative productivity matching Since the productivity of other ideators is observed by the participant, they might fall back to a common baseline level of productivity [GD87].

Pressure for cognitive conformity Participants might fear to propose ideas that do not conform to society's moral standards and norms [HK74, MKM84].

Personalization of issues Ideators consider personal matters instead of substantial and objective criteria [MKM84].

Social influences Individual participants might take a leading position that others subordinate to, although the individuals are not qualified to take this position [NDV+91, MKM84].

Free riding Ideators willingly withdraw from participation because they feel that their efforts are not needed or because they rely on the contributions of others due to their laziness (social loafing) [HJ85, PN03].

These negative effects magnify with growing group sizes, making classical brainstorming especially unfeasible for large-scale ideation.

2.3. Large-Scale Ideation

In the early 1980s, a new form of a digital communication platform, called group support systems (GSS), emerged. Among other things, they were used to provide a solution to the problems of classical brainstorming, and so, electronic brainstorming (EBS) was proposed [DW03]. Since many people can ideate simultaneously on computers without mutually blocking their resources, this approach was deemed to be more productive than verbal brainstorming and nominal brainstorming. Although this effect could be confirmed in several studies [GDC+92, DV93], other research challenged this view and suggested that EBS did not hold up to its expectations and could even be outperformed by nominal brainstorming [PBGH99].

IBM gained practical experience in this field by experimenting a lot with ideation sessions, which they internally called “innovation jams” [BW08]. It was shown that although it was easy to produce a high amount of ideas even with a large group of participants, the quality was poor (“a lot of garbage”) and hence the post-processing efforts (sorting, grouping, ranking) were very tedious. This is in accordance with the findings of other research, where most generated ideas are described as “simple, mundane and repetitive” and the micro-tasks needed for the post-processing as repetitive [SCDG16].

Hence, one of the most challenging problems of large-scale ideation systems is the production of obvious and unfavorable solutions, especially with its growing number of participants [KC15]. Although presenting abstracted solution paths in contrast to raw, previous ideas to the ideators has been described as a promising approach, its positive effects could not be shown [CDD16].

As an alternative, much research can be found about creativity-enhancing interventions, in particular, inspiration systems, that show previous ideas for the challenge as inspiration.

In these settings, showing highly diverse ideas can increase the level of creativity [KMMB18, MLH96]. To the contrary, showing a homogeneous set of ideas potentially increases depth, but might as well lead to conformity and fixation [GWB17, NSL02]. However, the categorization of ideas poses to be a hard problem. As the generated ideas are mostly very short snippets, automatic categorization, e.g., through topic models, perform poorly [CGW+09, BW08, SCDG16]. The Innovonto platform, however, proposed a promising way of partly automated classification utilizing semantic web technologies [KMMB18].

One recent and influential study was not so much concerned with which inspiration to show but rather when to show it: In Providing Timely Examples Improves the Quantity and Quality of Generated Ideas [SCGD15], Siangliulue et al. examined different ways of delivering inspirations. They compared the impact of providing inspirations when a user explicitly requests them as well as when a user is idle for 30 seconds against two baseline conditions: providing no inspirations at all and automatically providing inspirations at a fixed interval. In the experiment, 97 MTurk crowd workers ideated for 15 minutes. Then, the produced ideas were counted (*fluency*), and the mean of the idea quality (*novelty* and *value*) per session was assessed via ratings of other crowd workers. They found that ideators who requested ideas themselves had the highest novelty ratings for their ideation sessions, whereas participants, who received inspirations automatically when they appeared idle, generated the highest number of ideas. No significant differences were found for the novelty ratings. Since inspirations at a fixed interval generated the fewest ideas, they concluded that inspirations given at the wrong time hinder ideation while giving inspirations when a participant is prepared for it is beneficial for their creativity. As a next step, the authors highlight the possible "benefits of personalized examples" specific for different kinds of ideators. In the following section, approaches utilizing individual differences of ideators are presented.

2.4. Personalized Ideation

As mentioned before, taking the individual difference into consideration when designing a creativity system is an auspicious approach.

In that sense, Chirumbolo et al. found that the ideators with a high dispositional need for closure were less creative [CLM+04]. However, the effects of these individual differences were only investigated in a classical, in-situ brainstorming context.

Also, in a classical group interventions setting, Choo et al. classified professional designers according to the Myers-Briggs Type Indicator (MBTI) with a preceding survey [CLC+15]. They showed that the associated personality type impacted the ideation outcome in terms of quality and quantity.

In electronic brainstorming sessions, albeit still on-site, Garfield et al. used the MBTI to categorize ideators by individual differences [GTDS01]. They found a divergent tendency to produce paradigm-modifying ideas for ideators due to their different personality types.

In a similar setting, De Jonge et al. differentiated ideators in terms of their psychological needs for structure and autonomy [DJRVY18]. Participants that were either high in their need for autonomy or low in their need for structure showed higher levels of self-perceived creativity when exposed to novel input.

Also working with a user model, Gamper et al. investigated whether participants preferred getting feedback immediately during an in-situ EBS session or rather after it [GBD17]. Although classical feedback counts “deferred feedback” as one of its rules, the study found that participants with high motivation and ability felt distracted by immediate feedback while others benefited from it.

In an online study, Olteteanu et al. identified the tendency to reorient objects and also the need for closure as individual difference with an impact on the ideas [OS18]. In an alternative uses test, differences in fluency and flexibility of the ideation outcome were reported.

As the authors of one of the first studies on individual differences in large-scale ideation, Girotto et al. describe it as negligence of current creativity support systems that the “ideators’ individualities” are not considered [GWB19]. With CrowdMuse, they propose a personalized inspiration system that takes previously submitted ideas of participants into account. In two online studies, they showed that the width of the ideas increased with their adaptive ideation solution. For further research, they propose adapting to different “cognitive strategies of individual ideators”. However, the study only considered the content of the inspirational stimuli, not how they are presented.

Individual differences, especially with regard to inspirations, have been identified as a very recent and bright field of research. The next chapter presents an exploratory study to establish possible individual differences in this context.

3. Exploratory Study

As shown in the last chapter, individual differences are an active topic in the field of large-scale ideation. However, it is not clear what concrete individual differences could be relevant in the context of inspirational stimuli. In order to gain further insights into the mechanisms of inspiration, an exploratory brainstorming study was conducted. This chapter describes how by analyzing the ideators' behavior and survey results of three brainstorming sessions, different types of ideators in ideation processes could be identified.

3.1. Methodology

The goal of the exploratory study was a setup that is comparable to an electronic setting while at the same time, allowing for qualitative data of higher value. The method chosen provides ideators with the choice to behave differently regarding the usage of inspirations.

3.1.1. Study Design

The exploratory study has been designed according to the methodology of *participant observation*, which is the "observation of some social event (...), and explanations of its meaning" [BG57]. Therefore, it makes for "a tool for collecting data about people, processes, (...) in qualitative research" using "interviews, document analysis or surveys" [Kaw05]. Its goal is to "develop a holistic understanding of the phenomena under study" [MD10], although it has to be accepted that the description of the events will be "inevitably (...) selective" [Pes93].

An offline, in-situ approach has been chosen, because the user surveys are considered to be the most important artifact from the study and "web respondents might be less inclined to expend the necessary mental energy to answering survey questions" since they "might be multitasking, whereby they are engaged in several activities besides answering the survey questions" [Hee09]. Accordingly, crowd workers have been "shown to produce a higher *don't know* response rate, to differentiate less on rating scales, and to produce more item nonresponse than face-to-face survey respondents" [HL08].

Through a written-form, nonvocal brainstorming method, the conditions are ensured to still be very close to an actual EBS session on a computer although being face-to-face. Also, the chosen brainwriting pool method gives the ideators more freedom to express individual behavior as it is much less structured than other methods in which ideas are exchanged in fixed intervals like for example the 3-6-5 method.

For the exploratory study, three brainstorming sessions were conducted with five participants each (between-subject), see figure 3.1. The study began with an introduction to methodology and topic.

Figure 3.1.: Brainwriting pool session

Then, a 30-minute brainstorming session was performed for which the *brainwriting pool* method was used. In a brainwriting pool session, the participants sit together around a table and can execute the following two actions in any order and as often as they like:

1. Write a new idea on a sheet of paper (see Appendix figure A.2) and put it in the center of the table, called the *idea pool*.
2. Take one of the idea sheets submitted by another person from the idea pool and read it for inspiration; then put it back.

The execution of these actions was recorded on *execution sheets*, see Appendix figure A.1, so that the processed input and generated output by a participant can be reenacted. This brainstorming variant enables participants to gain inspiration from the other participants' ideas and guarantees that the individual actions of each participant are replicable. For all session audio recordings and photos were captured, for which all participants gave consent.

The study ended with a short survey, see Appendix figure A.3, in which the participant gave qualitative feedback on the session, the topic and the methodology.

3.1.2. Analysis

The generated ideas and surveys were digitized and processed as Excel sheets which are accessible as an open-source GitHub repository[1]. With these sheets, various statistical data (the number of submitted ideas, read inspirations etc.) and patterns

[1] https://github.com/kiwikern/master-thesis-ideator-types/tree/master/exploratory_study

of different actions (e.g. many idea submits in the beginning followed by alternating submits and inspirations) were analyzed.

On the ideas, an open coding was performed, based on grounded theory, a methodology "to facilitate an understanding" resting mainly on data [MT86]. As a Qualitative Data Analysis (QDA), it consists of three parts noticing, collecting and thinking about interesting things and is non-linear and recursive *[SC84]*. For the coding, the free version of the software Atlas TI[2] was used. Each brainstorming session was coded with a focus on the interaction with other participants in terms of reusage of their ideas as inspiration as well as the contentual relation of their own submitted ideas as a line-by-line (i.e. idea-by-idea) coding, while often coming back to already processed sessions to refine them. After all ideas were coded, the codes were grouped together in categories and subcategories by an axial coding.

The surveys were analyzed by focusing on how the participants perceived the study setup and its influence on their performance and differences in their behavior during the ideation session. This way, similar statements occurring multiple times were condensed into topics, and with these in mind, the surveys were reread to add all answers matching the grouping topics.

3.2. Results

After coding and refinement, 16 codes remained that were themselves again grouped by four non-exclusive dimensions, see table 3.1. First of all, it was differentiated whether an idea reused a previously seen or submitted idea, or if it was a context switch. The reusage of ideas was further classified into three subcategories:

- reusing an idea that was previously **seen vs. submitted**
- reusing an idea **directly vs. deferred**
- reusing an **entity vs. activity**

combining ideas	copying others idea	improving idea
reusing others application area	reusing others entity	reusing others older idea
reusing own activity	reusing own entity	reusing own older activity
reusing own older topic	reusing own topic	switching topic
reusing others activity	reusing others topic	reusing own older entity
using contextual cue		

Table 3.1.: List of codings of the ideas from the exploratory study

Overall, the 15 participants submitted a total of 225 ideas and read 220 ideas for inspiration. However, as can also be seen in the distribution of submits and inspirations (see figure 3.2), there was a range of different behaviors between the participants, which are examined more closely in this section.

[2]https://atlasti.com

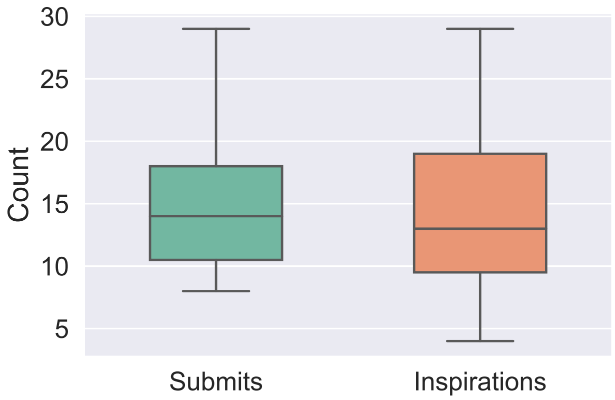

Figure 3.2.: Distribution of idea submits and read inspirations in the exploratory study

In the following section, participants will be labeled with a letter (A-E) and their session (1-3), e.g. C2.

3.2.1. Social Pressure

A couple of the participants felt pressured by the performance of other ideators:

C3 "At first I felt I had to hurry because I felt everyone were going faster than me, but after some minutes I decided to take my time to answer properly."

D2 "made me nervous how productive they seemed to be"

Others, however, felt positive reinforcement from the work atmosphere and liked that their ideas were not assessed during the session.

B3 "Everyone was working so you work"

B3 "evaluation for after (...) redeems you from the pressure (...) [so] you can access the good ideas in the end. Like a warm-up"

There was even a participant who missed getting feedback: "We don't get feedback of our ideas" (A2)

3.2.2. Idea Quality

Seven times ideas from others were copied with only slight changes directly after reading them, see table 3.2. This, of course, is very much in accordance with the rules of brainstorming in which the combination, improvement and reusage of ideas are even encouraged, see section 2.2. However, the added benefit stays small when the changes are minuscule.

"coat ceilings to get a night sky in your room" (D2)		"on the ceiling (...) so that we can see the sky" (A2)
"shirt that integrates (...) medical information" (E1)	$\Longrightarrow$	"clothes with (...) patients vital datas" (D1)
"Play solitaire with" (D3)		"play minesweeper" (B3)

Table 3.2.: Selection of copied ideas with only minor changes

When the context of the reused entity or action switches notably or is somewhat abstract, the derived idea might add value:

"mirror the sky for natural room light" (C1) $\Longrightarrow$ electronic "light-switcher" (E1)

3.2.3. Ideas of Others Considered Inspiring

13 of 15 participants had more than a third of their actions as reading. When participants reused concepts from the ideas they had seen, on average, 44% of it was based on the ideas of others. Participant D2, for instance, read 26 ideas for inspiration and reused these inspirations in two-thirds of their submitted ideas. Only five participants had less reusage of other ideators' ideas than the average (0% 0% 10% 20% 33%). This reusage was confirmed in the surveys as many participants found the ideas of others to be helpful, especially for showing them different use cases and enriching their train of thought:

B1	"showed me which areas were missing"
D1	"some of the other ideas broke me out of that thought-pattern"
E1	"showed very diverse use cases I was not thinking about"
B3	made them think "of an different area"
C3	"some of the ideas helped me think in a different way"
E3	"by providing topics that I hadn't thought before"
D2	"find new areas where to apply the technology"
E2	"I looked at them to change the area of application I was into"

For some participants, the opportunities to get inspirations even seemed not to go far enough. Participant A1 criticized that they "couldn't use the web for further inspiration". Similarly, participant E3 suggested: "it would help if we had more examples or even images that would stimulate/inspire us. Maybe then we could have had nicer ideas." Also, Participant E2 actively wished to "have an option to show random pictures".

3.2.4. Ideas of Others Considered Distracting

A1 read four ideas of other participants and did not reuse a single one of these. From the survey, it becomes clear that this was a rather deliberate choice:

A1 "No, I was too set in my ways. Didn't really want to change my methodology."

Although participants A2 and B2 had read a high number of ideas of other ideators (23, 14) they each reused only one of them and described the ideas as too similar or even distracting:

A2 "Some distract me because they require more than a magic coating to accomplish it."

B2 "Their ideas (...) distract me at some extent on the other hand, because their ideas just exchide [unreadable] the area which may comes to my mind later."

B2 "some ideas are similar and the inspiration of other ideas can be limited"

Ideator D3 reused only two ideas of other ideators, although 33% of their actions were reading ideas of others. Similarly, participants C2 and E2 each only reused one idea. Accordingly, they all described the ideas as too similar or distracting, especially when they contradicted their values:

D3 "When the ideas were way out of my imagination it distract me a bit."

D3 "Some ideas distract me because the use violated many of my thought and principle how people can be treated."

C2 "In some cases distracted me. They didn't sound right."

E2 "One did. Most of them were similar to mine."

E2 A software should not "show ideas which are similar"

C1 "often the ideas were pretty similar"

3.3. Conclusion

With the exploratory study, two different ideator types in respect to the usage of inspirations could be identified:

1. Ideators inspired by external input and actively looking for inspiration
2. Ideators distracted by external input and actively avoiding inspiration

Accordingly, one participant even stated it would be good when "you can choose if you would like to have other people brainstorming with you or if you feel disturbed by them" (D3).

In the results, the positive association of inspirations seemed to predominate, some, however, insistently stated they did not want to be distracted by the input from others. Why were there no ideators with no inspiration requests at all then? It can be expected that social pressure made the participants conform to the behavior of others. Also, participants tend to change their behavior to act according to the “true purposes of an experiment” which is called *demand characteristics* [Orn09]. Participants likely considered the usage of inspirations as expected of them as it was the most prominent feature of the experiment.

Siangliulue et al. already have shown that input provided in several distinct ways (none at all, on request or automatically) has varying impact on the overall quality (novelty) and fluency of the ideas [SCGD15]. However, it was not investigated whether the effects vary for different types of ideators. Siangliulue herself stated in her dissertation that future work should investigate “whether personalized inspiration could further help people benefit from inspirational examples” [Sia17].

With the findings of the exploratory study in mind, it is assumed that even stronger effects can be found when the ideators’ need for external input as inspiration is taken into consideration. As these findings were obtained in a classical brainstorming context, this assumption is tested quantitatively with the data from previous studies in an EBS setting in the following chapter.

