

IdeateRelate: An Examples Gallery That Helps Creators Explore Ideas in Relation to Their Own

XIAOTONG (TONE) XU, University of California San Diego, USA

ROSALEEN XIONG, The University of Texas at Austin, USA

BOYANG WANG, University of California San Diego, USA

DAVID MIN, University of California San Diego, USA

STEVEN P. DOW, University of California San Diego, USA

Creating truly original ideas requires extensive knowledge of existing ideas. Navigating prior examples can help people to understand what has already been done and to assess the quality of their own ideas through comparison. The creativity literature has suggested that the conceptual distance between a proposed solution and a potential inspiration can influence one's thinking. However, less is known about how creators might use data about conceptual distance when exploring a large repository of ideas. To investigate this, we created a novel tool for exploring examples called IdeateRelate that visualizes 600+ COVID-related ideas, organized by their similarity to a new idea. In an experiment that compared the IdeateRelate visualization to a simple list of examples, we found that users in the Viz condition leveraged both semantic and categorical similarity, curated a more similar set of examples, and adopted more language from examples into their iterated ideas (without negatively affecting the overall novelty). We discuss implications for creating adaptive interfaces that provide creative inspiration in response to designers' ideas throughout an iterative design process.

CCS Concepts: • **Human-centered computing** → Information visualization; *Computer supported cooperative work*; **Interaction design theory, concepts and paradigms**.

Additional Key Words and Phrases: creativity support tool, CST, design, examples, ideation, visualization, conceptual distance

ACM Reference Format:

Xiaotong (Tone) Xu, Rosaleen Xiong, Boyang Wang, David Min, and Steven P. Dow. 2021. IdeateRelate: An Examples Gallery That Helps Creators Explore Ideas in Relation to Their Own. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 352 (October 2021), 18 pages. <https://doi.org/10.1145/3479496>

1 INTRODUCTION

To generate truly creative ideas, one needs to have extensive knowledge of existing ideas [46, 50]. Exploring ideas by others can help designers on creative work by exposing new perspectives and modeling how work gets done. While some research indicates that such examples can create fixation effects [25, 38, 43], many designers have adopted a practice where they explore prior examples throughout the process [33, 41]. The rise of the Internet has supported this practice by giving people exposure to vast examples across difference domains [48, 55]. Despite increased availability

Authors' addresses: Xiaotong (Tone) Xu, University of California San Diego, La Jolla, California, USA, xt@ucsd.edu; Rosaleen Xiong, The University of Texas at Austin, Austin, Texas, USA, rosaleenmao@gmail.com; Boyang Wang, University of California San Diego, La Jolla, California, USA, b2wang@ucsd.edu; David Min, University of California San Diego, La Jolla, California, USA, dmin@ucsd.edu; Steven P. Dow, University of California San Diego, La Jolla, California, USA, spdw@ucsd.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2021 Copyright held by the owner/author(s).

2573-0142/2021/10-ART352

<https://doi.org/10.1145/3479496>

to this design knowledge, designers may still struggle to find examples they can relate to their own ideas. Online repositories allow creators to browse previous examples of creative work by theme, keywords, or domain-specific dimensions [1, 6, 23, 54, 55]. Research tools have emerged to help designers visually explore example galleries and creatively recombine features into new designs [16, 40, 42, 48, 54].

Example repositories have not generally been organized to highlight the commonalities and differences between examples and someone's newly proposed idea. Recommending examples based on one's newly proposed idea could help creators find inspiration that supports their conceptual exploration of a local design space. Researchers have explored the possibility of modeling and illustrating the conceptual distance between ideas to help designers compare and assess novelty [31, 59]. Identifying examples at an appropriate conceptual distance can trigger remote associations and distant memories and help designers generate more unique ideas [22]. Different strategies have emerged for leveraging conceptual distance during ideation, including showing examples from different content categories [50], using text analysis to find distal ideas [59], and supporting analogical reasoning [29]. While conceptual distance appears useful for selecting a particular example as potential inspiration [11], less is known about how people might use data about conceptual distance to explore a *collection* of examples.

To understand how conceptual distance might affect how creators navigate and curate prior ideas, we built a technology probe called IdeateRelate that presents existing solutions to a particular problem organized by the conceptual distance to a newly submitted idea. IdeateRelate performs Natural Language Processing (NLP) analysis on a user's idea in real-time and creates an adaptive visualization that shows prior ideas distributed around a user's idea. We populated IdeateRelate with over 600 project ideas from an open-source COVID-19 volunteer-matching website and developed computational approaches using NLP to calculate the conceptual distance between all ideas and the user's idea. The visualization highlights the conceptual distance between existing ideas and proposed ideas by using spatial arrangement and color to highlight relative differences. (Fig. 2).

We conducted an online experiment (N=45 participants) to evaluate the benefits and tradeoffs of the IdeateRelate system (Viz condition) compared to a simple list view (No-Viz condition). We sought to understand (1) How do creators leverage a visualization of conceptual distance to explore examples compared to a list? (2) How do people curate examples to help them improve their own ideas differently when using a visualization versus a list of ideas? As a study procedure, participants first generated their own idea for solving a COVID-related challenge. Then participants explored all 637 ideas — either using IdeateRelate (Viz condition) or a list view (No-Viz condition) — and curated at least seven ideas to help them refine their own proposals. Finally, participants generated a final concept while viewing the curated subset of examples. Using NLP approaches, we measured the relative similarity of the examples participants explored and curated. We also measured a shift in novelty by computing the change in similarity between participants' initial and final ideas compared to all the examples. Expert judges also rated the quality of participants' initial and final ideas.

Our study results indicate that information about conceptual distance is more salient to users when using the IdeateRelate visualization compare to a simple table. Viz participants leveraged both category and semantic distance data, in different ways, to hone in on primarily *similar* examples. Viz participants also adopted more language from the examples they curated. However, this increased influence of the examples did not mean users simply emulated them; we found no negative effect on the overall novelty. This paper's technology probe and study provides preliminary evidence for how we might generate data on conceptual distance to help designers find, relate, and get inspiration from a large repository of design examples. We discuss design implications for creating perceptual cues

for conceptual distance and supporting interactions around large example repositories throughout a design process.

This paper makes the following contributions:

- We offer a novel way to visualize prior examples of creative work by highlighting the conceptual distance, through indicators of semantic and categorical similarity, in relation to a new idea.
- We conduct a comparative evaluation that explores how users curate examples using both semantic distance and categories in relation to their own ideas. IdeateRelate users generally looked for inspiration among more similar examples, and infused more language from the examples without sacrificing the novelty, compared to using a list of examples.
- We discuss the design implications for leveraging perceptual cues of conceptual distance and for collectively interpreting examples to support idea generation and iteration.