4. Ideator Types

This chapter provides a more concrete definition of the individual differences found in the previous chapter. As these individual differences only apply to an ideation setting, they are called *ideator types* from here on. In order to transfer the types found in a classical group ideation setting into an EBS context, a quantitative analysis of data obtained in previous studies within the Ideas-to-Market project was conducted.

4.1. Definition

Gonçalves et al. described professional designers as *inspiration seekers* and *inspiration avoiders* [GCBS16] in a different context. These terms are used in this thesis to define the different ideator types as follows:

Inspiration Seekers are actively looking for inspiration during an ideation process and hence request more inspirations than the median of the number of inspiration requests.

Inspiration Avoiders feel distracted by external input and hence request inspirations at most once.

Ideators who cannot be clearly classified into these two categories are labeled with one of the following categories:

Unmotivated Ideators are not deeply invested in the given challenge and hence submit less than three ideas.

Undetermined Ideators lie between seekers and avoiders such they request inspirations more than once but less than the median of requested inspirations.

4.2. Applicability in an LSI Context

In order to ascertain whether the previously defined ideator types can also be found in an EBS setting, the data from three previous studies (chi19s2, chi19s3 and ac2) by Maximilian Mackeprang et al. [Mac18] have been analyzed. Although the studies differed in length, challenge and inspirations, all three brainstorming sessions enabled the participants to request inspirations. The frequencies of requested inspirations and submitted ideas were tracked over time per participant.

4.2.1. Methodology

The raw data of the three studies were available as a list of different events in JSON. It was transformed into 1-minute buckets with the number of events for each inspiration requests and idea submits.

For the following analysis, participants that submitted less than three ideas were filtered out as they were considered to have low motivation in participating in the study. Then, the inspiration requests of the participants at the very end as well as the chronological sequence of these were examined. The definitions of section 4.1 were applied to the data and the distribution at each point in time was compared to the distribution at the end of each study.

4.2.2. Results

In all three sessions, there were participants that requested a high number of inspirations (seekers) and participants that requested inspirations not more than once (avoiders), see figure 4.1. For all data sets, the inspiration requests are clearly not normally distributed. The distributions also show the tendency that avoiders are less frequent than seekers which confirms the expectation from the exploratory study.

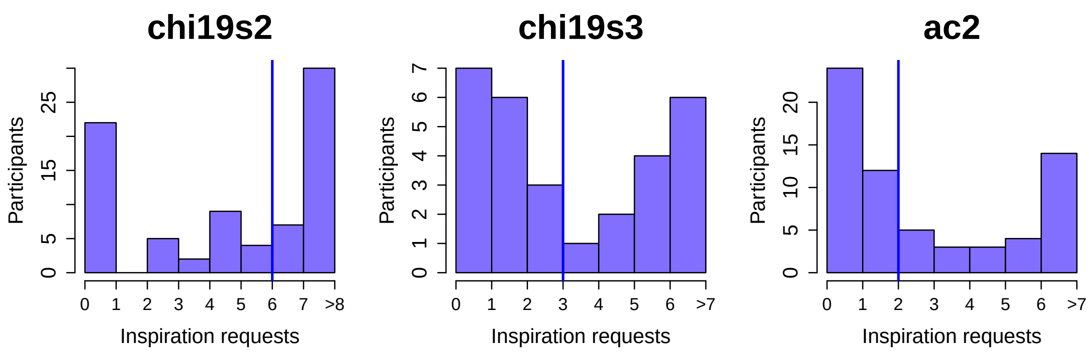

Figure 4.1.: Distribution of the number of inspiration requests showing motivated participants in absolute numbers. The median is depicted by a blue line.

The failure rates in figure 4.2 show how the number of wrongly classified ideators decreases over time compared to the classifications in the very end of each study.

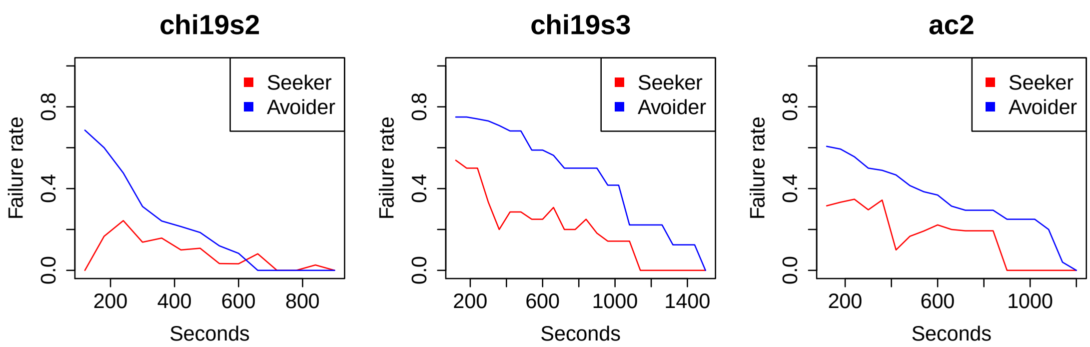

Figure 4.2.: Failure rates (false positives) for ideator type classification for motivated ideators only. For each point in time, the classified ideator types are compared to the results after the whole study duration.

4.2.3. Discussion

All three data sets could confirm the existence of the described ideator types in an electronic brainstorming context, while seekers appear to have a higher frequency than avoiders.

Hence, a brainstorming session in which participants have the opportunity to request inspirations can be used to classify them as seekers or avoiders. But how long should a brainstorming session be to classify by ideator type?

The failure rate analysis in figure 4.2 shows that at 10 minutes, the classification for every group is better than chance.

However, in the end, some participants seem to request inspirations merely out of boredom or because their creativity is exhausted and rather not because it is part of their ideation process. One participant in the chi19s3 study, for example, requested inspirations 19 times in the last 2 minutes. Therefore, the end might be biased against avoiders and the classifications at 10 minutes might actually be closer to the desired classification than the failure rate implies. Also, in general, the effects of boredom and exhausting are expected to be small for a shorter duration. Hence, as a good comprise between costs and performance, a 10-minute session was chosen.

4.2.4. Sequence Analysis

In order to test an alternative for the described classification, in collaboration with a researcher from the working group, the data were transformed into 1-minute buckets for idea submits and inspiration requests, and the ideator type was determined according to the definition in section 4.1. With this data, a model was fitted with a decision tree regressor[1], which predicts the ideator type by interpreting the sequence of idea submit and inspiration request events for each participant. Then, the model was visualized as a tree, see figure 4.3. The variables x_i represent the buckets as follows:

$$x_i = \begin{cases} i < 10 & idea\, submits\, in\, minute\, i+1 \\ i \geq 10 & inspiration\, requests\, in\, minute\, i-9 \end{cases}$$

A value of 1.0 represents a seeker; a value of 0.0 an avoider; *samples* represents the number of participants in the current node.

[1] https://scikit-learn.org/stable/modules/generated/sklearn.tree.DecisionTreeRegressor.htm

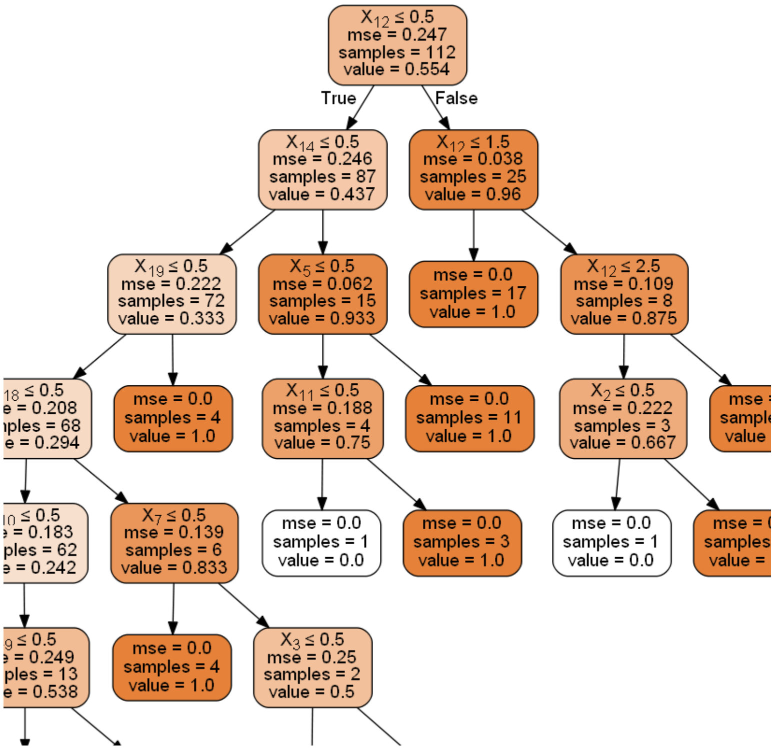

Figure 4.3.: Clipped illustration of a decision tree regressor model from the previous studies

The visualization shows x_{12}, the number of inspiration request in the third minute, as the most important feature in the classification. Only by looking at this single data point, the model can distinctly classify 17 seekers of the 112 participants. Further iterations with the model confirmed that the inspiration requests in the minutes 3 and 4 seem to be the most important feature for the classification.

Hence, this seems to be a very promising approach to get a classification quicker. The advance was not followed further, though, since the elicitation of the training data was too different from the setup in the upcoming quantitative study.

5. Quantitative Study

In the following study, the findings from the exploratory study from chapter 3 are tested quantitatively: Do different ideator types in terms of their need for inspiration exist? If so, does the ideator type have an impact on the quantity and quality of the produced ideas under different forms of exposure to external input? The results are also compared to the findings of the Timely Examples study [SCGD15].

5.1. Methodology

This section describes the study design as well as the statistical methods used to analyze the gathered data. As Siangliulue et al. could already show the influence of the way inspirations are provided, this study's methodology is based on theirs. However, it has to be extended to also allow for showing the impact of the ideator types at the same time.

5.1.1. Study Design

A between-subject 2x3 full factorial study design was used that was kept correspondent to the design of the Timely Examples study where possible so that, on the one hand, the results of the study can be replicated, and additionally, the impact of ideator types in combination with the inspiration conditions can be measured.

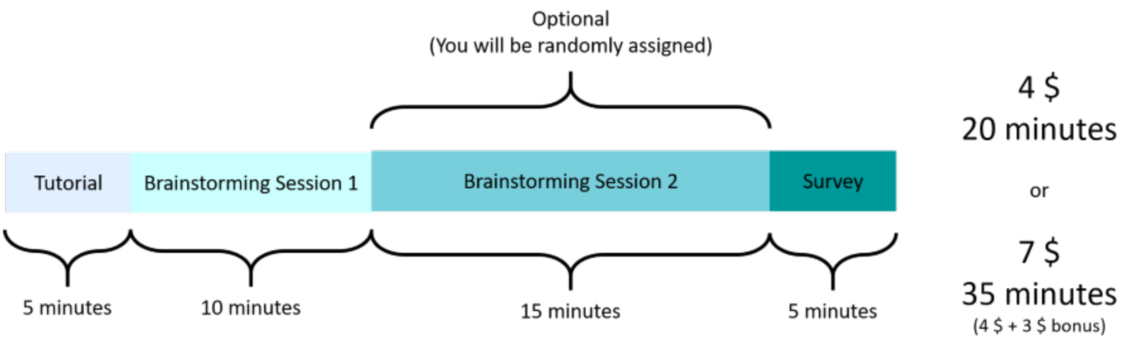

Figure 5.1.: Illustration of the study design presented to the crowd workers

Procedure

The study began with a short introduction of the task and the brainstorming tool. Then the participant worked on the first brainstorming session with the challenge *bionic radar*[1] for 10 minutes in which participants were able to request inspiration (on-demand). With the results of the first session, the ideator type of the participants was determined as defined above (see chapter 4.1). Only if they were assigned either the type *seeker* or *avoider* they continued with the second brainstorming session; otherwise, it was omitted. The second brainstorming session had the challenge *fabric display*[2] and a duration of 15 minutes.

In the end, the participant filled out a short survey including information about their demographics, inspiration usage and mental effort.

In total, it was aimed to recruit 120 participants that could be identified as either avoider or seeker. Accordingly, 20 avoiders and 20 seekers are sought for each of the three inspiration conditions. Since the median of inspiration requests, which is needed for the classification, is not known a priori, an ideator is considered a seeker when they request inspirations at least five times and are motivated (submits ≥ 3).

Only Amazon Mechanical Turk (MTurk) crowd workers who had participated in at least 1000 HITs, were U.S. native and had a minimum approval rate of 95% were allowed to participate in the study. For the shorter version, they received a compensation of $4 (estimated duration 20 minutes), for the longer one $7 (estimated duration 35 minutes) which corresponds to an hourly wage of $12. For an overview of the study design and compensation, which was also shown to the crowd workers, see figure 5.1.

Conditions

For the second session, the ideator was randomly assigned one of the following conditions that are taken from Siangliulue et al. [SCGD15]:

baseline No seed ideas are shown to the participant at all.

on-demand The participant can actively request a new seed idea to be shown for inspiration (one at a time).

on-idle Three different seed ideas are automatically shown to the participant when they do not type anything into the idea text field for at least 30 seconds.

The condition *on-interval*, which shows a new seed idea to the participants every 60 seconds, has been left out because it has proven to distract ideators during ideation.

[1] *A technology can perceive the movement of an object, such as humans, living beings or objects, like a bat. Once remembered, the technology can subsequently recognize it like a mother can pick her child out of a crowd. By comparing the object's movement with the movement of other objects, the technology can recognize movement patterns, like a good doctor can diagnose a knee disease by observing a specific walking motion. The technology is approximately hand-sized and can be used anywhere.*

[2] *Imagine there was a touch-sensitive 'fabric display' that could render high resolution images and videos on any fabric through a penny-sized connector. Brainstorm product ideas for this technology.*

Inspirations

In order to obtain the inspirations, ideas for both challenges from previous studies were manually rated for quality (in terms of novelty and value) by two persons (author and supervisor). The inspirational ideas were sorted by the sum of novelty and quality so that only high-quality ideas were shown as inspiration.

Implementation For the brainstorming, an application used for previous studies [MM19] was extended by the second session, the different conditions and the classification mechanism, see figure B.2. The changes affected the Spring Boot backend and its h2 database schema as well as the brainstorming application written in ClojureScript embedded in the instructions and surveys written as hbs templates.

The ratings of the generated ideas were executed with a rating tool developed for this purpose, see figure B.1. The tool was developed with Angular 8 and, again, Spring Boot 2.1 as an open-source project[3].

Both applications included a management console for Amazon MTurk and a JSON export interface for the collected data.

Metrics

After all brainstorming sessions were carried out, the generated ideas of the second session were counted per session (*fluency*) and rated by other MTurk workers in terms of their *novelty* ("Consider how novel, original or surprising the idea is") and their *value* ("Consider how useful the product idea is and how practical the idea sounds assuming the 'fabric display' technology is real"), both on a 7-point Likert scale, with the rating tool.

Each worker had to rate 15 ideas. Before rating each idea on its own, they first had to read all ideas. In the middle of the rating task, the workers had to answer a simple attention question. The duration of one session was estimated to be 6 minutes and offered a reward of $1. If a participant finished the task in less than 2.5 minutes, failed the attention question or rated each idea exactly the same, the rating session was not used.

All generated ideas were rated by at least three crowd workers each. Then, the ratings were normalized into z-scores between the rating sessions, and for each idea, the average of the ratings was calculated. This way, every idea obtained one averaged rating for each novelty and value. As an example, table B.2 shows three ideas with both the highest and lowest scores for each novelty and value.

The ideas were then grouped by brainstorming session, and the mean and maximum rating per session were calculated. The maximum ratings were additionally computed, although they were not part of the Timely Examples study. The maximum appears to be the more valid measure since brainstorming puts quantity before quality, see section 2.2, resulting in a high number of obvious or low-quality ideas

[3]https://github.com/kiwikern/RatingToolFrontend

and only a few surprising, valuable ideas. This phenomenon is also described by the serial order effect [BS12], which, in its essence, says that ideators first have to get all the bad and apparent ideas out of their system before they reach their best, like a warming-up process. It is particularly unlikely to produce many novel ideas because, by definition, they are unheard of and categorically far from other ideas.

5.1.2. Analysis

The data were exported from the brainstorming application as JSON files and then transformed with a node.js v10.15.3 script. For the statistical analysis of the data, both python jupyter notebooks[4] v4.4.0 and R with RStudio[5] v1.2.1335 were used. All scripts and data files can be found at the open-source repository[6].

The analysis was done with linear regression models since the data had multiple factors; the *lm* package of the R core library was used [R C19].

All linear models were tested against their assumptions according to Field et al. [FMF12], see Appendix B.2:

- assumption of independent errors tested with the durbin watson test
- assumption of no multicollinearity tested with the variance inflation factor
- assumption of homoscedasticity of the residuals visually confirmed with a Q-Q plot and the fitted values against the residuals

For idea submits, the data were log-transformed to ensure a normal distribution of model residuals as suggested by the Box-Cox test [BC64]. A probable cause for the positively skewed distribution is that ideators with a low number of submits were sorted out as *unmotivated*.

The contrasts were coded as described by Schad et al. [SVHK18]: The ideator type was coded as a sliding difference contrast, and the condition was coded as a custom contrast with on-demand vs. baseline and on-idle vs. baseline so that the estimates can be interpreted and the conditions are always compared to the baseline condition.

5.2. Results

In total, 414 MTurk crowd workers completed the study. Of these, 134 participants continued with the second session, from whom 122 sessions were considered usable. The participants that were sorted out had either submitted only non-sense ideas (e.g. “this a machine to know underwater scenery that makes the more deeps goes that due to make the changes”) or completely misunderstood the challenge (e.g. “Any item to make clipping cats nails easier”). The aimed numbers of 60 sessions per ideator type, 40 sessions per condition and 20 sessions for the product of type and condition were roughly met, see table 5.1.

[4] https://jupyter.org/
[5] https://www.rstudio.com/
[6] https://github.com/kiwikern/master-thesis-ideator-types/tree/master/quantitative_study

		Conditions			
		on-idle	on-demand	baseline	Σ
Ideator Types	Seeker	21	21	21	63
	Avoider	20	20	19	59
	Σ	41	41	40	122

Table 5.1.: Usable participants per type and condition in session 2

On average, the generated 1132 ideas were rated by five workers, while no idea had less than three ratings, see table 5.2.

Number of ratings	Number of ideas
3	125
4	147
5	566
6	226
7	37
8	22
9	4
10	5
Σ	1132

Table 5.2.: Number of ratings per idea (e.g. 125 ideas were rated by 3 workers)

In table B.2, three ideas with both the highest and lowest scores can be seen for each novelty and value. These ratings are considered valid, as short and obvious ideas received low ratings and elaborate and surprising ideas obtained high ratings.

5.2.1. Comparison to Timely Examples Study

In this section, the results are compared to the Timely Examples study. Hence, only the conditions are used as the independent variable. As they did not pay attention to individual differences of the participants, the ideator types are left out for the moment.

Fluency

There was no significant effect of the conditions on fluency, as can be seen in table 5.3. Siangliulue et al. found a significant main effect, but the pairwise comparisons only found a significant difference between on-idle and on-interval. Since on-interval was left out in this study, this effect could not be replicated.

Predictors	Estimates	SE	Statistic	p
Grand Mean	2.10	0.05	45.22	<0.001
on-demand vs. baseline	-0.15	0.11	-1.36	0.177
on-idle vs. baseline	-0.01	0.11	-0.05	0.964
Observations	122			
R^2/ adjusted R^2	0.020 / 0.003			

Table 5.3.: Linear regression model predicting fluency based on condition (log-transformed)

In terms of the means, the values in the Timely Examples paper (baseline=10.9, on-demand=10.9, on-idle=13.8) are slightly higher than the means found in session 2 (9.5, 8.4, 9.8), see *Overall* in figure 5.2. Here, on-idle has the highest fluency as well, on-demand has, however, a lower fluency than baseline. The maximum difference between the means (1.4) was lower than in the replicated study (2.9).

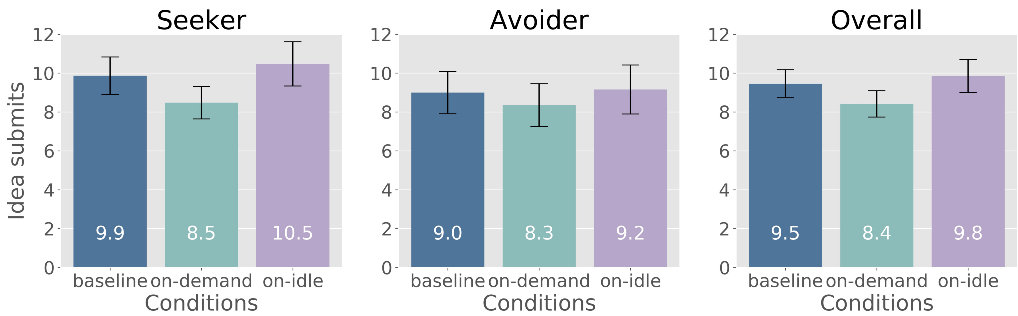

Figure 5.2.: Fluency per ideator type and condition in session 2. Overall are seekers and avoiders together. The error bars represent the standard error.

Mean Novelty

The Timely Examples study could show a significant main effect on the mean novelty of ideas and in the pairwise t-tests between on-demand (mean = 0.18) and on-idle (mean = -0.01) as well as on-demand and baseline (mean = -0.18). In this study, no significant main effect was found on mean novelty, see table 5.4. The highest mean novelty ratings were achieved in the on-demand condition, followed by on-idle and baseline, see figure 5.3a under *Overall*. The maximum difference between the means (0.04) was much lower than in the replicated study (0.36).

Predictors	Estimates	SE	Statistic	p
Grand Mean	-0.10	0.03	-3.29	0.001
on-demand vs. baseline	0.02	0.07	0.30	0.761
on-idle vs. baseline	0.02	0.07	0.23	0.820
Observations	122			
R^2/ adjusted R^2	0.001 / -0.016			

Table 5.4.: Linear regression model predicting mean novelty based on condition

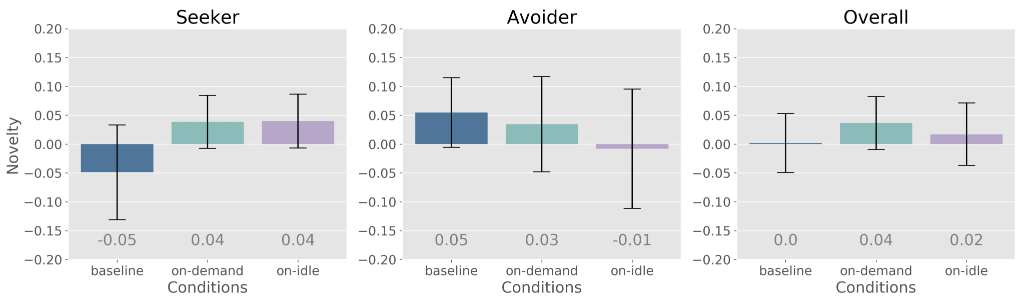

(a) Mean novelty ratings

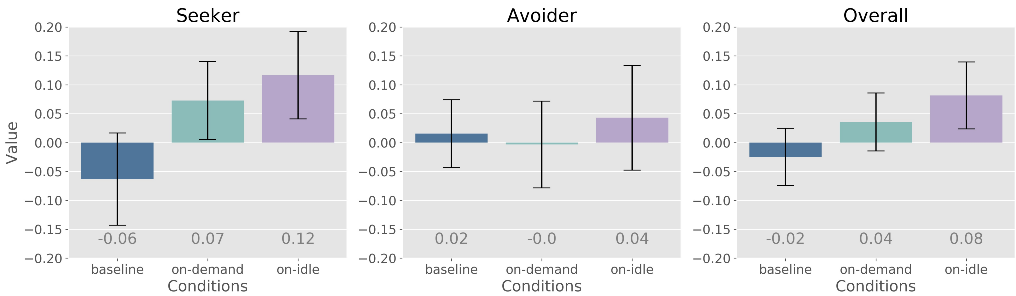

(b) Mean value ratings

Figure 5.3.: Mean idea ratings per ideator type and condition in session 2. Overall are seekers and avoiders together. The error bars represent the standard error.

Mean Value

Siangliulue et al. could not show a significant main effect on the mean value ratings; neither did this study yield a significant effect, see table 5.5. The maximum difference between the means (0.15) of baseline (-0.105), on-demand (0.049) and on-idle (0.015) is slightly smaller than in this study, see figure 5.3b under *Overall*. In both studies, baseline yielded the least valuable ideas.

Predictors	Estimates	SE	Statistic	p
Grand Mean	-0.08	0.03	-2.47	0.015
on-demand vs. baseline	0.08	0.07	1.05	0.297
on-idle vs. baseline	0.11	0.08	1.44	0.152
Observations	122			
R^2/ adjusted R^2	0.018 / 0.002			

Table 5.5.: Linear regression model predicting mean value based on condition

5.2.2. Existence of Ideator Types

In the first session of the MTurk study, both ideator types as defined in the beginning of chapter 4 could be identified. As expected from the results of the exploratory study and the pre-study, the number of seekers exceeded the number of avoiders distinctly; almost four times as many seekers as avoiders were found, see figure 5.4a.

60% of the avoiders in the first session requested inspiration once; the seekers requested an average of 12.8 inspirations, see figure 5.4b.

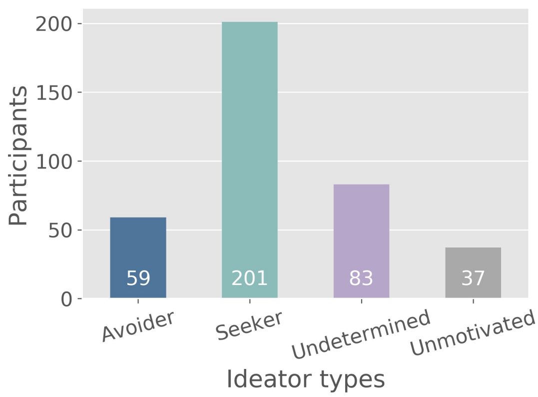

(a) Number of participants per ideator type

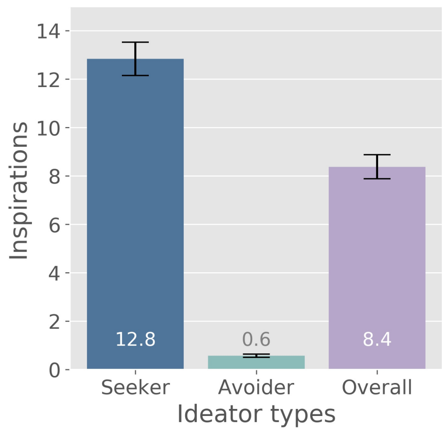

(b) Number of inspiration requests per ideator type

Figure 5.4.: Ideator types in session 1. Overall are seekers and avoiders together. The error bars represent the standard error.

Also in the second study, where in the on-demand conditions both types again had the opportunity to request inspiration, they mostly confirmed their associated behavior, such that seekers again requested a high number of inspirations (mean = 25.0) whereas avoiders used the inspirations button only rarely (mean = 3.5), see figure 5.5. A linear model predicting the number of inspiration requests for ideators in the on-demand condition shows a highly significant difference between seekers and avoiders (p-value < 0.001, $r = 0.412$, $r^2 = 0.402$, estimate $= 21.402$).

Following the ideator type definitions, all seekers in the on-demand condition would again be classified as seekers, whereas of the avoiders in the on-demand condition, 11 of 20 can be classified as avoiders again. The rest would be classified as undetermined (5), seekers (3) or unmotivated (1).

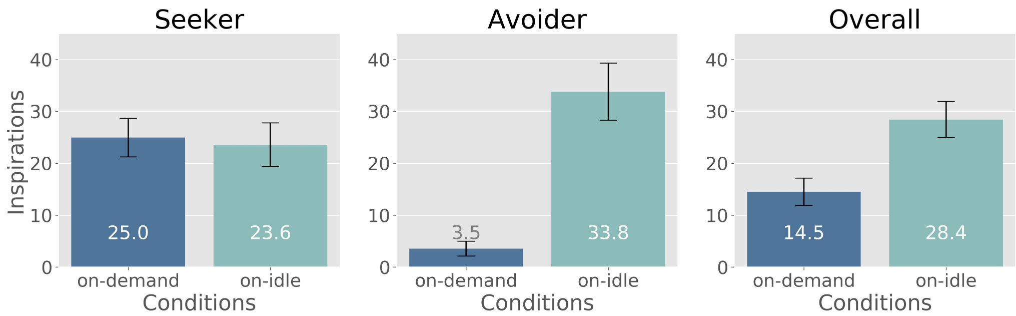

Figure 5.5.: Number of inspirations per ideator type and condition in session 2 (condition baseline omitted). Overall are seekers and avoiders together. The error bars represent the standard error.

5.2.3. The Impact of Ideator Types

The following section looks at the impact of the ideator type on the chosen metrics *fluency*, *value* and *novelty*. The latter two quality ratings are aggregated by calculating both the mean and the maximum for each metric.

Fluency

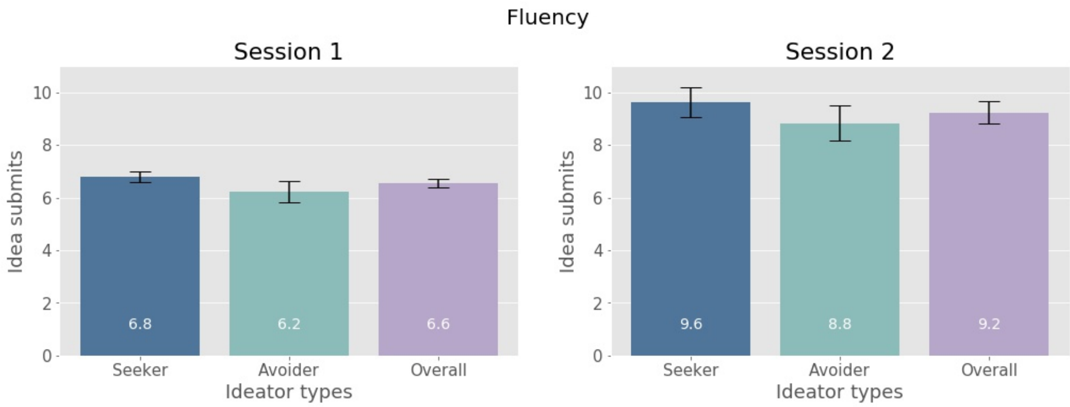

Figure 5.6.: Fluency per ideator type in both sessions. Overall are seekers and avoiders together. The error bars represent the standard error.

There were no main or interaction effects in the multiple linear regression on fluency between ideator types and conditions, see table 5.6. In both sessions, the means were slightly lower for avoiders than seekers, albeit non-significant, see figure 5.6. The largest difference between the means could be found between avoider:on-demand (8.4) and seeker:on-idle (10.5), see figure 5.2. The pairwise Tukey HSD comparisons show no significant differences, with the lowest p-value of 0.670 for the comparison between avoider:on-demand and seeker:baseline, see table B.3.

Predictors	Estimates	SE	Statistic	p
Grand Mean	1.78	0.04	42.42	<0.001
Seeker vs. Avoider	0.11	0.08	1.27	0.208
on-demand vs. baseline	-0.15	0.10	-1.48	0.141
on-idle vs. baseline	-0.08	0.10	-0.81	0.417
Seeker vs. Avoider : on-demand vs. baseline	-0.01	0.20	-0.04	0.966
Seeker vs. Avoider : on-idle vs. baseline	0.12	0.21	0.59	0.558
Observations	122			
R^2 / adjusted R^2	0.036 / -0.006			

Table 5.6.: Multiple linear regression model predicting fluency based on ideator type and condition (log-transformed)

Mean Novelty

There were no main or interaction effects in the multiple linear regression on mean novelty between ideator types and conditions, see table 5.7. Seekers have a slightly lower mean novelty (mean = 0.01) than avoiders (mean = 0.03) under all conditions, see figure 5.7. Especially in the baseline condition, the avoiders (mean = 0.05) outperform the seekers (mean = -0.05), see figure 5.3a; the difference, however, is not significant as the pairwise comparison with Tukey HSD shows (adjusted p-value = 0.911), see table B.4.

Predictors	Estimates	SE	Statistic	p
Grand Mean	0.02	0.03	0.63	0.532
Seeker vs. Avoider	-0.02	0.06	-0.29	0.772
on-demand vs. baseline	0.03	0.07	0.47	0.641
on-idle vs. baseline	0.01	0.07	0.18	0.859
Seeker vs. Avoider : on-demand vs. baseline	0.11	0.14	0.75	0.456
Seeker vs. Avoider : on-idle vs. baseline	0.15	0.14	1.05	0.296
Observations	122			
R^2/ adjusted R^2	0.013 / -0.030			

Table 5.7.: Multiple linear regression model predicting mean novelty based on ideator type and condition

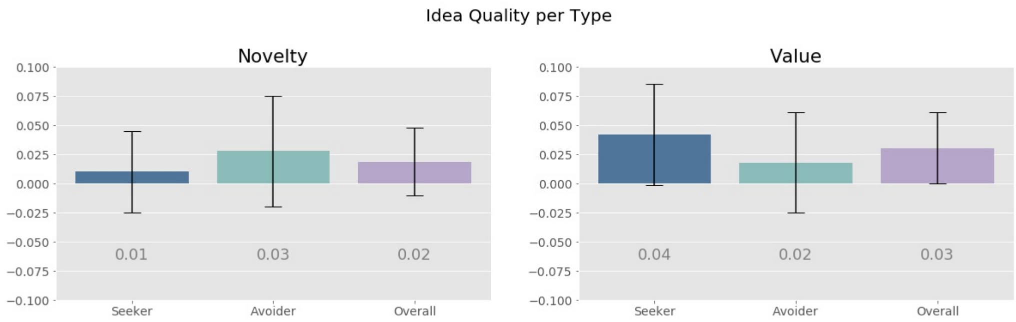

Figure 5.7.: Mean ratings per ideator type. Overall are seekers and avoiders together. The error bars represent the standard error.

Mean Value

There were no main or interaction effects in the multiple linear regression on mean value between ideator types and conditions, see table 5.8. Seekers have a slightly higher mean value than avoiders in all conditions (difference = 0.02), see figure 5.7. Seekers seem to have more valuable ideas in the on-idle condition (mean = 0.12) than the baseline condition (mean = -0.05), see figure 5.3b; the difference, however, is not significant as the pairwise comparison with Tukey HSD shows (adjusted p-value = 0.519), see table B.5.