2 RELATED WORK

To motivate our research, we review the literature on the roles of examples in design, tools to support exploration of large-scale collections of design examples, and the importance of conceptual distance for creativity.

2.1 The role of examples in design

Prior work has offered mixed evidence of how examples affect people during creative activities [33, 61]. Examples can serve as design stimuli that activate distant memory to provoke divergent thinking, metaphors, and analogical reasoning [29, 49]. On the other hand, designers may simply borrow from examples by summarizing topics and borrowing low-level features [24]. Examples can instigate *design fixation*, where people constrain their thinking around particular features or functions [38, 49]. Multiple examples can afford people to explore existing ideas without the cost of implementation and provide representative exemplars and solutions [20].

To understand how people might navigate and gather information from a large repository of examples, we draw from Pirolli et al.'s *information foraging theory* [51], "an approach to understanding how strategies and technologies for information seeking, gathering, and consumption are adapted to the flux of information in the environment." Information foraging theory emphasizes the interaction between the perceptual environment and the discovery process. Chi et al. (2001), for example, used the theory to describe how exploring the Internet is a relationship between the expected value of following a link and the cost of performing actions [17]. Web pages give off an 'information scent' that people use to sense the value of a web link based on prior knowledge and perceptual cues within the navigation environment [52]. To explore the relationship between information foraging theory and design inspiration, we investigate how creators use perceptual cues to navigate and curate examples.

2.2 Exploring examples in creativity-support tools

Researchers have invented a variety of tools to help designers get inspiration [39, 48, 55, 57]. For example, CrowdBoard [4] and IdeaGens [10] support ways for facilitators and observers to provide additional information and examples that provide inspiration during ideation [41, 60]. D.Tour allows users to find relevant web design examples by styles visually and verbally [55]. Paragon uses an interactive example gallery to compliment design feedback [39]. Research on online design communities also leveraged existing data about human ratings, such as comments or applause, to help scout examples of good quality [2, 3, 36, 44].

Digital platforms have made vast numbers of existing ideas accessible and developed to help find examples more easily [1, 23, 37, 40]. A number of research tools have been created to help designers explore prior examples to create their designs [35, 42, 48, 54]. Most example exploration systems leverage categories, which is central to human cognition [8, 27, 56], to organize examples by common topics or characteristics [65]. Contemporary online galleries, such as Behance[1] and Dribbble[23], categorize examples by the type of designs (i.e., graphic, interaction, etc.). Behance also supports search by color composition.

The HCI research community has explored how to help users find examples based on the deeper structure of a design, such as locating websites and graphic designs by design styles and page layout [42, 55]. RecipeScape attempts to model the salient features of cooking instructions and visualize the connections between those features [16]. Research on crowdsourcing has demonstrated how crowds can generate categories, tags, and other forms of metadata on large sets of examples [5, 18, 30, 67]. This ability to assign any kind of metadata to an example repository rejuvenates questions about how to best support exploration. In this paper, we offer a technology probe that creates metadata to help designers explore examples in relation to their own ideas, and we investigate the efficacy of this designer-centered way of presenting examples.

2.3 Modelling conceptual distance in creativity-support tools

Conceptual distance has long been a subject of research for creativity-support systems [11, 58, 66]. The relative distance of two ideas — whether semantically near or far — can affect creativity [12, 13, 26]. Research indicates that semantically-far examples generally offer more diverse perspectives and can influence designers to create more novel ideas, while semantically-near examples serve more to help designers move towards more useful ideas [31]. However, when examples are too conceptually distant, designer may struggle to apply insights to a current problem; examples may need to be somewhat relatable in order to tap into cognitive processes of associative memory [11, 14, 15]. Hence, the ease of combining ideas depends on how people relate two ideas and whether they see a connection. How people explore ideas — whether they explore more semantically near or far examples — could provide insight on how the creative mind works.

Research has demonstrated the possibility to automatically mine the similarity between examples, particularly in domains of visual design [42, 48, 55]. In text-based domains, where ideas are captured as written descriptions, researchers have generally relied on crowdsourcing techniques to collect data on the similarity of examples [30, 47, 59]. For instance, Siangliulue et al. infers the affinities between ideas based on how users position ideas on a 2D pane [59]; Giroto et al. explored how to collect idea comparison data as a peripheral task alongside an ideation task [30]; Mackeprang et al. allowed users to add similarity data about their own proposals after producing ideas [47]. Researchers have also attempted to integrate these data collection tasks into users' natural workflow, such as sketching concept maps when listening to lectures or tagging existing examples before proposing new designs [45, 59], or clustering similar ideas before getting recommendations [53, 62, 63]. These human approaches to capturing similarity data highlight their importance in creative contexts; however, crowdsourcing this data can be costly and time-consuming, especially when dealing with a large repository of examples [18].

Our work explores how we might automatically produce similarity data based on text-based computational approaches. This paper explores how we might model and visualize two forms of conceptual distance. First, we model semantic similarity between any two ideas in our example repository and produce an affinity matrix that defines how close or far all examples are to each other. Second, we leverage human-produced categories to determine if two examples are in the same or different categories. While semantic similarity is a continuous variable between 0 and 1, category is a nominal variable.




Fig. 1. IdeateRelate users (1) read a problem statement about a brainstorming challenge, (2) propose an initial solution, (3) explore examples in relation to their own idea while selecting and annotating them for later consideration, and (4) review the curated examples with highlights and notes and iterate.

Our work not only models the conceptual distance between existing examples, we apply these same methods in real-time to new ideas offered by users. This allows us to visualize all the existing examples in relationship to a designer's new idea. The empirical work in this paper investigates whether this visualization approach improves how people navigate and apply insights from examples, compared to simply viewing examples in a list.

3 SYSTEM DESIGN AND IMPLEMENTATION

IdeateRelate is a web-based system that helps users visualize and reflect on prior design examples and iterate on their own ideas. This section introduces the IdeateRelate user experience and system architecture.