Predictors	Estimates	SE	Statistic	p
Grand Mean	0.03	0.03	0.98	0.327
Seeker vs. Avoider	0.02	0.06	0.39	0.701
on-demand vs. baseline	0.06	0.07	0.79	0.433
on-idle vs. baseline	0.10	0.08	1.38	0.170
Seeker vs. Avoider : on-demand vs. baseline	0.15	0.15	1.04	0.302
Seeker vs. Avoider : on-idle vs. baseline	0.15	0.15	1.01	0.313
Observations	122			
R^2/ adjusted R^2	0.030 / -0.012			

Table 5.8.: Multiple linear regression model predicting mean value based on ideator type and condition

Maximum Novelty

In the multiple linear regression on maximum novelty between ideator types and conditions, there was a significant main effect for ideator types ($p < 0.1$) and significant interaction effects between ideator types and conditions ($p < 0.05$ and $p < 0.1$) see table 5.9. On average, an avoider's most novel idea had a higher rating than a seeker's most novel idea (difference = 0.11), see figure 5.9. But more importantly, the interaction effect significantly shows, that avoiders increase their novelty ratings when moving from on-demand or on-idle to baseline, whereas it is the other way around for seekers, see figure 5.8. Interestingly, all means of avoiders surpass the means of seekers, see figure 5.8b, when the avoiders in the on-demand condition are filtered out that would be classified as seekers in the second session, see section 5.2.2.

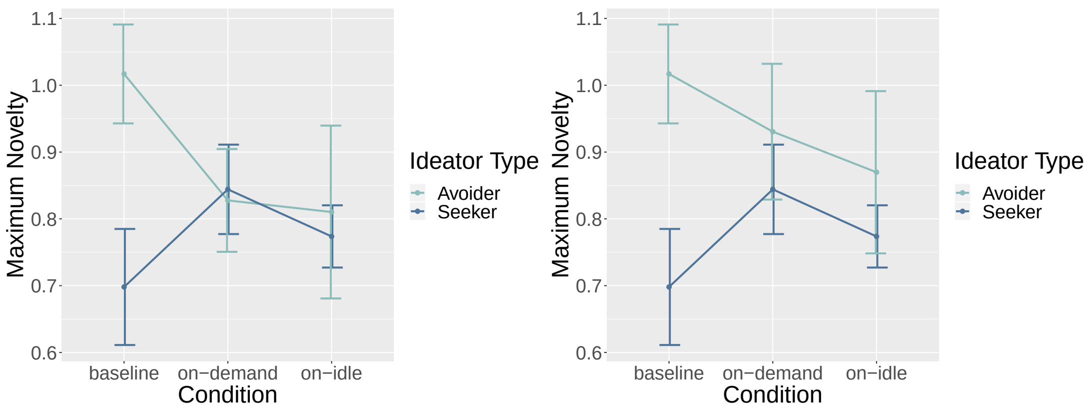

(a) Interaction for all avoiders

(b) Interaction filtered for confirmed avoiders in second on-demand session

Figure 5.8.: The interaction effect between ideator type and condition for maximum novelty. In figure (a), all data for session 2 are depicted, whereas in figure (b), in the on-demand condition the avoiders are filtered out that would have been classified as seekers in the second session, see section 5.2.2. The error bars represent the standard error.

The largest difference can be found between avoiders and seekers in baseline (difference = 0.32), see figure 5.10a; the pairwise comparison with Tukey HSD shows the difference between these two groups is the only significant one (adjusted p-value = 0.074), see table B.6.

Predictors	Estimates	SE	Statistic	p	
Grand Mean	0.83	0.03	24.61	<0.001	***
Seeker vs. Avoider	-0.11	0.07	-1.68	0.096	.
on-demand vs. baseline	-0.02	0.08	-0.26	0.793	
on-idle vs. baseline	-0.07	0.08	-0.79	0.430	
Seeker vs. Avoider : on-demand vs. baseline	0.34	0.16	2.04	0.043	*
Seeker vs. Avoider : on-idle vs. baseline	0.28	0.17	1.71	0.090	.
Observations	122				
R^2/ adjusted R^2	0.067 / 0.026				

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Table 5.9.: Multiple linear regression model predicting maximum novelty based on ideator type and condition

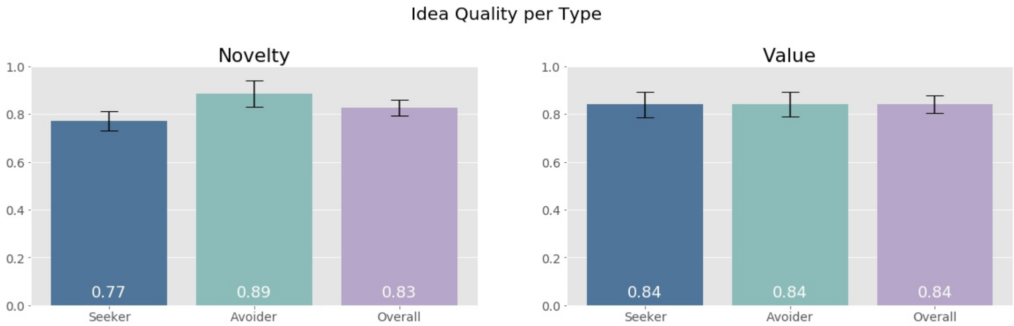

Figure 5.9.: Maximum ratings per ideator type in session 2. Overall are seekers and avoiders together. The error bars represent the standard error.

Maximum Value

There were no main or interaction effects in the multiple linear regression on maximum value between ideator types and conditions, see table 5.10. No difference at all was found in the means between the two ideator types, see figure 5.9. The biggest difference in means was found for seekers between on-demand (mean = 0.9) and on-idle (mean = 0.81), see figure 5.10b; the difference, however, is not significant as the pairwise comparison with Tukey HSD shows (adjusted p-value = 0.979), see table B.7.

Predictors	Estimates	SE	Statistic	p
Grand Mean	0.84	0.04	22.09	<0.001
Seeker vs. Avoider	-0.00	0.08	-0.01	0.993
on-demand vs. baseline	0.04	0.09	0.40	0.690
on-idle vs. baseline	0.01	0.09	0.15	0.878
Seeker vs. Avoider : on-demand vs. baseline	0.12	0.19	0.63	0.532
Seeker vs. Avoider : on-idle vs. baseline	-0.03	0.19	-0.14	0.891
Observations	122			
R^2/ adjusted R^2	0.007 / -0.036			

Table 5.10.: Multiple linear regression model predicting maximum value based on ideator type and condition

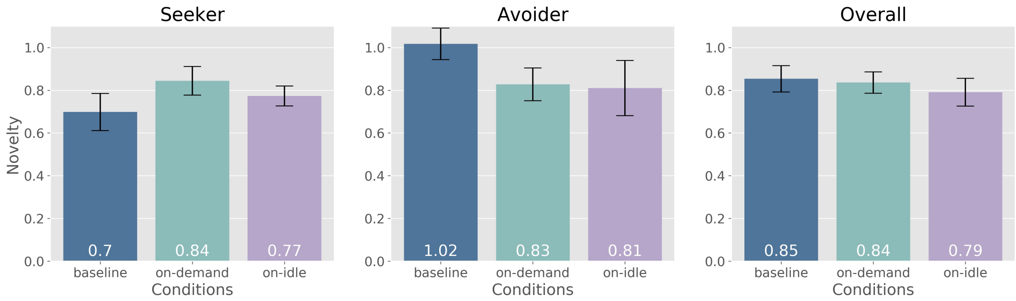

(a) Max novelty ratings

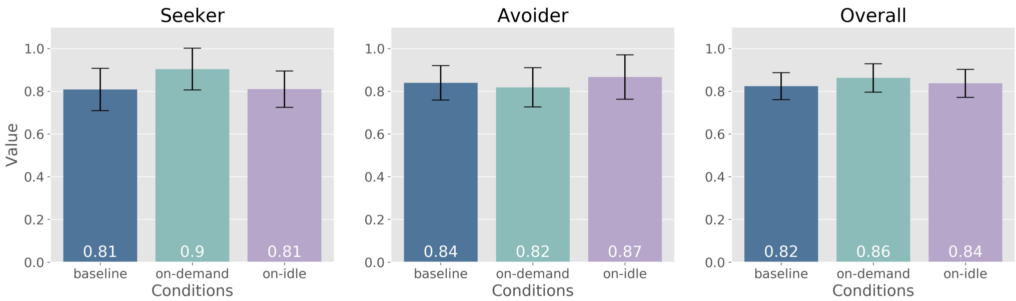

(b) Max value ratings

Figure 5.10.: Maximum idea ratings per ideator type and condition in session 2. Overall are seekers and avoiders together. The error bars represent the standard error.

5.2.4. Surveys

The survey at the end of the session inquired about the usage of inspiration, the overall task load and free text qualitative feedback. A selection of the free text feedback and the demographics can be found in the appendices B.4 and B.5.

Usage of Inspirations

The survey questions concerning the usage of inspirations showed the tendency that seekers viewed inspirations as more beneficial to their ideation process compared to avoiders, see figure 5.11. Avoiders had a higher ratio of neutral answers than seekers.

Furthermore, the participants were asked when they requested inspirations. Overall, avoiders confirmed that they requested inspirations less frequently than seekers, see table 5.11. The differences for the answers “in the middle”, “out of ideas” and “did not” were highly significantly different (p-value < 0.001); no significant difference were found for the answers “bored” and “before”. Only when they were bored, avoiders seemed more likely to request inspiration, albeit not significantly.

	before	in the middle	out of ideas	bored	did not
Seeker	43%	41%	74%	4%	0%
Avoider	32%	15%	37%	8%	24%

Table 5.11.: Survey answers for “When did you request inspiration?” per ideator type. Multiple answers were possible.

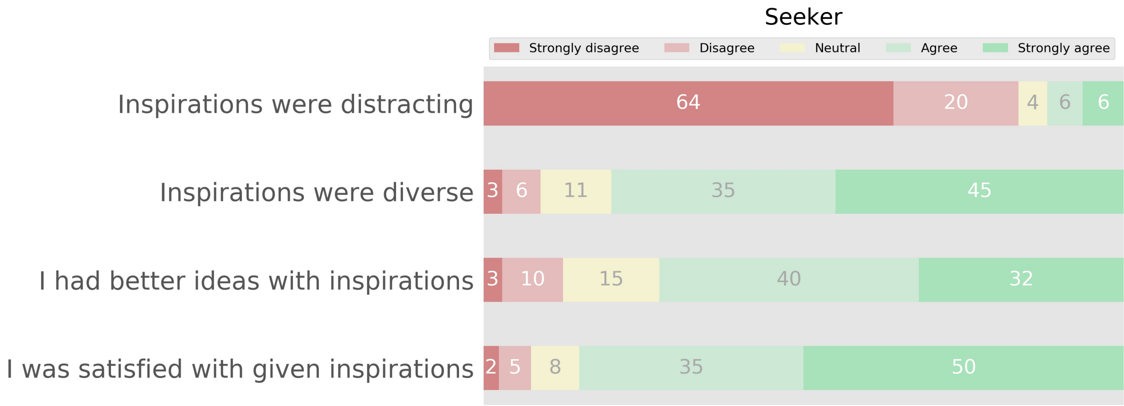

(a) Seeker survey

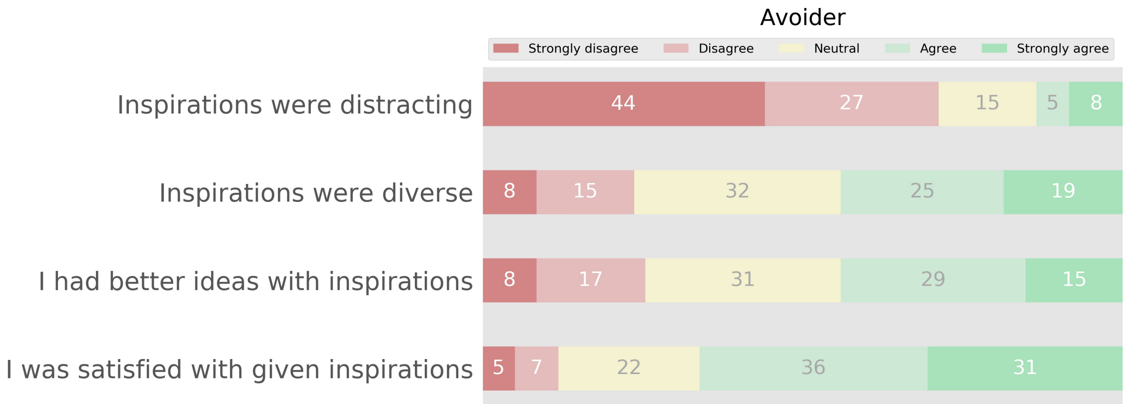

(b) Avoider survey

Figure 5.11.: Inspiration survey results per ideator type in percent. The questions were answered on a 5-point Likert scale.

Task Load Index

In the NASA Task Load Index (TLX) [HS88] survey questions, seekers had a significantly higher frustration level[7] and temporal demand[8] than avoiders, see table 5.12 and figure 5.12.

[7] How insecure, discouraged, irritated, stressed and annoyed were you?
[8] How hurried or rushed was the pace of the task?

These results are backed by the qualitative feedback of seekers who reported about temporal stress, e.g., “I get nervous when a time limit is shown.”, and high frustration, e.g., “I sometimes thought that my ideas were dumb and I was afraid that I would get rejected.”.

(a) TLX frustration levels (b) TLX temporal demand

Figure 5.12.: TLX survey answers per ideator type on a 7-point Likert scale where 0 stands for very low and 7 for very high. For the survey questions, see footnotes 7 and 8 on the preceding page.

TLX Indicator	Seeker - Avoider	SE	Statistic	p	R^2 / adjusted R^2	
Frustration	0.638	0.244	2.618	0.009	0.026 / 0.022	*
Temporal	0.491	0.247	1.986	0.048	0.015 / 0.011	.
Effort	0.111	0.181	0.612	0.541	0.001 / -0.002	
Performance	0.078	0.199	0.391	0.696	0.001 / -0.003	
Mental	-0.048	0.192	-0.251	0.802	0.000 / -0.004	
Observations	259					

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Table 5.12.: Linear regression models predicting TLX indicators based on ideator type

5.3. Discussion

In the last section, the results of the quantitative study were reported. This section offers interpretations of these results. Firstly, it is investigated why the results of the Timely Examples study could not be replicated [SCGD15]. Secondly, the impact of ideator types on the ideation outcome is evaluated, and possible explanations for the existence of ideator types are discussed.

5.3.1. Comparison to Timely Examples Study

Surprisingly, the significant effects found in the Timely Examples study could not be replicated. Here, the authors reported that the highest mean novelty ratings were achieved in the on-demand condition, whereas the highest fluency was observed in the on-idle condition. This thesis found no significant differences between the conditions.

Although the study conditions were very similar to the Timely Examples study in terms of length, inspiration delivery conditions and challenge, there were some relevant differences that might explain the different results:

Participants Most importantly, in this study, only seekers and avoiders were allowed to continue with the second brainstorming session. This selection excluded undetermined and unmotivated ideators and also distorted the distribution of seekers and avoiders because, in an unfiltered setting, seekers are almost five times as frequent as avoiders, see figure 5.4a. Meanwhile, in the Timely Examples study, only participants who requested inspiration more than once were used in the on-demand condition, which is effectively an exclusion of avoiders.

Duration and Setup In contrast to this study, the Timely Examples study began with an explicit, 3-minute warm-up task. The first task of this study could also be seen as a warm-up task; however, it was longer and had a different challenge. The longer total duration of the study can also explain the lower means in fluency of this study compared to Timely Examples (see section 5.2.3) as it may have led to higher levels of fatigue which is especially relevant in the high-pace setting of an MTurk HIT.

Furthermore, a different brainstorming application was used, and different seed ideas were given as inspirations as the seed ideas from the Timely Examples paper could not be obtained.

Analysis Further differences can be found in the analysis of the data, which was done with linear regression models, see section 5.1.2. Although the Timely Examples study used an ANOVA in combination with pairwise t-tests, linear models with Tukey HSD tests were preferred because they offer more insights and are statistically more valid as they correct for the cumulative type I error.

In the Timely Examples study, the sessions were only compared by their mean ratings; the maximum ratings per session were not reported and therefore, could not be compared with the results of this study. The authors of the Timely Examples study were contacted, but, unfortunately, the study data could not be obtained to calculate the maximum ratings retrospectively.

5.3.2. Existence of Ideator Types

The analyses have confirmed the expectations from the exploratory study and pre-study that there are ideators requesting a high number of inspirations during brainstorming, *inspiration seekers*, and at the same time, there are participants that explicitly do not want to see inspiration, *inspiration avoiders*, see section 5.2.2. The question remains why they choose to behave so differently.

Possible answers to this question might revolve around the ideators' frustration tolerance: Avoiders might be more resilient towards the pressure to be productive. This might be based on the ideator's confidence that is either task- or personality-related as the following paragraphs expand on.

Ideator's Confidence A possible explanation for the different behaviors is that avoiders might be more confident with the challenge at hand or even with brainstorming in general. Seekers, on the other hand, might be unsure if they understood the task correctly or question the quality of their ideas and hence look for confirmation in the inspirations.

This reasoning is supported by the fact that seekers had significantly higher task load indices in two dimensions: The higher frustration level ("insecure, discouraged, irritated", see footnote 7 on page 33) might arise when the participant does not know what to do or is unsatisfied with their own performance. One seeker expressed it as follows: "I sometimes thought that my ideas were dumb and I was afraid that I would get rejected."

Because of temporal pressure, as was reported by several seekers in their feedback, see Appendix B.4.2, as well as reflected in the significantly higher temporal demand ratings, see figure 5.12, the participants might feel rushed. This might encourage them to try to force their ideation and be less tolerant of longer periods of thinking. In both cases, ideators might try to start off their creativity with inspirations or compare themselves to the performance of others to get a better assessment of their own performance.

As avoiders reported a significantly lower task load, they might feel more confident in their own performance, and hence, they do not rely on external input. Their belief in a positive outcome allows them to have a higher tolerance for periods without output.

Personality- vs. Task-based Ideator Types It is yet unclear if the ideator type classification relates to general personality traits, or is rather influenced by the specific task at hand, i.e., a participant might be in general very confident in their performance but feels overwhelmed by the challenge or the other way around. The switch of three avoiders to seekers in the second session might either be explained by the different challenge or the limitations of the instructions, see section 6.2.

This study presented the participants with two similar, rather technical tasks and therefore can give little insight in this regard. It would be interesting to explore whether the ideator type classification stays stable when two very different challenges are compared.

5.3.3. Differences between Ideator Types

The analysis showed differences in the ideation outcome between the ideator types. After presenting a short theoretical background, possible explanations of these differences are given.

Theoretical Background

According to the SIAM model [NS06, NSL02], there are two different cognitional loops in the ideation process. In the first loop, the ideator searches their long-term memory (LTM) for images that match the given challenge. Images, in this context, are understood as an abstract unit representing a concept and its associated content. Once such a concept is identified, it is loaded into the short-term, working memory (STM) which can only hold one image at a time. Now, the second loop begins: For the currently loaded image, the ideator tries to come up with ideas relating to the image. When the ideator successfully finds an idea, they continue; when they fail, they may continue with the currently loaded image, or they discard the currently loaded image and start over with the first loop, see figure 5.13 illustrating the loops. Discarding the current image becomes more likely the more often the ideator has already failed to come up with a new idea for the image as the frustration of the ideator rises. Reading an idea for inspiration is considered to replace the currently loaded image as “inspirations can break an ideator’s train of thought” [SCGD15, NSL02].

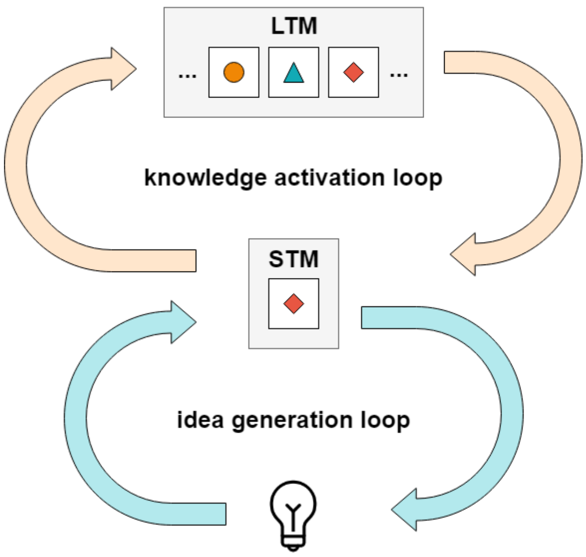

Figure 5.13.: The SIAM model includes two cognitional loops: The knowledge activation loop searches for a related image in the long-term memory (LTM) and loads it into the short-term memory (STM) when found. Then, in the idea generation loop, ideas are generated for the current image. When the concept is exhausted, the first loop starts again.

The serial-order effect [BS12], however, states that the quality and especially novelty of the ideas increases over the course of the ideation session. According to the spreading-activation theory [CL75], the reason for this increase lies in the organization of the mind: Related concepts (respectively images) are stored close to each other. Thus, very close and obvious links are found first, whereas more original and surprising links are found only when thinking longer about a concept. Although this is not yet explicitly verified, it can be assumed that the continuous improvement of idea quality does not only hold true for an ideation session as a whole, but also for the second loop in SIAM when trying to come up with new ideas within *one* category, as illustrated in figure 5.14. Therefore, it becomes more likely to produce ideas of higher novelty, the longer an ideator stays within one category.

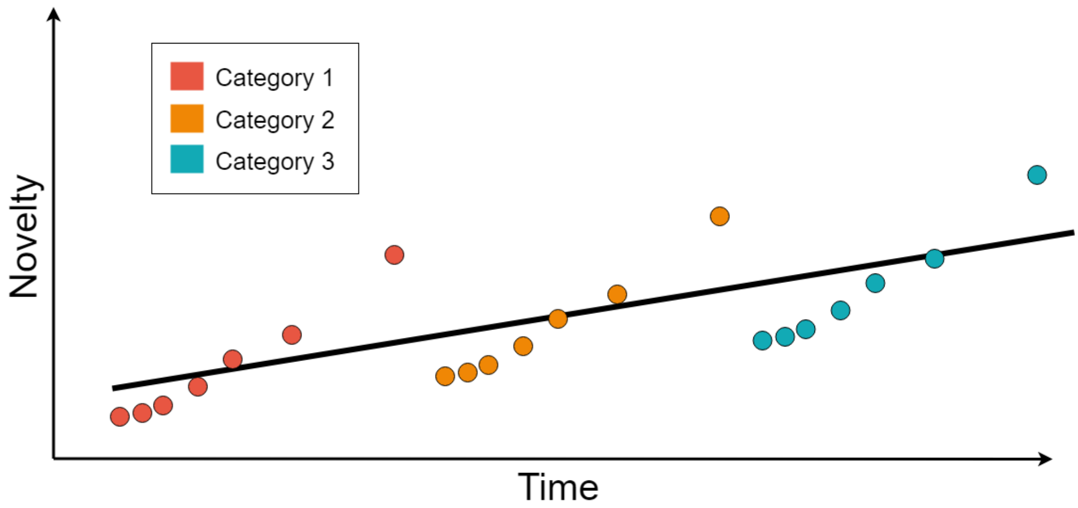

Figure 5.14.: Illustration of the serial-order effect per category with synthetic data. The dots represent ideas that are submitted over the course of a session with their novelty on the y-axis. The ideas are colored by their category (resp. image). The black line illustrates the overall increase in idea quality over time. The same increase is expected to happen within one category while obvious, lower-quality ideas are submitted first, and high-quality ideas are submitted after longer, fruitless periods without submits.

Higher Novelty Ratings for Avoiders

Throughout all conditions, avoiders, without the ones behaving likes seekers in the second session, have higher means in the maximum novelty ratings than seekers, see 5.8b. The higher ratings might be linked to the reduced task load and therefore, higher frustration tolerance, as explained in the previous section 5.3.2.

Also, the exposure to ideas of others' can lead to a fixation on certain categories [KS11] which, as novelty per definition is something rather unheard of, is counterproductive. As avoiders in the on-demand condition did not request inspiration and might have chosen to ignore the shown inspirations in the on-idle condition, fixation might have been less relevant for avoiders than seekers.

Interaction Between Ideator Type and Condition

The significant interaction effect between inspiration conditions and ideator types shows that avoiders have higher ratings in baseline compared to on-demand or on-idle, whereas for seekers, it is precisely the other way around, see figure 5.8b.

Hence, the self-selected exposure of inspirations of the two groups makes sense; it reflects in their performance. Seekers do not only complain about the missing inspirations in baseline ("I would have like have inspiration for the second tasks." (Seeker, baseline)), they actually perform worse when they do not have access to inspirations; a limitation of the ratings might even have weakened the actual effect, see section 6.3.1.

There are a number of factors that might contribute to the unexpected lower performance of avoiders in the on-demand condition than in baseline: The mere existence of the inspiration button might be distracting because mental capacity is needed to decide whether to ask for inspiration or not.

Furthermore, the button might seduce avoiders to rely on external stimulation, as it provides an easy way out of an unsuccessful thinking period. In this way, it might lower the frustration tolerance of the avoider and cut short their current trail of thought, which might still have produced novel ideas.

In her dissertation, Siangliulue offers a possible explanation of why in the Timely Examples study, on-idle outperformed on-demand in terms of novelty ratings:

> "To benefit from inspirational stimuli, being stuck is not enough — you must also know that you are stuck." [Sia17]

This explanation might also work the other way around:

> To not be distracted from inspiration, having the potential to discover more ideas is not enough — you must also know that you are not stuck.

In terms of SIAM, knowing that you are not stuck translates to staying in the second loop, the idea generation phase, for longer before discarding the currently loaded image. As described in the previous subsection Theoretical Background, the serial-order effect and the spreading-activation theory might suggest that the novelty of ideas increases when the ideator stays in the idea generation phase longer. Hence, a lowered frustration tolerance might lead to less novel ideas.

As seekers are expected to have a lower frustration tolerance in the first place, they, therefore, generate ideas with lower novelty. Especially, feeling a high temporal pressure does not allow for enduring a longer, unproductive phase within a category. For them, however, having a good starting point from an inspiration and knowing there is the opportunity to receive inspiration when stuck might lower the tension and allow them to stay within one category for longer.

No Effects for Mean Ratings or Maximum Value

No significant effects were found for ideator types and conditions on mean novelty or maximum and mean value.

The mean novelty and value ratings exhibit a large standard error, meaning that the mean ratings are highly dispersed. This variance might be related to the limited number of participants in this study. However, more importantly, this dispersion seems to be an inherent characteristic of brainstorming: As it values quantity over quality by definition, a large variety in the quality of ideas is expected with a few or even no high-quality ideas per session.

The lack of difference in maximum value is less easily attributed to the characteristic of brainstorming but might be related to differences between the dimensions novelty and value: Since obvious, unexciting and therefore non-novel ideas can still be highly valuable, it is considered easier to come up with valuable ideas than novel ones. Also, as a limitation, novelty might be easier to rate by novices than value, see section 6.3.1.

5.3.4. Practical Application

The interaction effect between the ideator types and the conditions clarifies that a one-size approach does not fit all. Thus, it makes sense to distinguish between different ideator types, especially when the ideation aims for highly novel ideas.

There are different ways of incorporating these findings into a brainstorming application: First, the ideator types have to be determined. This can be done by also adding a preceding classification session or by dynamically classifying the ideator type during the session and then either promoting inspirations for seekers or removing inspirations for avoiders. It's also conceivable that the ideator type is derived from survey answers before or during the session, for example about the proficiency in the topic and brainstorming in general or by inquiring the helpfulness of a displayed inspiration.

As avoiders had the highest novelty ratings overall, another possibility is to exclude seekers from the session and only allow avoiders to continue. Especially for long sessions, this might be an opportunity to save resources.

6. Limitations and Future Work

In this thesis, the existence of ideator types was derived from qualitative insight and confirmed by the analysis of data from previous studies. A user study then revealed their effects on the ideation outcome that have relevance for further research and practical development. However, the study was subject to some limitations of the study design and implementation, which this chapter reflects on while shedding light on some possible enhancements.

6.1. Interpersonal Differences

In the first brainstorming session, all participants ideated under the exact same condition: they could request inspirations by clicking a button. In the second session, they then ideated in one of three different conditions. In figure 6.1, the data of the first brainstorming session are grouped by the conditions of the *second session*. Here, although no difference between the groups is expected, a similar pattern in the differences can be found for the fluency in the first session, compared to the fluency of the second session, see figure 5.2: In both sessions, the on-demand groups submitted the lowest number of ideas, although, in the first session, there was no difference in the setup between the groups.

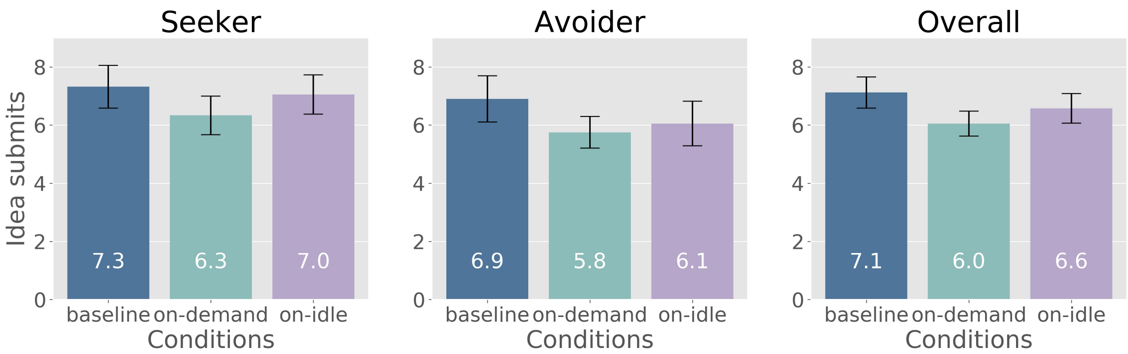

Figure 6.1.: Fluency per type in session 1 grouped by the condition of session 2

This unexpected difference suggests that there are interpersonal differences between the ideators that influence the outcome of the brainstorming session, possibly due to their different levels of creativity, motivation or familiarity with the topic. Future research could investigate whether the first classification session can be used to control for the differences between the ideators. If a low value in the first session predicts a similarly low value in the second session, for instance, a generalized linear mixed model (GLMM) could be used to control for these differences as the interpersonal differences (i.e. the participants) are treated as *random effects*. A GLMM could be especially interesting in combination with the quality ratings. Differences between the idea quality of the first session could again be treated as random effects. With these, the GLMM could correct for the general, setup-independent performance of an ideator towards the means.

6.2. Unclear Instructions

One potential impact on the results of the ideation is the participants' understanding of the procedure and task of the study. This section points out two possible weak spots and corresponding correctives.

6.2.1. Inspirations Mechanism

Although the study included introductory instructions and a short tutorial of the interface, it is not clear that all avoiders deliberately chose not to use inspirations or whether they might not have understood the inspirations mechanism.

One avoider in the on-demand condition stated in the survey "If it was more of a chat with several people, maybe they would have thought of more things by hearing the other ideas.", still they did not use the inspiration button in either of the two sessions. However, since the majority of the avoiders requested inspirations once, see figure 5.4b, and some explicitly stated in the survey that they did not want inspiration: "I did not use the inspiration. I wanted to use my own ideas and brainstorming." (avoider, on-demand), or that they just wanted to try out the possibility: "The use of the "inspiration" was curiosity. :)" (avoider, baseline), it is expected that the deliberate choice prevails.

In an improved version of the brainstorming application, an interactive tutorial might discard any lingering doubts that the inspirations mechanism is unclear.

6.2.2. Challenge Description

Some participants might have used the inspirations only to understand the challenge better and not to boost their creativity, as several survey answers show: "I was not clear at all from carefully reading the directions about what I was supposed to do. Reading the example at the start of the experiment, clarified this fully and increased my confidence. I looked for an example for the second part but the first part was adequate." (avoider, baseline) and "The description of the technology was somewhat ambiguous, so I used the inspiration to make sure I was interpreting it correctly. Once I was confident I'd understood the prompt, it was fun coming up with ideas." (Undetermined, none)

As a consequence, the challenge description could be tested for its clarity and maybe include a seed example.

6.3. Inspirations

Inspirations played a central role in this study. The way they were provided and how they were constituted had a substantial influence on the ideation outcome. Hence, it offers a potential opportunity for improvement.

6.3.1. Quality Ratings

Especially in large-scale *innovation*, the main goal of ideation is to generate ideas that add value (resp. novelty) to the knowledge that already exists. However, as observed in the exploratory study, sometimes inspirations are merely copied or slightly changed so that the added value is negligible, see section 3.2.2.

The same behavior is expected to be found in the quantitative study as well. Yet, this constitutes an inherent positive bias for seekers over avoiders and the conditions on-idle and on-demand over baseline. Especially since only highly rated ideas were shown as inspirations, an ideator that merely copied all inspirations would attain high ratings, which is not intended.

Even when ideas are not directly copied but instead combined or improved, as it is explicitly desired in classical brainstorming, see section 2.2, it is still up to debate how the rating of an improved idea can be compared to the rating of a completely new idea. The rating of a derived idea will most likely only slightly deviate from the rating of the original idea if the added benefit is small. If the rating of an idea that was created without inspiration is lower than the one of the derived idea, does this really mean that the derived idea is the better outcome for the innovation session? In practice, the added benefit of the new idea might be larger although it received lower ratings.

This description can be understood as a general critique of the practices in this field. The listed state-of-the-art related work did not mention a detection mechanism for copy-pasted inspirations or a better way to determine the *added benefit* instead of comparing mere ratings.

In order to compensate for this kind of bias, the produced ideas would need to be rated in terms of the added value to the existing ideas, respectively, the inspirations. For example, the provided inspirations in the on-demand and on-idle conditions could be shown next to the generated ideas in the rating sessions. This way, the raters can include the external influence in their rating decisions.

Future work should place additional focus on the influence of inspirations in terms of its integration into the generated ideas and its systemic influence on the quality ratings.

Furthermore, the ratings were conducted by novices and not domain experts, which renders them less conclusive [KBC+13]. Recognizing a very *novel* idea, something that has not been seen before, is considered to be easier and to require less domain knowledge than evaluating its real costs and benefits, its *value*.

6.3.2. On-idle Heuristic

In the condition on-idle, participants were considered idle when they did not type anything for 30 seconds, see section 5.1.1.