3.1 User Experience for IdeateRelate

The activity is divided into four main stages: **1. Problem brief.** Users are introduced to the problem statement and the ideation task. For our exploratory study, we asked participants to brainstorm solutions to COVID-19-related challenges because of the high level of familiarity in the general public and the wide availability of preexisting ideas for solutions. **2. Initial ideation.** Users submit an initial idea for a potential solution, including a name and description, in different text boxes. They are also asked to pick a category from a drop-down menu that best represents their idea. **3. Idea exploration.** Users then move onto the idea exploration interface (either the visualization or a simple list) and are instructed to curate ideas that might help them iterate or change their own ideas. IdeateRelate users can hover on the dots in the visualization and/or in the list to read ideas. They can also click on any idea to see more information, to take notes, to annotate words within the idea description, and to mark them as favorites. Users can see instructions at the top (Fig. 2a) and can minimize them to provide more space on screen for the data visualization below. **4. Idea iteration.** Finally, in the idea iteration phase, users are again given access to the list of curated ideas, together with the annotations they created using the main interface, so that users can integrate the information from their notes into their new idea. Users again write a new idea in a text box and select an idea category from a drop-down menu of the categories.

3.2 Visualization Design

The main visualization illustrates a distribution of ideas in relation to the users initial idea as a dot plot where each dot represents an idea from our database of preexisting ideas (Fig. 2b). Two features of similarity are highlighted in this case. (1) The visualization arranges ideas by their semantic similarity to the initial proposal submitted by the user, i.e., with more similar ideas further to the

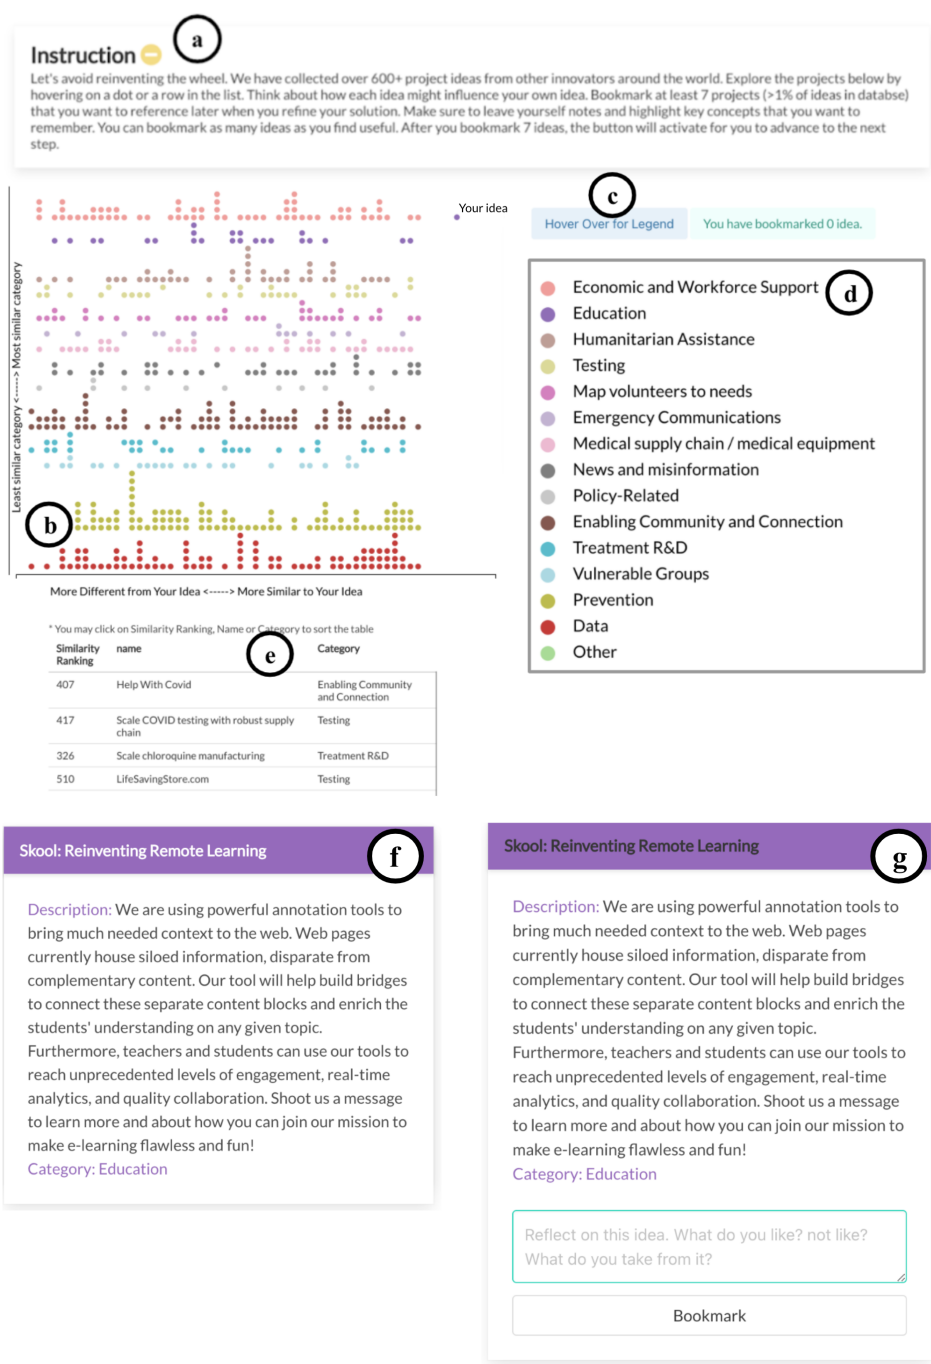


Fig. 2. The IdeateRelate main interface, including (a) the task instructions, (b) the similarity score visualization with colored categories, (c) the legend tab that produces (d) the legend of categories and colors, (e) the table containing idea name and category, (f) the “idea card” section that displays an idea’s card upon hovering and (g) a idea card upon clicking.

right of the graph and less similar ideas to the left. (2) The ideas are also color coded based on which of the 14 idea categories they belong to. The order of categories was based on the average similarity of the examples from each category. More similar categories are placed further to the top of the graph closer to the users' idea and less similar categories to the bottom. By hovering over the legend button (Fig. 2c), users can see how colors map onto different categories (Fig. 2d). A complimentary table below the visualization contains similarity rankings, idea names, and categories. The ideas are placed into a random order and can be sorted by clicking on each column head (Fig. 2e).

On the right of the visualization is the "idea card" area where users can read the ideas (Fig. 2f). The idea card includes an idea title, an idea description, and an idea category, which users can access by hovering over an idea in the visualization. If users click on a dot or row, the corresponding idea card will then change to 'sensemaking mode', where users can highlight words in the example descriptions, write down notes in a text box, and save the idea for later inspiration (Fig. 2g). Once saved, the idea is outlined in black in the visualization so that users can easily see which ideas they have already saved.

3.3 Example database

We collected existing ideas for our example database from HelpWithCovid.com [19], an open-source website described as "a clearing house for COVID-19 related projects and volunteers". As of June 2020, it had accumulated 700+ project ideas (Fig. 2f. shows an example project idea). To prepare the repository, our research team removed ideas that were solely recruiting advertisements, ideators' personal information, non-English, or those with very short descriptions. The remaining dataset included descriptions ranging between 10 to 359 words, with a median of 65 words. The research team leveraged the predefined project categories from the COVID-19 Innovation Hub [28] to manually assign a category to each one the ideas (see Appendix 1. for definitions). As a result of cleaning and categorizing the data, we ended up with 637 COVID-related ideas organized into 14 categories: Data, Economic and workforce support, Education, Emergency communications, Enabling community and connection, Humanitarian assistance, Map volunteers to needs, Medical supply chain / medical equipment, News and misinformation, Policy-related, Prevention, Testing, Treatment R&D, and Vulnerable groups.

3.4 System implementation

The front end of the IdeateRelate web application was coded using Javascript and HTML/CSS. The data visualizations and data interactions were created using the D3 Javascript library. All of the user input data were collected through the front-end interface, converted into JSON format, and stored in a MySQL database using API calls. The back end of IdeateRelate was built using Flask, a web framework written in Python. It was hosted using NGINX and Gunicorn. The semantic similarity between ideas is computed using the spaCy Transformer and sklearn Python libraries. User data were stored in a MySQL database using the Flask-SQLAlchemy extension to model the database structure and commit data from the front end to the database.

To shorten computation time when processing the user's idea, we pre-processed data about the existing ideas in our repository. When generating semantic similarity between any two ideas, we included both the idea's title and description. We process these texts by removing stop words and unrecognizable words. For each idea, we use existing NLP tools to create a word-embedding vector that numerically represents the example idea in context. We use python NLTK [7] to process texts (remove stopwords, lemmatize texts, etc.), and spaCy [34] as a pipeline for pre-trained transformer BERT(Bidirectional Encoder Representations from Transformers) [21] to extract 768 features of word vectors in the text. A SpaCy transformer tokenizes words, transforms those word tokens into token vectors, and outputs the last hidden state as the feature vector. We sum up the feature

vectors of idea description keywords to represent the idea. We then stored these data in a CSV file on our server so that they could quickly be used and accessed when processing the user's idea. This pre-processing step allows us to skip the steps where example texts are transformed into feature vectors for semantic distance calculation when comparing the examples to the user's initial idea.

3.4.1 Visualizing examples in relation to a new idea. Once the user has submitted their idea, we use the same set of NLP tools to get word vectors for 768 features from the idea. We then use sklearn's [9] spatial distance processing tool to calculate the cosine distances between the word-embedding vector for the user's idea and each of the word-embedding vectors for the examples in the repository. We invert the cosine distance metric to compute a similarity score between 0 and 1 that represents how near or far the user's idea is to all examples in the repository.

To prepare for the data visualization phase, we take the cosine similarity score for individual examples and rank them within the same category. We also calculate an average cosine similarity score for ideas within the same category. For each user's initial idea, we assign calculated similarity scores to each example based on the similarity data that was just generated. The algorithm ensures that examples with the same x coordinate, i.e. examples in the same column, are assigned adjacent y coordinates so that multiple ideas do not occupy the same X-Y coordinate.

The coordinate data are then stored in a CSV file on the server for use in the main visualization. The data describing each idea's name, description, category, similarity score, and coordinates are accessed through the CSV file created in the initial ideation phase. The main visualization is a dot plot created by simply plotting the coordinates generated in the ideation phase and stored in the CSV file. Below the visualization, a table displays the same ideas as a simple list with the name, description and category pulled directly from the CSV file.

3.4.2 Interacting with examples. Both the visualization and table support two interactions, hovering and clicking. When a hover event is triggered — either by hovering over a dot in the visualization or a row in the table — the idea card is made visible and filled with the corresponding idea name, description, and category from the CSV data. When a click event is triggered on a dot or row, the persistent annotation card is displayed containing the same information, as well as a text box for user-generated notes and a *bookmark* button. The user's annotations are stored locally until the bookmark button is clicked, which triggers an event that converts the annotation data into JSON format. Once the data are in JSON format, they can be injected into an API call to our back end, which formats and commits the data to the MySQL database. During the "idea iteration" phase, the curated examples are displayed again along with the user's annotations and notes.

3.5 Validating the similarity algorithm using human ratings

To validate whether the algorithm's calculations of semantic similarity align with human assessments, we generated a triplet comparison task where human judges select the two most similar ideas out of the three displayed [64]. Our judges were presented with a seed idea from early piloting studies and two randomly-ordered example ideas deemed to be more similar to the seed idea.

To avoid the potential effects of comparing across categories, each triplet comparison task only sampled ideas from the same category. We selected examples with different computed similarities (one from each quartile) compared to the seed idea. These four examples created six different pairwise combinations; judges compared each seed idea with all six different pairs in a random order. Overall we collected ratings from two judges on eight seed ideas with six triplet comparisons for each seed idea, for a total of 96 ratings.

Overall, 72.9% of the human judgements matched what the algorithm calculated on that same triplet comparison. That is, humans generally judged the two computationally 'nearest' ideas to be the most similar to each other. Human raters agreed on 62.5% of triplet ratings. This is generally

high agreement; for example, on a similar judgement task, Chan et al. (2015) found a correlation of 51% between human raters and an algorithm.

4 EVALUATION STUDY

In order to understand how do people leverage the visualization of conceptual distance to explore and curate examples, we conducted a between-subjects online experiment. Our primary research question explores whether the presence or absence of the relational visualization affects how user explore and curate examples, and to what extent this influences subsequent ideation. In addition, we were interested to what extent users paid attention to categories and the semantic distance data when exploring and curating examples. We hypothesize that IdeateRelate visualization will emphasize conceptual distance, and make this data more salient when exploring and getting inspiration from examples. It remains an open question whether people will leverage this information and to what extent they prefer relative semantic distance or categorical information and how this might influence the downstream ideation process.

4.1 Methodology

4.1.1 Study design. Participants were randomly assigned to one of two conditions: Half the participants had access to the IdeateRelate tool with the visualization and also a table of examples listed below it (Viz). The other half of participants only saw the table of examples showing the idea name and category (No-Viz).