Siangliulue et al. already reported that their on-idle heuristic, which was reused in this thesis, might have room for improvement. Their analysis of the idle duration in between submitting an idea and requesting inspiration revealed a substantial variance over the course of an ideation session [SCGD15].

In this study, it was observed that some participants in the on-idle condition were exposed to a high number of inspirations compared to others who did not see any inspiration while at the same time not exhibiting a high fluency. Some even complained about the on-idle timer in the survey: "I found it more helpful in the first task to be able to choose when I get helpful ideas rather than to struggle and wait for suggestions." (Seeker, on-idle) A possible explanation is a relevant difference in typing speed as well as idea length per participant. The longer the participants needed to type out an idea, the less likely it was for them to be in such an idle period.

In order to account for these limitations, the on-idle mechanism could take the session timer and personal characteristics of the participants (e.g. mean time between idea submits) into consideration and thereby provide a more dynamic solution to predict idleness. For the development of such a heuristic, more data has to be gathered, for instance, the mean time between inspiration requests, idea submits and keystrokes.

6.4. Serial-Order Effect Within Categories

The serial-order effect offered an intuitive understanding of the higher novelty ratings for avoiders. In section 5.3.3, the explanation went beyond the insights that the serial-order effect gives: It was assumed that the quality of ideas does not only increase over the course of a session but also when ideating within a specific concept, see figure 5.14.

It could be very interesting to check this assumption with a data analysis. For this, the ideas need to be grouped in abstract, higher-level concepts. Since ideas are short snippets and state-of-the-art categorization mechanisms perform poorly on them [CGW$^+$09], more sophisticated approaches like Interactive Concept Validation [MM19] need to be considered.

With such a categorization, it can be tested, whether the quality within the same concept rises over time.

6.5. Survey Data

In both the exploratory study and the quantitative study, the survey feedback proved to be very useful.

In the first study, especially the free text feedback lead to the discovery of inspiration seekers and inspiration avoiders. The in-situ approach seems to have been a good choice, as the participants filled out the survey with care and gave very elaborate answers.

In the MTurk study, the qualitative feedback was less elaborate but still substantiated several statements, see Appendix B.4. In this context, the other survey questions gave valuable insights about the ideator's personality and state, e.g., the task load index.

Future research, especially research on individual differences, should lay a special focus on the survey questions and include questions about indicators such as the proficiency in the topic and brainstorming in general, the motivation, etc., both as quantitative and free-text questions. This might support the intuitive understanding of the encountered effects.

6.6. Sequence Analysis

The length of the first brainstorming session for the ideator type classification was determined by the analysis of the data of previous studies, see section 4.2. Since a session just for the ideator type classification demands a considerable amount of resources, it is interesting whether different ways exist to determine the ideator types and whether even a short duration suffices for a reliable result. As a first experiment, a sequence analysis was conducted with a decision tree regressor that showed promising results. For the analysis, the study duration (10 minutes) was segregated into ten 1-minute buckets. The visualization of the regressor's model illustrates that in most iterations of the model generation, the number of inspiration requests in minute 3 or 4 was the most important feature to predict the ideator type, which would allow for a much shorter classification session duration or even a dynamic classification during a session, see figure 4.3. With this, a dynamic inspiration system could even during a session automatically show and hide the possibility to request inspiration to adapt it to the ideators' needs interactively. Therefore, sequence analysis can be considered as a promising field for further research.

7. Conclusion

In the field of creativity research, many studies have been conducted on the relevance of inspirations during ideation. These studies have focused mainly on the characteristics of these inspirations and the way they are provided. Thus far, the effect of individual differences between ideators in their usage of inspiration during the ideation process has primarily been left unstudied.

This thesis first conducted an exploratory, in-situ brainstorming study to identify possible individual differences that are relevant regarding the usage of inspirations. As a result, two types of ideators were identified: inspiration seekers that actively look for inspirational stimuli and inspiration avoiders that feel distracted by external input and try to avoid it. As these results were found in a study setting with pen and paper, the data of previous studies from the Ideas-to-Market project were analyzed to confirm that the ideator type classification can be transferred into an EBS context as well.

In order to explore the impact that inspiration has on the ideation outcome of inspiration seekers and inspiration avoiders, a quantitative MTurk study has been conducted. The study consisted of two brainstorming sessions. The first one was used to classify the ideator as one of the two types; the second session had three different inspiration provision conditions taken from Siangliulue et al. and was evaluated in terms of fluency and quality (in this case defined as novelty and value).

First, the results were compared to the results of Siangliulue et al.: Surprisingly, the significantly higher mean novelty when ideators were able to request inspirations compared to when they received no inspirations at all, could not be replicated. This might be attributed to slight differences in the study design. Further analyses were not possible because the original data of the reference study were not obtainable.

The ideator types from the exploratory study were found again in the quantitative study. Seekers were almost five times as numerous as avoiders, all of which confirmed their assigned classification in the second brainstorming sessions. A few of the avoiders switched their behavior and requested inspirations in the second session. The ideator type classification likely depends both on general personal characteristics like confidence as well as task-specific factors like proficiency in the topic.

This study found significant differences between seekers and avoiders. A significantly higher maximum novelty was found in the ideas of avoiders compared to seekers. Novel, less obvious ideas are considered to be more likely, the longer a participant ideates within a train of thought or category. Staying within a category, however, might require a higher level of frustration tolerance, since the ideator has to endure a potentially fruitless period before coming up with an unexpected idea. As seekers reported a significantly higher task load than avoiders, it seems likely that avoiders have higher confidence in their performance, thus can stay within one category for longer.

Interestingly, avoiders had higher maximum novelty ratings without inspirations, whereas seekers benefited the most from inspirations on request. Consequently, inspirations might have a detrimental effect on the ideation outcome, based on the personal characteristics of the ideator. Avoiders might be distracted by the mere opportunity to request inspirations.

These results show that the incorporation of individual differences is a very promising approach for electronic brainstorming and large-scale ideation in general. Especially when highly novel ideas are the desired outcome of an ideation session, it is beneficial to distinguish between different types of ideators in terms of inspirations. It is up for future research to determine whether a quicker or dynamic classification is possible and if other characteristic factors like personality types and surveys can be used to identify the ideators' need for external stimuli. The sequence analysis described in section 4.2.4 showed first, auspicious results for the dynamic detection of ideator types and can be further developed in future studies. When different types of ideators are identified, application developers have to focus their attention on adapting interfaces and task descriptions to the different needs of these types. These systems can be envisioned as personalized inspiration systems.

Bibliography

[AI02] Daniele Archibugi and Simona Iammarino. The globalization of technological innovation: definition and evidence. *Review of International Political Economy*, 9(1):98–122, 2002. 1

[BBLK10] Ivo Blohm, Ulrich Bretschneider, Jan Marco Leimeister, and Helmut Krcmar. Does collaboration among participants lead to better ideas in it-based idea competitions? an empirical investigation. In *2010 43rd Hawaii International Conference on System Sciences*, pages 1–10. IEEE, 2010. 1

[BC64] George EP Box and David R Cox. An analysis of transformations. *Journal of the Royal Statistical Society: Series B (Methodological)*, 26(2):211–243, 1964. 5.1.2

[BG57] Howard Becker and Blanche Geer. Participant observation and interviewing: A comparison. *Human organization*, 16(3):28–32, 1957. 3.1.1

[BS12] Roger E Beaty and Paul J Silvia. Why do ideas get more creative across time? an executive interpretation of the serial order effect in divergent thinking tasks. *Psychology of Aesthetics, Creativity, and the Arts*, 6(4):309, 2012. 2.1, 5.1.1, 5.3.3

[BW08] Osvald M Bjelland and Robert Chapman Wood. An inside view of ibm's' innovation jam'. *MIT Sloan management review*, 50(1):32, 2008. 1, 2.3

[CDD16] Joel Chan, Steven Dang, and Steven P Dow. Comparing different sensemaking approaches for large-scale ideation. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, pages 2717–2728. ACM, 2016. 2.3

[CGW+09] Jonathan Chang, Sean Gerrish, Chong Wang, Jordan L Boyd-Graber, and David M Blei. Reading tea leaves: How humans interpret topic models. In *Advances in neural information processing systems*, pages 288–296, 2009. 2.3, 6.4

[CJV90] Terry Connolly, Leonard M Jessup, and Joseph S Valacich. Effects of anonymity and evaluative tone on idea generation in computer-mediated groups. *Management science*, 36(6):689–703, 1990. 1, 2.2

[CL75] Allan M Collins and Elizabeth F Loftus. A spreading-activation theory of semantic processing. *Psychological review*, 82(6):407, 1975. 2.1, 5.3.3

[CLC+15] Pui Kun Choo, Zhi Ning Lou, Bradley A Camburn, Kristin L Wood, Ben Koo, and Francois Grey. Ideation methods: a first study on measured outcomes with personality type. In *ASME 2014 International*

Design Engineering Technical Conferences and Computers and Information in Engineering Conference. American Society of Mechanical Engineers Digital Collection, 2015. 2.4

[CLM+04] Antonio Chirumbolo, Stefano Livi, Lucia Mannetti, Antonio Pierro, and Arie W Kruglanski. Effects of need for closure on creativity in small group interactions. *European Journal of Personality*, 18(4):265–278, 2004. 2.4

[DFS+11] Steven Dow, Julie Fortuna, Dan Schwartz, Beth Altringer, Daniel Schwartz, and Scott Klemmer. Prototyping dynamics: sharing multiple designs improves exploration, group rapport, and results. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2807–2816. Acm, 2011. 2.1

[DJRVY18] Kiki MM De Jonge, Eric F Rietzschel, and Nico W Van Yperen. Stimulated by novelty? the role of psychological needs and perceived creativity. *Personality and Social Psychology Bulletin*, 44(6):851–867, 2018. 2.4

[DS87] Michael Diehl and Wolfgang Stroebe. Productivity loss in brainstorming groups: Toward the solution of a riddle. *Journal of personality and social psychology*, 53(3):497, 1987. 1, 2.2

[DV93] Alan R Dennis and Joseph S Valacich. Computer brainstorms: More heads are better than one. *Journal of applied psychology*, 78(4):531, 1993. 2.3

[DW03] A Dennis and M Williams. Electronic brainstorming. *Group creativity: Innovation through collaboration*, pages 160–178, 2003. 2.3

[FMF12] Andy Field, Jeremy Miles, and Zoë Field. *Discovering statistics using R*. Sage publications, 2012. 5.1.2

[GBD17] Veronika Gamper, Andreas Butz, and Klaus Diepold. Sooner or later?: immediate feedback as a source of inspiration in electronic brainstorming. In *Proceedings of the 29th Australian Conference on Computer-Human Interaction*, pages 182–190. ACM, 2017. 1, 2.4

[GCBS16] Milene Gonçalves, Carlos Cardoso, and Petra Badke-Schaub. Inspiration choices that matter: the selection of external stimuli during ideation. *Design Science*, 2, 2016. 4.1

[GD87] George R Goethals and John M Darley. Social comparison theory: Self-evaluation and group life. In *Theories of group behavior*, pages 21–47. Springer, 1987. 2.2

[GDC+92] R Brent Gallupe, Alan R Dennis, William H Cooper, Joseph S Valacich, Lana M Bastianutti, and Jay F Nunamaker Jr. Electronic brainstorming and group size. *Academy of Management Journal*, 35(2):350–369, 1992. 2.3

[GR99] Alexander Gerybadze and Guido Reger. Globalization of r&d: recent changes in the management of innovation in transnational corporations. *Research policy*, 28(2-3):251–274, 1999. 1

[GTDS01] Monica J Garfield, Nolan J Taylor, Alan R Dennis, and John W Satzinger. Modifying paradigms—individual differences, creativity techniques, and exposure to ideas in group idea generation. *Information Systems Research*, 12(3):322–333, 2001. 2.4

[GWB17] Victor Girotto, Erin Walker, and Winslow Burleson. The effect of peripheral micro-tasks on crowd ideation. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 1843–1854. ACM, 2017. 1, 2.3

[GWB19] Victor Girotto, Erin Walker, and Winslow Burleson. Crowdmuse: Supporting crowd idea generation through user modeling and adaptation. In *Proceedings of the 2019 on Creativity and Cognition*, C&C '19, pages 95–106, New York, NY, USA, 2019. ACM. 1, 2.4

[Hee09] Dirk Heerwegh. Mode differences between face-to-face and web surveys: an experimental investigation of data quality and social desirability effects. *International Journal of Public Opinion Research*, 21(1):111–121, 2009. 3.1.1

[HJ85] Stephen G Harkins and Jeffrey M Jackson. The role of evaluation in eliminating social loafing. *Personality and Social Psychology Bulletin*, 11(4):457–465, 1985. 1, 2.2

[HK74] J Richard Hackman and Robert E Kaplan. Interventions into group process: An approach to improving the effectiveness of groups. *Decision Sciences*, 5(3):459–480, 1974. 2.2

[HL08] Dirk Heerwegh and Geert Loosveldt. Face-to-face versus web surveying in a high-internet-coverage population: Differences in response quality. *Public opinion quarterly*, 72(5):836–846, 2008. 3.1.1

[HS88] Sandra G Hart and Lowell E Staveland. Development of nasa-tlx (task load index): Results of empirical and theoretical research. In *Advances in psychology*, volume 52, pages 139–183. Elsevier, 1988. 5.2.4

[Kaw05] Barbara B Kawulich. Participant observation as a data collection method. In *Forum Qualitative Sozialforschung/Forum: Qualitative Social Research*, volume 6, 2005. 3.1.1

[KBC+13] James C Kaufman, John Baer, David H Cropley, Roni Reiter-Palmon, and Sarah Sinnett. Furious activity vs. understanding: How much expertise is needed to evaluate creative work? *Psychology of Aesthetics, Creativity, and the Arts*, 7(4):332, 2013. 6.3.1

[KC15] Mark Klein and Gregorio Convertino. A roadmap for open innovation systems. *Journal of Social Media for Organizations*, 2(1):1, 2015. 2.3

[KG15] Mark Klein and Ana Cristina Bicharra Garcia. High-speed idea filtering with the bag of lemons. *Decision Support Systems*, 78:39–50, 2015. 1

[KMMB18] Abderrahmane Khiat, Maximilian Mackeprang, and Claudia Müller-Birn. Innovonto: An enhanced crowd ideation platform with semantic annotation (hallway test). 2018. 2.3

[KS11] Nicholas W Kohn and Steven M Smith. Collaborative fixation: Effects of others' ideas on brainstorming. *Applied Cognitive Psychology*, 25(3):359–371, 2011. 5.3.3

[Mac18] Maximilian Mackeprang. Fub-hcc/innovonto-data: 2018 smart coating v1. Aug 2018. 4.2

[MD10] Kathleen Musante and Billie R DeWalt. *Participant observation: A guide for fieldworkers.* Rowman Altamira, 2010. 3.1.1

[MJS91] Brian Mullen, Craig Johnson, and Eduardo Salas. Productivity loss in brainstorming groups: A meta-analytic integration. *Basic and applied social psychology*, 12(1):3–23, 1991. 2.2

[MKGP06] Jonas Matthing, Per Kristensson, Anders Gustafsson, and A Parasuraman. Developing successful technology-based services: the issue of identifying and involving innovative users. *Journal of Services Marketing*, 20(5):288–297, 2006. 1

[MKM84] Joseph E McGrath, Janice R Kelly, and Diane E Machatka. The social psychology of time: Entrainment of behavior in social and organizational settings. *Applied social psychology annual*, 1984. 2.2

[MLH96] Richard L Marsh, Joshua D Landau, and Jason L Hicks. How examples may (and may not) constrain creativity. *Memory & cognition*, 24(5):669–680, 1996. 2.3

[MM19] Maximilian Stauss Maximilian Mackeprang, Claudia Müller-Birn. Discovering the sweet spot of human—computer configurations: A case study in information extraction. *Proc. ACM Hum.-Comput. Interact. 3, CSCW*, 2019. 5.1.1, 6.4

[MT86] Patricia Yancey Martin and Barry A Turner. Grounded theory and organizational research. *The journal of applied behavioral science*, 22(2):141–157, 1986. 3.1.2

[NDV+91] Jay F. Nunamaker, Alan R Dennis, Joseph S Valacich, Douglas Vogel, and Joey F George. Electronic meeting systems. *Communications of the ACM*, 34(7):40–61, 1991. 2.2

[Nij00] Bernard Arjan Nijstad. *How the group affects the mind: Effects of communication in idea generating groups.* Universiteit Utrecht, 2000. 1, 2.2

[NS06] Bernard A Nijstad and Wolfgang Stroebe. How the group affects the mind: A cognitive model of idea generation in groups. *Personality and social psychology review*, 10(3):186–213, 2006. 2.1, 2.2, 5.3.3

[NSL02] Bernard A Nijstad, Wolfgang Stroebe, and Hein FM Lodewijkx. Cognitive stimulation and interference in groups: Exposure effects in an idea generation task. *Journal of experimental social psychology*, 38(6):535–544, 2002. 1, 2.1, 2.2, 2.3, 5.3.3

[Orn09] Martin T Orne. Demand characteristics and the concept of quasi-controls. *Artifacts in behavioral research: Robert Rosenthal and Ralph L. Rosnow's classic books*, 110:110–137, 2009. 3.3

[OS18] A-M Olteţeanu and LH Shu. Object reorientation and creative performance. *Journal of Mechanical Design*, 140(3):031102, 2018. 2.4