4.1.2 Procedure. Users first read the instructions and propose an initial idea in the idea initiation phase. Then they see the main interface where they interact with (read, curate and annotate) examples to explore what has been done previously in the design space; users either see the Viz condition or the No-Viz condition, as outlined above. Finally, users see their curated examples and the notes/annotations they created in the previous stage, and they are instructed to use the examples to improve or pivot on their idea. At the end of study, participants are directed to an external survey to answer questions related to their experience. After the study, to assess how the IdeateRelate tool affected ideation behaviors, we recruited three teaching assistants from a project-based design course to independently rate the users' initial and final ideas based on novelty (1 not novel to 7 very novel) and usefulness (1 not useful to 7 very useful).

4.1.3 Participants. We had a total of 45 participants (27 female) after excluding four from the analysis for not interacting with the system. Participants were recruited through a university email list, and ranged in age from 18 to 64. Participants were paid at \$15/hour for an average of forty minutes of time. Sixteen participants reported some prior experience in design. The majority of participants (N=44) reported that they were at least moderately familiar with COVID-19 challenges.

4.1.4 Analysis. We began by examining how participants navigated the examples during the idea exploration stage. We then further analyze the traits of examples that participants explored and curated. We also calculated the semantic distance between the examples and users' ideas using similar NLP approaches (by calculating cosine distance between the word vectors of a user's idea and the existing examples). We used the last hidden state as the feature vector for semantic similarity calculation of texts and measured cosine similarity between sums of the feature vectors of the extracted noun phrases. This allowed us to understand to what extent users adopted the language from an example into their revised idea.




Fig. 3. Viz participants tended to explore examples more similar to their own whereas No-Viz participants explored more randomly.

4.2 Results

This study investigates how people explore, curate and take inspiration from examples when navigating them in a list (No-Viz) compared to using a visualization (Viz condition) that visually arranges the examples in relation to the semantic and categorical similarity to a newly-submitted idea, as described above.

4.2.1 Exploration: Viz participants focused on finding examples similar to their own. Participants in the Viz condition viewed $M = 64$ ($SD = 80.1$) examples compared to $M = 44$ ($SD = 33.2$) in the No-Viz condition. In terms of interacting with examples, Viz participants, on average, hovered over 42 ideas ($SD = 24$) in the visualization to read and 50 rows ($SD = 106$) in the table. Some of these ideas were viewed in both ways, so the user got exposed to 44 unique ideas on average. The average time for reading each idea card took 5.1s in Viz condition, and 3.9s in the No-Viz condition.

Participants in the Viz condition explored more ideas that were closer to their initial idea, 65% of were semantically closer. The average similarity ranking of explored examples in the Viz condition ($M = 249$, $SD = 185$) was higher than the No-Viz condition ($M = 299$, $SD = 184$) ($t(43) = 4.57$, $p < 0.01$) (see Fig. 3).

From the qualitative survey filled out at the end of the study, many participants expressed a desire to improve their idea as the most common rationale for why they curated particular examples.

I looked for ideas that can be subcomponents of the idea that I was proposing and incorporated some of their elements to the final result. (P6, No-Viz)

Participants seemed to evaluate examples based on how they might associate them with their own. Participants reported that the more similar examples gave them specific directions for improving their own idea, while far examples gave them a broader understanding of the the topic space.

I looked through the ideas most similar to mine and ones that were most dissimilar, and I also explored random dots in each color category. (P5, Viz)

Other Viz participants talked about how the similarity cues helped them get a broader sense of the problem space:

... I explored topics outside of my original idea to gain a sense of the problem space people are trying to tackle regarding covid. (P3, Viz)

Initial proposal (P9):

A lot of students are lonely only attending zoom classes and want to see and meet others and their professor **in person**. If the weather is good enough then classes could optionally meet outside and there could be seating arrangements **marked outside** and a projector could project slides on the side of a building.

Curated examples (P9):

Example#586: Our **virtual event** space allows people to join conversations the same way they would at a networking event. Networking is a key ingredient to developing solutions to problems like COVID-19. Currently, networking is hindered by COVID-19 and our current solutions are not optimal. Our product filters audio based on your position in virtual space and the conversation you are currently in.

Example#104: The **Social Hour** is a daily 1-hour session where we can spend an hour to put our minds off of Covid-19 and use this chance to learn intellectual/personal topics that we didn't know before.

Example#222: A dating service for friends. Great for **matching** seniors, overwhelmed parents, and anyone feeling particularly isolated with others who want to chat, find an activity or generally reach out.

Example#45: A really quick and easy Mural board for anyone wanting to state problems and **contribute ideas** in this new space.

Iterated proposal (P9):

There could be a map where it shows which areas have the most exposure to covid and which areas are exposure-free zones. Depending on the zone you're in and your personal preferences, you can choose to attend optional **in-person** social and educational **events** hosted by your school. The school can collect data from these events to inform whether their strategy is safe and to inform the map of covid exposure. There could be **markings outside** that show where students can sit and which way they have to face (kind of like what is at a grocery store). For students who are less comfortable with this or who may be exposed to covid, there can be a **virtual version of the event**, but with options for people to have individual conversations like they would have in person. After the event there can be a **social hour** where people can choose from a list of (non-covid or covid) topics and be **matched** with people to discuss the topics with. Afterwards students can **submit ideas** about how to make the events better or safer in the future and these ideas can get reviewed by a panel of experts in safety and design.

Fig. 4. A participant in the Viz Condition took inspiration from curated examples to build on their initial ideas. The iterated proposal contained components like "in person" and "markings outside" from the initial proposal and took inspirations like "virtual event", "social hour", "matching (people)" and "submit ideas" from the examples.

In both conditions, participants mentioned that 'category' was a major cue for them to look for relevant examples.

I first looked at what category the innovation was in, and filtered through those. (P25, Viz)
I looked under the category I started with, which helped to narrow down the direction I wanted to go. Then, if I saw something that would help, I tried to use that to match up with my idea. (P38, No-Viz)

Some participants in the No-Viz condition initially grappled with the meaning of the similarity data. As P43 mentioned:

I looked at the similarity ranking and at first I didn't know what that meant, but as I spent more time on it I thought that the lower the ranking, the more similar the idea was to what I had written down. (P43, No-Viz)

One participant specifically mentioned that they "initially looked at the ranking of similarity but it was not helpful". Instead, many of those participants used the sort-by-category mechanism or reflected on how "interesting titles would catch my eye (P47, No-Viz)". The No-Viz participants seemed to leverage titles and categories, more than the similarity ranking, to decide what to read in more depth. As a participant mentioned, "I clicked on titles that seemed interesting and on topics I wanted to know more about (P29, No-Viz)".

4.2.2 Curation: Viz participants curated more examples based on conceptual distance. Overall, participants selected to curate 343 ideas including 239 unique ideas. The top selected idea was selected by six participants. 293 ideas (46.1% of ideas) were not curated at all among participants. Viz participants were more likely to select unique examples (146 unique ideas) compared to No-Viz participants who viewed the same default random ranking of the ideas (128 unique ideas).

Not only did participants evaluate alignment between examples and their own ideas, they would also assess the quality of examples. "I decided to bookmark ideas that shared either (1) an interesting/new means of applying my original [idea] or (2) aligned with the values proposed in my

original. (P3, Viz)" Participants also looked for novel aspects of ideas that could be integrated into their own proposal, as one person said:

I then bookmarked ideas that seemed novel and were at least slightly related to the idea that I had initially proposed. (P25, Viz)

Participants in the Viz condition used both semantic similarity and categorical similarity as cues to curate a diverse set of examples. Aligning with what they explored, more Viz participants selected ideas from the "similar" side of visualization (the top 50% percentile ranking of ideas on semantic similarity to user proposals), where it was perceptually closer to their own ideas. Participants in the Viz condition curated primarily semantically-near examples, while only four did in the No-Viz condition. Individual participants tended to curate examples of diverse categories in both conditions. In contrast to the Viz condition where participants could explore examples in the same categories more seamlessly, the No-Viz participants could sort by category but then could only explore them in a more randomized order.

4.2.3 Iteration: Viz participants adopted more language from curated examples. Participants in both conditions reported integrating 3 to 4 examples on average from their curated list into their final ideas ($M = 3.7$, $SD = 2.1$). Even though there was no self-reported difference in the average number of examples borrowed between conditions, we still observed that participants in the two conditions used the examples differently to iterate on their ideas. The iterated ideas in the Viz condition became more semantically similar to the curated examples on average than those in the No-Viz condition ($t\text{-test} = 2.09$, $p = 0.037$). In other words, the ideas adopted more language from the curated set of ideas, a sign of creative contagion. Fig. 4 illustrates how one Viz participant adopted language from the curated examples into their final idea iteration.

Across all participants, there was no significant difference between ratings of the initial and iterated ideas. Overall, participants who scored lower were more likely to improve their ideas after interacting with examples than those who scored higher ($cor = -0.517$, $p < 0.01$) indicating a ceiling effect.

5 DISCUSSION

We created the IdeateRelate system as a design probe to understand the value of surfacing aspects of conceptual distance. Our evaluation shows that participants using this visualization explored and curated examples that were more similar to their own compared to using a list with simple sorting mechanisms. Viz participants integrated more language and features from their curated examples, although not at the cost of reducing the novelty, compared to No-Viz participants.

Participants selected ideas for different reasons. Some participants tend to look for different features to combine in semantically-near examples, and some participants tend to look for commonalities among semantically-far examples. Participants reported a range of behaviors during curation, including annotating the important features of examples, pointing out the connections between ideas and their proposals, and providing constructive feedback or critical thinking on the examples.

5.1 Surfacing cues for conceptual distance to support creative recombination

The literature often portrays creativity as the recombination of existing knowledge into new patterns [32, 68]. Our study results show that participants look for "relevant and divergent" examples to combine with their initial proposals. They either look for different features in semantically-near examples or common themes in semantically-far examples. A few different factors may play into the thought process of seeking 'relevant' and 'divergent' examples. The first factor is to see connections between the examples and the user's initial proposal. The visual cues on similarity help participants

look for associations with their own proposals and, in turn, this led them to more easily combine new concepts into their idea. The second factor is the ability to combine features in a novel way by providing cues on both semantic and categorical similarity. Having multiple ways to calibrate conceptual similarity seems to help people better estimate the difficulty and novelty of creative recombination.

Prior literature has grappled with how far or near examples should be to influence creative thinking [31, 58]. Exploring related categories of ideas in depth seems to improve the originality of a new idea [50]. semantically-far examples activate remote association so that people could creatively combine elements from the examples with their own ideas [15]. However, examples that are too semantically far from the target domain may place a cognitive burden that reduces creative production [11, 15]. Participants in our study tended to curate ideas from various categories. They seem to search for examples that combine similar word vectors (semantic similarity) with divergent categories (categorical diversity). We observed participants evolving their ideas, adopting language from the examples, but we did not measure a significant change in the novelty of their final ideas compared to their original idea. Perhaps allowing people to interact with the visualization through multiple rounds of iteration will lead to more significant effects. A longer engagement might also reveal exploration patterns like searching for far inspiration first and then honing in on more nuanced differences with nearby examples, or vice versa.

5.2 Supporting interactions for collective interpretation of examples

In our evaluation study, most participants who used the visualization to explore examples edited the ideas more than those who predominantly used the list view. One reason could be those participants who mostly rely on using the built-in browser functionality to search key terms to navigate are already fixated on the initial ideas by searching relevant information. The search query helped them focus on relevant information or common characteristics among examples rather than different features. Future designs may incorporate user needs on searching key terms but provide an expansive list of words to search.

Each participant had their own agenda, as very few ideas were selected by multiple individuals. This indicates that participants selected examples in relation to their own idea; there was no zeitgeist or best ideas that made it into everyone's curated stacks. If more people were using the system, we would anticipate more people would curate the same examples for different reasons. Our study also found that people took very different types of notes on the examples they curated. Some annotations provided constructive feedback or reemphasized the examples' interesting points. Could these diverse annotations useful for other creators? Sharing previously generated annotations with future creators could give access to other considerations of dimensions of the problem space that were not previously considered.

5.3 Limitations and future directions

The current version of IdeateRelate showcases ideas from a COVID-related database, but it can incorporate any text-based examples. Different example databases will have different categories or hierarchies to organize examples, i.e., the examples may be labeled differently or they may belong to more than one category. We may also enhance the system in the future by offering other ways to find examples. For instance, we can make semantic associations more explicit by highlighting similar words or by crowdsourcing additional metadata according to an expert-derived schema [67].