[Osb48] Alex Osborn. *Your creative power.* Read Books Ltd, 1948. 1, 2.2

[Osb53] Alex F Osborn. Applied imagination. 1953. 1, 2.2

[PBGH99] Alain Pinsonneault, Henri Barki, R Brent Gallupe, and Norberto Hoppen. Electronic brainstorming: The illusion of productivity. *Information Systems Research*, 10(2):110–133, 1999. 1, 2.2, 2.3

[Pes93] Alan Peshkin. The goodness of qualitative research. *Educational researcher*, 22(2):23–29, 1993. 3.1.1

[PN03] Paul B Paulus and Bernard A Nijstad. *Group creativity: Innovation through collaboration.* Oxford University Press, 2003. 1, 2.2

[PS94] Andrea L Patalano and Colleen M Seifert. Memory for impasses during problem solving. *Memory & Cognition*, 22(2):234–242, 1994. 2.1

[R C19] R Core Team. *R: A Language and Environment for Statistical Computing.* R Foundation for Statistical Computing, Vienna, Austria, 2019. 5.1.2

[SC84] John V Seidel and Jack A Clark. The ethnograph: a computer program for the analysis of qualitative data. *Qualitative sociology*, 7(1-2):110–125, 1984. 3.1.2

[SCDG16] Pao Siangliulue, Joel Chan, Steven P Dow, and Krzysztof Z Gajos. Ideahound: improving large-scale collaborative ideation with crowd-powered real-time semantic modeling. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, pages 609–624. ACM, 2016. 1, 2.3

[SCGD15] Pao Siangliulue, Joel Chan, Krzysztof Z Gajos, and Steven P Dow. Providing timely examples improves the quantity and quality of generated ideas. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*, pages 83–92. ACM, 2015. 1, 2.3, 3.3, 5, 5.1.1, 5.3, 5.3.3, 6.3.2

[Sch75] Helmut Schlicksupp. *Grundlagen der Ideenfindung und Problemlösung, Page 186.* 1975. 2.2

[Sia17] Kanya Pao Siangliulue. *Supporting Effective Collective Ideation at Scale.* PhD thesis, 2017. 3.3, 5.3.3

[SVHK18] Daniel J Schad, Shravan Vasishth, Sven Hohenstein, and Reinhold Kliegl. How to capitalize on a priori contrasts in linear (mixed) models: A tutorial. *arXiv preprint arXiv:1807.10451*, 2018. 5.1.2

Appendix

A. Exploratory Study

A

EXECUTION SHEET

Please always write down an idea's number

1. *when you write an idea and put it on the table*
2. *when you take an idea from the table to read it*

1)	**22)**	**33)**
2)	**23)**	**34)**
3)	**24)**	**35)**
4)	**25)**	**36)**
5)	**26)**	**37)**
6)	**27)**	**38)**
7)	**28)**	**39)**
8)	**29)**	**40)**
9)	**30)**	**41)**
10)	**31)**	**42)**
11)	**32)**	**43)**
12)	**33)**	**44)**
13)	**34)**	**45)**
14)	**35)**	**46)**
15)	**36)**	**47)**
16)	**37)**	**48)**
17)	**38)**	**49)**
18)	**39)**	**50)**
19)	**40)**	**51)**
20)	**41)**	**52)**
21)	**42)**	**53)**

Figure A.1.: Execution sheet

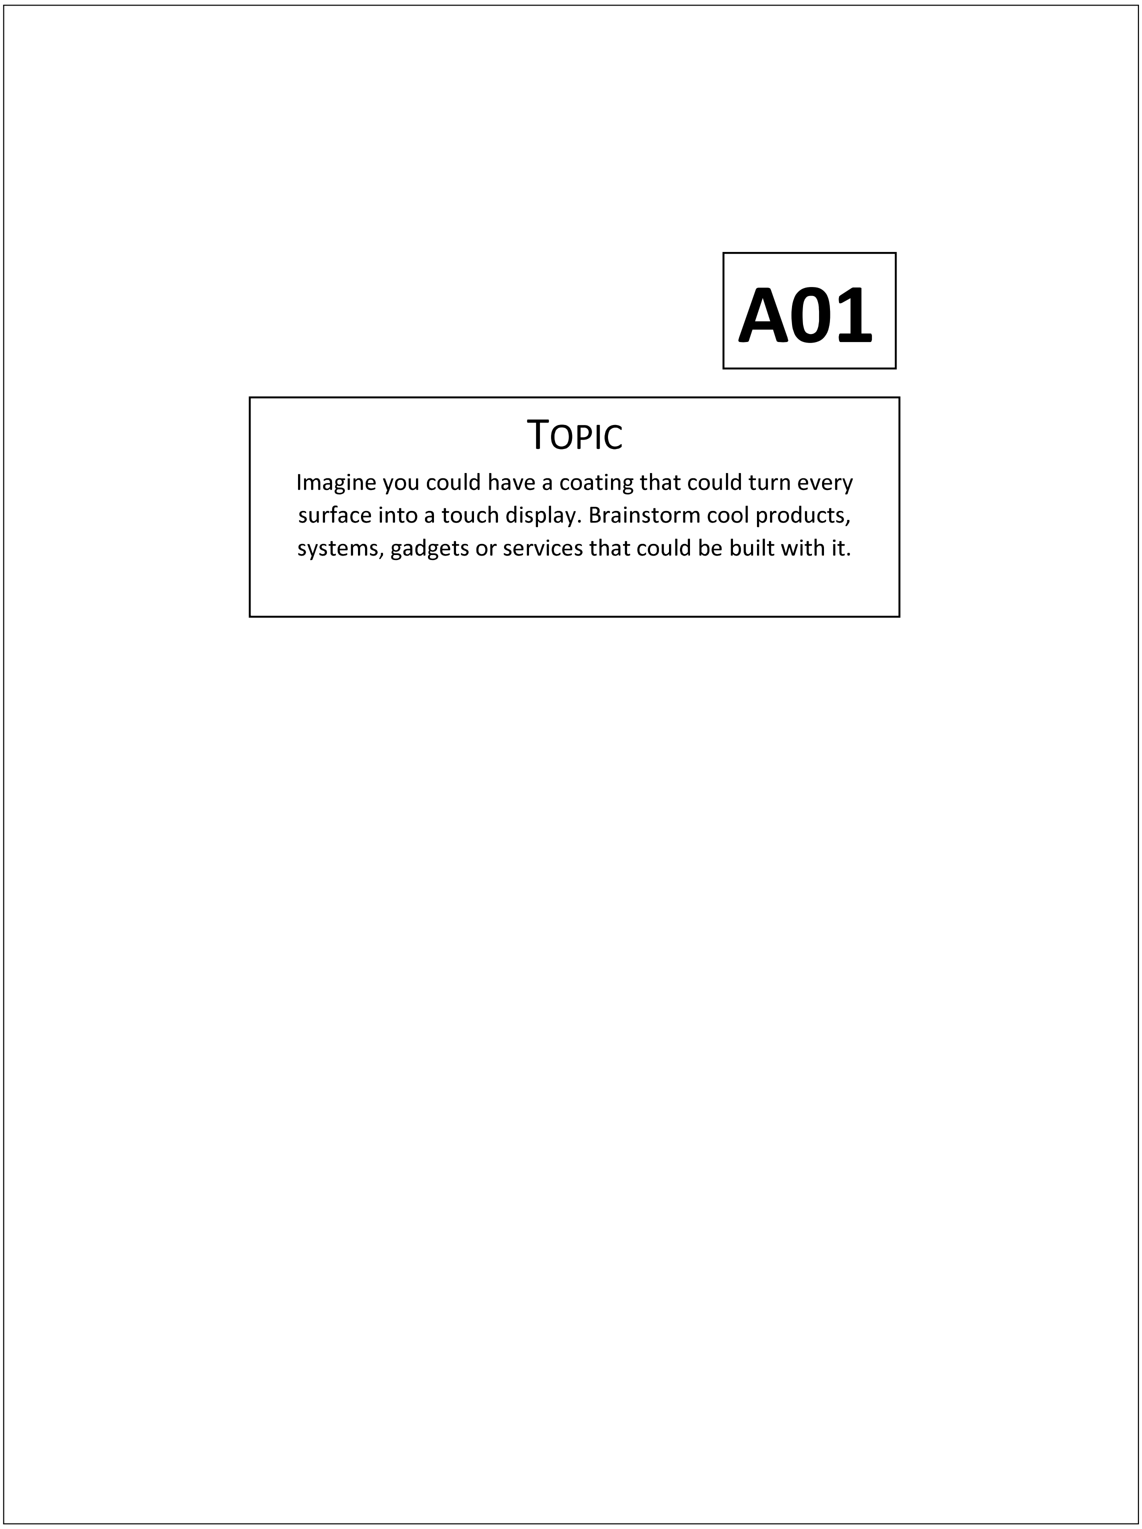

Figure A.2.: Idea template

A

SURVEY

When did you get ideas from the pool?

You can check multiple answers.

- ☐ When I was bored.
- ☐ When I ran out of ideas.
- ☐ Before I started generating my own ideas.
- ☐ In the middle of coming up with new ideas.
- ☐ Other:

Which things did you like or dislike about the Brainwriting Pool method?

Please mark your overall rating:

Very Good	Good	OK	Bad	Very Bad

Pros	Cons

How do you brainstorm in other contexts?

Can you compare it to the Brainwriting Pool method?

Figure A.3.: Survey (first page)

B. Quantitative Study

B.1. Tools

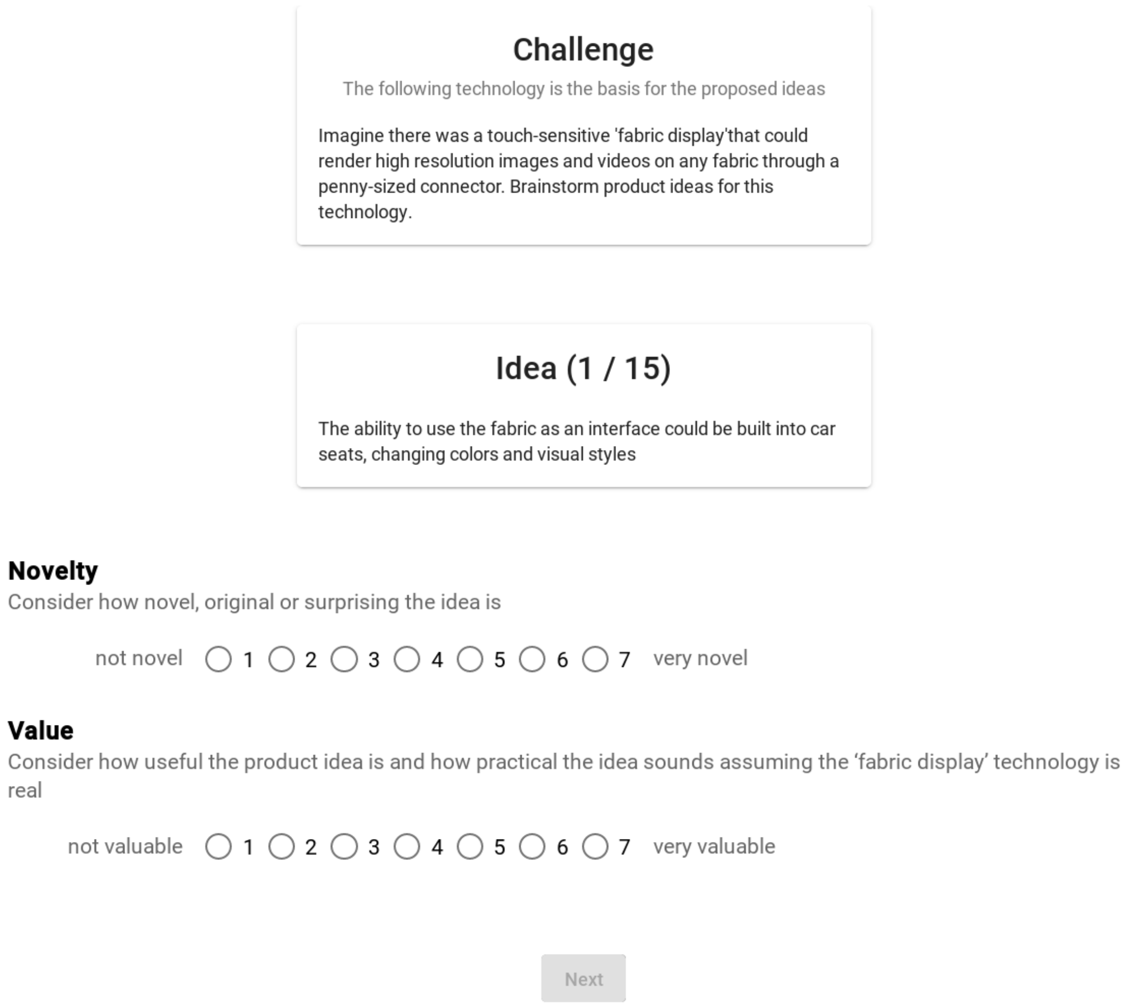

Figure B.1.: Rating tool to assess the quality of ideas by crowd workers

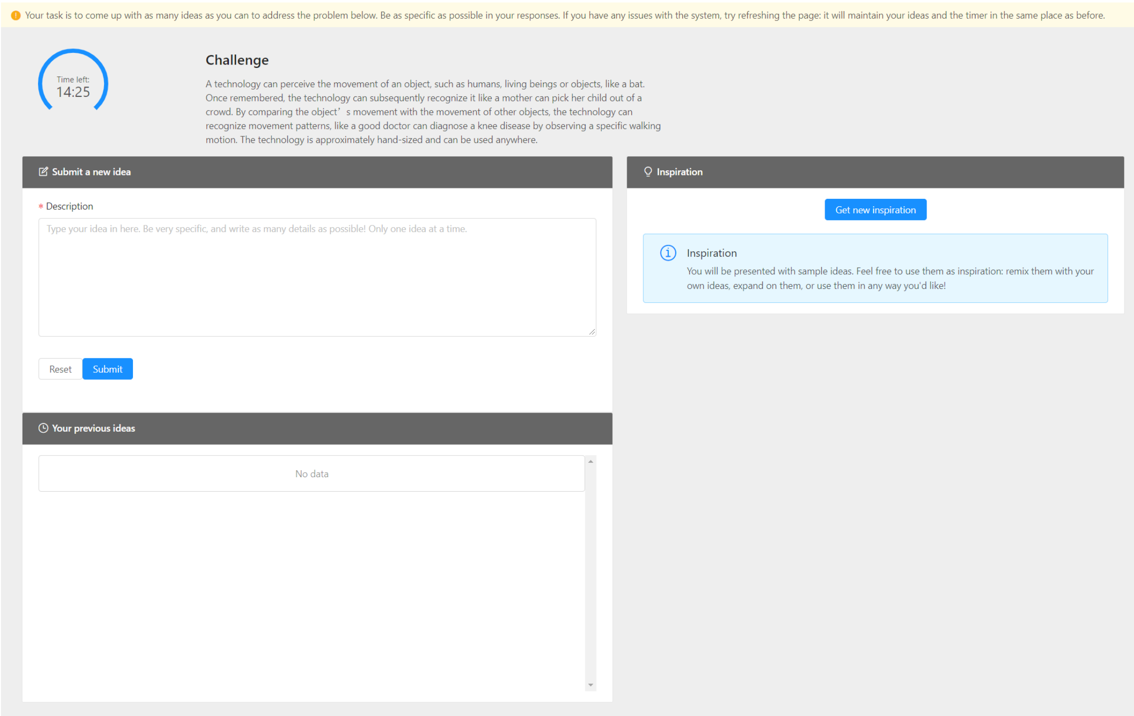

Figure B.2.: The interface of the brainstorming software. It showed the current challenge and a timer on the top, offered inspirations (depending on the active condition) and let the workers submit new ideas and view previously submitted ideas.