Information foraging theory would suggest that the visualization decreased the cost of search by surfacing better informational cues compared to scrolling through idea titles in a list [17, 51, 52]). Considering ideas were never read by participants, filtering mechanisms may be offered on

visualization to allow participants to conduct localized search on a subset of a larger repository of ideas. Hiding some of less relevant ideas can also save space for interacting with more useful examples.

Our study found a ceiling effect where we did not observe a significant increase in the novelty of ideas produced after one iteration. One future direction is to provide adaptive visualizations of examples after idea iteration so that the inspirations evolve along with an individual's ideas.

6 CONCLUSION

This paper explores how to visualize existing ideas within a particular problem space organized by conceptual distance in relation to a newly submitted idea. We created the IdeateRelate system as a design probe to study how ideators explore large repositories of examples with two types of conceptual distance – semantic similarity and conceptual similarity – centered around their own ideas. We conducted an online experiment with forty-five participants with an ideation task, where they either viewed examples in the visualization or in a simple list. The IdeateRelate visualization allows users to explore and curate ideas that are more similar to their own ideas, and borrow more perspectives into their own idea iterations than participants who used a list view. Future repositories may offer multi-faceted methods of navigating conceptual distances as perceptual cues for people exploring examples.

ACKNOWLEDGMENTS

We thank the volunteers and contributors at helpwithcovid.com for making data open and accessible. We thank Prof. Judith Fan, Prof. Scott Klemmer, Prof. Julian McAuley, and Prof. Haijun Xia for providing insightful comments and suggestions at various stages of the research project, and undergraduate research assistants Tiffany Manuel, Vincent Chu, and Michael Wroblewski for their contribution to the early work. The National Science Foundation funded this work under grant #1122206 and #1821618.

REFERENCES

- [1] Adobe. 2020. *Behance*. Retrieved October 15, 2020 from <https://www.behance.net/>
- [2] Faez Ahmed and Mark Fuge. 2017. Capturing winning ideas in online design communities. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 1675–1687.
- [3] Faez Ahmed, Mark Fuge, and Lev D Gorbunov. 2016. Discovering diverse, high quality design ideas from a large corpus. In *ASME 2016 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. American Society of Mechanical Engineers Digital Collection.
- [4] Salvatore Andolina, Hendrik Schneider, Joel Chan, Khalil Klouche, Giulio Jacucci, and Steven Dow. 2017. Crowdbord: augmenting in-person idea generation with real-time crowds. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. 106–118.
- [5] Paul André, Aniket Kittur, and Steven P Dow. 2014. Crowd synthesis: Extracting categories and clusters from complex data. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. 989–998.
- [6] Brian P. Bailey and Eric Horvitz. 2010. What's Your Idea? A Case Study of a Grassroots Innovation Pipeline within a Large Software Company. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (*CHI '10*). Association for Computing Machinery, New York, NY, USA, 2065–2074. <https://doi.org/10.1145/1753326.1753641>
- [7] Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc."
- [8] Lera Boroditsky. 2007. Comparison and the development of knowledge. *Cognition* 102, 1 (2007), 118–128.
- [9] Lars Buitinck, Gilles Louppe, Mathieu Blondel, Fabian Pedregosa, Andreas Mueller, Olivier Grisel, Vlad Niculae, Peter Prettenhofer, Alexandre Gramfort, Jaques Grobler, Robert Layton, Jake VanderPlas, Arnaud Joly, Brian Holt, and Gaël Varoquaux. 2013. API design for machine learning software: experiences from the scikit-learn project. In *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*. 108–122.

- [10] Joel Chan, Steven Dang, and Steven P Dow. 2016. IdeaGens: enabling expert facilitation of crowd brainstorming. In *Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion*. 13–16.
- [11] Joel Chan, Steven Dow, and Christian Schunn. 2014. Conceptual distance matters when building on others' ideas in crowd-collaborative innovation platforms. In *Proceedings of the companion publication of the 17th ACM conference on Computer supported cooperative work & social computing*. 141–144.
- [12] Joel Chan, Steven P Dow, and Christian D Schunn. 2018. Do the best design ideas (really) come from conceptually distant sources of inspiration? In *Engineering a Better Future*. Springer, Cham, 111–139.
- [13] Joel Chan, Katherine Fu, Christian Schunn, Jonathan Cagan, Kristin Wood, and Kenneth Kotovsky. 2011. On the benefits and pitfalls of analogies for innovative design: Ideation performance based on analogical distance, commonness, and modality of examples. (2011).
- [14] Joel Chan and Christian Schunn. 2015. The impact of analogies on creative concept generation: Lessons from an in vivo study in engineering design. *Cognitive Science* 39, 1 (2015), 126–155.
- [15] Joel Chan, Pao Siangliulue, Denisa Qori McDonald, Ruixue Liu, Reza Moradinezhad, Safa Aman, Erin T Solovey, Krzysztof Z Gajos, and Steven P Dow. 2017. Semantically far inspirations considered harmful? accounting for cognitive states in collaborative ideation. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. 93–105.
- [16] Minsuk Chang, Léonore V Guillain, Hyeungshik Jung, Vivian M Hare, Juho Kim, and Maneesh Agrawala. 2018. Recipescape: An interactive tool for analyzing cooking instructions at scale. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [17] Ed H. Chi, Peter Pirolli, and James Pitkow. 2000. The Scent of a Site: A System for Analyzing and Predicting Information Scent, Usage, and Usability of a Web Site. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (The Hague, The Netherlands) (CHI '00). Association for Computing Machinery, New York, NY, USA, 161–168. <https://doi.org/10.1145/332040.332423>
- [18] Lydia B. Chilton, Greg Little, Darren Edge, Daniel S. Weld, and James A. Landay. 2013. Cascade: Crowdsourcing Taxonomy Creation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) (CHI '13). Association for Computing Machinery, New York, NY, USA, 1999–2008. <https://doi.org/10.1145/2470654.2466265>
- [19] Help With Covid. 2020. *Help With COVID 19*. Retrieved July 20, 2020 from <https://helpwithcovid.com/>
- [20] Denise Dellarosa Cummins. 1992. Role of analogical reasoning in the induction of problem categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 18, 5 (1992), 1103.
- [21] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [22] Arne Dietrich. 2004. The cognitive neuroscience of creativity. *Psychonomic bulletin & review* 11, 6 (2004), 1011–1026.
- [23] Dribbble. 2020. *Dribbble*. Retrieved October 15, 2020 from <https://dribbble.com/>
- [24] Claudia Eckert and Martin Stacey. 1998. Fortune favours only the prepared mind: Why sources of inspiration are essential for continuing creativity. *Creativity and Innovation Management* 7, 1 (1998), 9–16.
- [25] Ronald A Finke, Thomas B Ward, and Steven M Smith. 1992. Creative cognition: Theory, research, and applications. (1992).
- [26] Katherine Fu, Joel Chan, Jonathan Cagan, Kenneth Kotovsky, Christian Schunn, and Kristin Wood. 2013. The meaning of “near” and “far”: the impact of structuring design databases and the effect of distance of analogy on design output. *Journal of Mechanical Design* 135, 2 (2013).
- [27] Dedre Gentner and Kenneth J Kurtz. 2005. Relational categories. (2005).
- [28] Global Innovation Exchange (GIE). 2020. *COVID-19 Categories*. Retrieved October 15, 2020 from <https://covid19innovationhub.org/categories>
- [29] Karni Gilon, Joel Chan, Felicia Y Ng, Hila Liifshitz-Assaf, Aniket Kittur, and Dafna Shahaf. 2018. Analogy mining for specific design needs. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [30] Victor Giroto, Erin Walker, and Winslow Burleson. 2017. The Effect of Peripheral Micro-Tasks on Crowd Ideation. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 1843–1854. <https://doi.org/10.1145/3025453.3025464>
- [31] Kosa Goucher-Lambert and Jonathan Cagan. 2019. Crowdsourcing inspiration: Using crowd generated inspirational stimuli to support designer ideation. *Design Studies* 61 (2019), 1–29.
- [32] Andrew Hargadon and Robert I Sutton. 1997. Technology brokering and innovation in a product development firm. *Administrative science quarterly* (1997), 716–749.
- [33] Scarlett R Herring, Chia-Chen Chang, Jesse Krantzler, and Brian P Bailey. 2009. Getting inspired! Understanding how and why examples are used in creative design practice. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 87–96.