B.2. Accuracy Testing for Regression Model

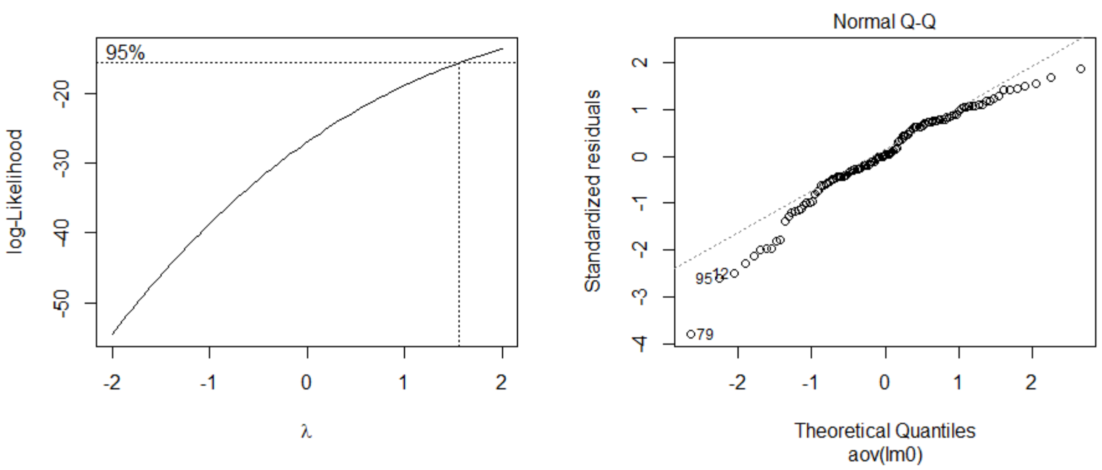

(a) Box-Cox test showing a normal distribution

(b) Q-Q plot for testing the assumption of homoscedasticity

Figure B.3.: Assumption testing for a linear model of ideator type and condition predicting mean novelty

Test	Value
Durbin Watson Test	2.026
Mean Variance Inflation Factor	1.223
Variance Inflation Factor Tolerance	0.888

Table B.1.: Assumption testing for a linear model of ideator type and condition predicting mean novelty

B.3. Results

Idea	Value
It could be used to help certain disabled people communicate how they're feeling and if they need something or someone.	1.71
A display to such a high resolution caliber would be a great asset to the medical community (surgeons) and medical professionals that are doing research for new cures of cancers and other diseases. This will allow for a clear and full visual in any space and any time.	1.66
This would be cool for people that are on some kind of medical surveillance. The fabric could display vital health information on a wristband or on a t-shirt that may be important for them to know. Maybe blood sugar info for diabetics.	1.66

(a) Ideas with highest value

Idea	Novelty
There can be a touch sensitive shirt that if you need to take your vitals it will do it at the touch of a button. All you have to do is tap on certain parts on your body and it will read the vitals for that area. If you tap the arm then it will have a built in blood pressure cup to take your pressure.	1.70
This type of technology could be used in the emergency area of hospitals, where the doctors would display a color or image on the patient's clothing depending on the severity of their condition, so it's easily identifiable and rapidly treated.	1.56
Furniture that can change styles/pattern with a tap. (...) Turn all your furniture from wood brown to all white or something on a whim. Changes up a whole room without moving a thing.	1.46

(b) Ideas with highest novelty

Idea	Value
A tablecloth for your dining room table.	-1.94
cover yourself in it a do some weird stuff like on each screen have a different video	-1.89
Office pigs could use it to point at their employees back and monitor what they're working on, to "increase productivity."	-1.84

(c) Ideas with lowest value

Idea	Novelty
Billboard	-2.41
Ceiling mounted tv over bed	-2.00
Of course this technology could be used communication devices such as laptops and phones	-1.95

(d) Ideas with lowest novelty

Table B.2.: Ideas with the highest and lowest novelty and value ratings. All ratings were assessed by at least three crowd workers as normalized means by calculating the z-scores per rating session.

	diff	lwr	upr	p_adj
Seeker:baseline - Avoider:baseline	0.06873	-0.350	0.487	0.997
Avoider:on-demand - Avoider:baseline	-0.14671	-0.570	0.277	0.916
Seeker:on-demand - Avoider:baseline	-0.08677	-0.505	0.332	0.991
Avoider:on-idle - Avoider:baseline	-0.14409	-0.573	0.285	0.925
Seeker:on-idle - Avoider:baseline	0.04537	-0.373	0.464	1.000
Avoider:on-demand - Seeker:baseline	-0.21544	-0.634	0.203	0.670
Seeker:on-demand - Seeker:baseline	-0.15550	-0.569	0.258	0.884
Avoider:on-idle - Seeker:baseline	-0.21282	-0.637	0.211	0.693
Seeker:on-idle - Seeker:baseline	-0.02335	-0.437	0.390	1.000
Seeker:on-demand - Avoider:on-demand	0.05995	-0.358	0.478	0.998
Avoider:on-idle - Avoider:on-demand	0.00262	-0.426	0.432	1.000
Seeker:on-idle - Avoider:on-demand	0.19209	-0.226	0.610	0.767
Avoider:on-idle - Seeker:on-demand	-0.05733	-0.481	0.367	0.999
Seeker:on-idle - Seeker:on-demand	0.13214	-0.281	0.545	0.939
Seeker:on-idle - Avoider:on-idle	0.18947	-0.234	0.613	0.787

Table B.3.: Pairwise comparisons with TukeyHSD for a multiple linear regression model predicting fluency based on ideator type and condition

	diff	lwr	upr	p_adj
Seeker:baseline - Avoider:baseline	-0.104	-0.398	0.191	0.911
Avoider:on-demand - Avoider:baseline	-0.020	-0.318	0.278	1.000
Seeker:on-demand - Avoider:baseline	-0.016	-0.311	0.278	1.000
Avoider:on-idle - Avoider:baseline	-0.063	-0.365	0.239	0.991
Seeker:on-idle - Avoider:baseline	-0.014	-0.309	0.280	1.000
Avoider:on-demand - Seeker:baseline	0.083	-0.211	0.378	0.963
Seeker:on-demand - Seeker:baseline	0.087	-0.203	0.378	0.953
Avoider:on-idle - Seeker:baseline	0.040	-0.258	0.339	0.999
Seeker:on-idle - Seeker:baseline	0.088	-0.202	0.380	0.950
Seeker:on-demand - Avoider:on-demand	0.004	-0.291	0.299	1.000
Avoider:on-idle - Avoider:on-demand	-0.042	-0.345	0.259	0.998
Seeker:on-idle - Avoider:on-demand	0.005	-0.289	0.300	1.000
Avoider:on-idle - Seeker:on-demand	-0.046	-0.345	0.252	0.997
Seeker:on-idle - Seeker:on-demand	0.001	-0.290	0.292	1.000
Seeker:on-idle - Avoider:on-idle	0.048	-0.250	0.347	0.997

Table B.4.: Pairwise comparisons with TukeyHSD for a multiple linear regression model predicting mean novelty based on ideator type and condition

	diff	lwr	upr	p_adj
Seeker:baseline - Avoider:baseline	-0.079	-0.385	0.227	0.976
Avoider:on-demand - Avoider:baseline	-0.019	-0.328	0.291	1.000
Seeker:on-demand - Avoider:baseline	0.057	-0.249	0.363	0.994
Avoider:on-idle - Avoider:baseline	0.027	-0.286	0.341	1.000
Seeker:on-idle - Avoider:baseline	0.101	-0.205	0.407	0.930
Avoider:on-demand - Seeker:baseline	0.060	-0.246	0.366	0.993
Seeker:on-demand - Seeker:baseline	0.136	-0.166	0.438	0.782
Avoider:on-idle - Seeker:baseline	0.106	-0.204	0.416	0.919
Seeker:on-idle - Seeker:baseline	0.179	-0.122	0.482	0.519
Seeker:on-demand - Avoider:on-demand	0.076	-0.230	0.382	0.979
Avoider:on-idle - Avoider:on-demand	0.046	-0.267	0.360	0.998
Seeker:on-idle - Avoider:on-demand	0.119	-0.186	0.426	0.866
Avoider:on-idle - Seeker:on-demand	-0.029	-0.340	0.280	1.000
Seeker:on-idle - Seeker:on-demand	0.043	-0.259	0.346	0.998
Seeker:on-idle - Avoider:on-idle	0.073	-0.237	0.384	0.983

Table B.5.: Pairwise comparisons with TukeyHSD for a multiple linear regression model predicting mean value based on ideator type and condition

	diff	lwr	upr	p_adj
Seeker:baseline - Avoider:baseline	-0.319	-0.655	0.0176	0.074
Avoider:on-demand - Avoider:baseline	-0.189	-0.530	0.1512	0.593
Seeker:on-demand - Avoider:baseline	-0.173	-0.509	0.1637	0.673
Avoider:on-idle - Avoider:baseline	-0.207	-0.552	0.1383	0.511
Seeker:on-idle - Avoider:baseline	-0.243	-0.580	0.0932	0.297
Avoider:on-demand - Seeker:baseline	0.130	-0.207	0.4660	0.874
Seeker:on-demand - Seeker:baseline	0.146	-0.186	0.4785	0.798
Avoider:on-idle - Seeker:baseline	0.112	-0.229	0.4532	0.931
Seeker:on-idle - Seeker:baseline	0.076	-0.257	0.4080	0.986
Seeker:on-demand - Avoider:on-demand	0.017	-0.320	0.3530	1.000
Avoider:on-idle - Avoider:on-demand	-0.017	-0.362	0.3276	1.000
Seeker:on-idle - Avoider:on-demand	-0.054	-0.390	0.2826	0.997
Avoider:on-idle - Seeker:on-demand	-0.034	-0.375	0.3070	1.000
Seeker:on-idle - Seeker:on-demand	-0.071	-0.403	0.2619	0.990
Seeker:on-idle - Avoider:on-idle	-0.037	-0.377	0.3044	1.000

Table B.6.: Pairwise comparisons with Tukey HSD for a multiple linear regression model predicting maximum novelty based on ideator type and condition

	diff	lwr	upr	p_adj
Seeker:baseline - Avoider:baseline	-0.031	-0.412	0.350	1.000
Avoider:on-demand - Avoider:baseline	-0.021	-0.406	0.364	1.000
Seeker:on-demand - Avoider:baseline	0.064	-0.316	0.445	0.996
Avoider:on-idle - Avoider:baseline	0.027	-0.363	0.417	1.000
Seeker:on-idle - Avoider:baseline	-0.029	-0.410	0.351	1.000
Avoider:on-demand - Seeker:baseline	0.010	-0.371	0.390	1.000
Seeker:on-demand - Seeker:baseline	0.095	-0.281	0.471	0.977
Avoider:on-idle - Seeker:baseline	0.058	-0.327	0.444	0.998
Seeker:on-idle - Seeker:baseline	0.002	-0.374	0.378	1.000
Seeker:on-demand - Avoider:on-demand	0.085	-0.295	0.466	0.987
Avoider:on-idle - Avoider:on-demand	0.048	-0.342	0.439	0.999
Seeker:on-idle - Avoider:on-demand	-0.008	-0.389	0.372	1.000
Avoider:on-idle - Seeker:on-demand	-0.037	-0.423	0.349	1.000
Seeker:on-idle - Seeker:on-demand	-0.094	-0.470	0.282	0.979
Seeker:on-idle - Avoider:on-idle	-0.057	-0.442	0.329	0.998

Table B.7.: Pairwise comparisons with TukeyHSD for a multiple linear regression model predicting maximum value based on ideator type and condition

B.4. Qualitative Survey Feedback

B.4.1. Instructions unclear

- “The inspirations were about applications of the technology, but the banner on the top seemed to me to request ways the technology could work. Confusing.” (Unmotivated, n/a)
- “The first ideas are related to possible product improvements. After getting inspiration my ideas change a bit.” (Seeker, n/a)
- “The description of the technology was somewhat ambiguous, so I used the inspiration to make sure I was interpreting it correctly. Once I was confident I’d understood the prompt, it was fun coming up with ideas.” (Undetermined, n/a)
- “Task was complicated as I was offered minimal instructions. The idea offered was very specific and was hard to apply.” (Seeker, n/a)
- “It would be nice to see any examples of how the technology, or one very similar to it, is currently used. ” (Undetermined, n/a)
- “Maybe include a stand in image of the device just so people have something to latch onto mentally. Otherwise was an interesting survey/ thought experiment. “ (Seeker, n/a)
- “Maybe provide a mock-up image of the technology to help give the person a better idea of it.” (Seeker, baseline)
- “the beginning of the task was a bit intimidating. Maybe soften the interface a bit.” (Seeker, n/a)

- "I would've liked slightly longer descriptions of the items we were brainstorming about. Maybe that's by design to keep it open, but I wasn't sure I could exactly picture the touch-activated fabric thing." (Avoider, on-demand)
- "I had to read a lot of the inspirations in order to understand what the idea was in the first place." (Seeker, n/a)

B.4.2. Temporal Pressure

- "I feel like if I had more time I could have come up with lots of or better ideas." (Avoider, on-idle)
- "I think there should have been a little more time to reflect on possible ideas. I had formulated other ideas but had not couched them in words." (Unmotivated, n/a)
- "I was a little surprised in just how good the suggestions that I saw were. I wish I had chosen to view them sooner" (Undetermined, n/a)
- "I just wish I had more time to come up with new ideas because once I got going ideas started to flow." (Seeker, n/a)
- "More time would have made me feel less pressure and enabled me to come up with more ideas." (Undetermined, n/a)
- "i thought wasnt enough time to finish the task i felt rushed and therefore it impacted the number and quality of the ideas i came up to" (Unmotivated, n/a)
- "I think you need a longer timer! I really wanted to come up with more (and even better) ideas, but the 10 minute timer just wasn't quite enough time to read the inspirations, formulate an idea, and get it all written down coherently. I was only able to complete two of them." (Unmotivated, n/a)
- "I think time is the limiting factor here. Once we have thought of the idea we have to write it down in an understandable manner which makes it take a bit long per idea." (Undetermined, n/a)
- "I think more time would lead to a better quality brainstorming session." (Undetermined, n/a)

B.4.3. Social pressure

- Don't show timer because "I get nervous when a time limit is shown." (Seeker, n/a)
- "The timer up at the left corner added a little layer of stress, because I was trying to submit as many ideas as possible, but I also wanted to make sure they were of good quality. " (Seeker, n/a)
- "with a timer , I felt a lot pressure (...) tough to be imaginative with the pressure" (Unmotivated, n/a)

- "I sometimes thought that my ideas were dumb and I was afraid that I would get rejected." (Seeker, n/a)

B.4.4. Inspirations

- "A tiny bit more variance in the suggestions would have been better, though." (Seeker, n/a)
- "coming up with original ideas was more difficult on the second round because the inspirations were similar to ideas i came up with" (Seeker, on-demand)
- "the inspirations were of great help" (Seeker, n/a)
- "I was avoiding accepting this hit because I thought I wouldn't be able to generate idea but the inspirations were so helpful and I really hope my ideas weren't too lame. I had fun thinking of different ideas. " (Seeker, n/a)
- "I would have like have inspiration for the second tasks." (Seeker, baseline)
- "the inspirations were of great help" (Seeker, n/a)
- "I was avoiding accepting this hit because I thought I wouldn't be able to generate idea but the inspirations were so helpful and I really hope my ideas weren't too lame. I had fun thinking of different ideas. " (Seeker, n/a)
- "the inspiration made it more challenge since they were all over the place." (Unmotivated, n/a)
- "the inspirations were so well thought out and I felt like a failure" (Seeker, n/a)
- "In the second exercise, I couldn't summon inspiration on demand like I did in the first task. When I got stuck, I was really hoping I could hit the button to see other ideas. Instead, they seemed to only pop up after I had not typed anything for a period of time. I would have preferred to have on-demand access to the inspiration tool." (Seeker, on-idle)
- "I was surprised how useful the inspiration ideas especially with how much it helped me as a starting point." (Undetermined, n/a)
- "I was not clear at all from carefully reading the directions about what I was supposed to do. Reading the example at the start of the experiment, clarified this fully and increased my confidence. I looked for an example for the second part but the first part was adequate." (Avoider, baseline)
- "I really didn't use the inspiration option, I just clicked on it to see what it would do." (Unmotivated, n/a)
- "the inspirations were both helpful and reassuring." (Seeker, n/a)

B.5. Demographics

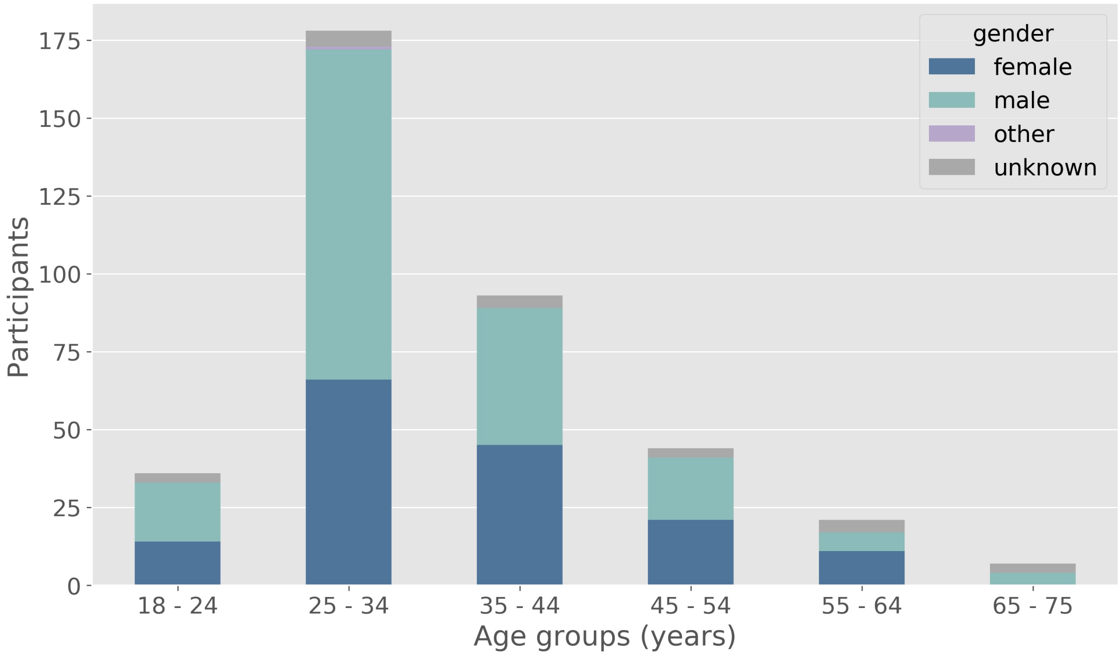

Figure B.4.: The demographics showing the number of participants per gender and age groups