- [34] Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. (2017). To appear.
- [35] Gaoping Huang and Alexander J. Quinn. 2017. BlueSky: Crowd-Powered Uniform Sampling of Idea Spaces. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition* (Singapore, Singapore) (C&C '17). Association for Computing Machinery, New York, NY, USA, 119–130. <https://doi.org/10.1145/3059454.3059481>
- [36] Yi-Ching Huang, Chun-I Wang, and Jane Hsu. 2013. Leveraging the crowd for creating wireframe-based exploration of mobile design pattern gallery. In *Proceedings of the companion publication of the 2013 international conference on Intelligent user interfaces companion*. 17–20.
- [37] Invision. 2020. *Muzli Design Inspiration*. Retrieved October 15, 2020 from <https://muz.li/>
- [38] David G Jansson and Steven M Smith. 1991. Design fixation. *Design studies* 12, 1 (1991), 3–11.
- [39] Hyeonsu B Kang, Gabriel Amoako, Neil Sengupta, and Steven P Dow. 2018. Paragon: An online gallery for enhancing design feedback with visual examples. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [40] Laura J Kornish and Karl T Ulrich. 2011. Opportunity spaces in innovation: Empirical analysis of large samples of ideas. *Management Science* 57, 1 (2011), 107–128.
- [41] Chinmay Kulkarni, Steven P Dow, and Scott R Klemmer. 2014. Early and repeated exposure to examples improves creative work. In *Design thinking research*. Springer, 49–62.
- [42] Ranjitha Kumar, Arvind Satyanarayan, Cesar Torres, Maxine Lim, Salman Ahmad, Scott R. Klemmer, and Jerry O. Talton. 2013. Webzeitgeist: Design Mining the Web. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) (CHI '13). Association for Computing Machinery, New York, NY, USA, 3083–3092. <https://doi.org/10.1145/2470654.2466420>
- [43] Keelin Leahy, Shanna R Daly, Seda McKilligan, and Colleen M Seifert. 2020. Design Fixation From Initial Examples: Provided Versus Self-Generated Ideas. *Journal of Mechanical Design* 142, 10 (2020).
- [44] Brian Lee, Savil Srivastava, Ranjitha Kumar, Ronen Brafman, and Scott R Klemmer. 2010. Designing with interactive example galleries. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 2257–2266.
- [45] Ching Liu, Juho Kim, and Hao-Chuan Wang. 2018. ConceptScape: Collaborative concept mapping for video learning. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [46] Brian J Lucas and Loran F Nordgren. 2015. People underestimate the value of persistence for creative performance. *Journal of Personality and Social Psychology* 109, 2 (2015), 232.
- [47] Maximilian Mackeprang, Abderrahmane Khayat, and Claudia Müller-Birn. 2018. Concept Validation during Collaborative Ideation and Its Effect on Ideation Outcome. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI EA '18). Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3170427.3188485>
- [48] J. Marks, B. Andalman, P. A. Beardsley, W. Freeman, S. Gibson, J. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. 1997. Design Galleries: A General Approach to Setting Parameters for Computer Graphics and Animation. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97)*. ACM Press/Addison-Wesley Publishing Co., USA, 389–400. <https://doi.org/10.1145/258734.258887>
- [49] Richard L Marsh, Joshua D Landau, and Jason L Hicks. 1996. How examples may (and may not) constrain creativity. *Memory & cognition* 24, 5 (1996), 669–680.
- [50] Bernard A Nijstad, Carsten KW De Dreu, Eric F Rietzschel, and Matthijs Baas. 2010. The dual pathway to creativity model: Creative ideation as a function of flexibility and persistence. *European Review of Social Psychology* 21, 1 (2010), 34–77.
- [51] Peter Pirolli and Stuart Card. 1999. Information foraging. *Psychological review* 106, 4 (1999), 643.
- [52] Peter Pirolli and Wai-Tat Fu. 2003. SNIF-ACT: A model of information foraging on the World Wide Web. In *International Conference on User Modeling*. Springer, 45–54.
- [53] Samuel Rhys Cox, Yunlong Wang, Ashraf Abdul, Christian von der Weth, and Brian Y Lim. 2021. Directed Diversity: Leveraging Language Embedding Distances for Collective Creativity in Crowd Ideation. *arXiv e-prints* (2021), arXiv–2101.
- [54] Christoph Riedl, Steffen Wagner, Jan Marco Leimeister, and Helmut Krcmar. 2010. Exploring large collections of ideas in collaborative settings through visualization. In *Proceedings of 20th Annual Workshop on Information Technologies and Systems*.
- [55] Daniel Ritchie, Ankita Arvind Kejriwal, and Scott R Klemmer. 2011. d. tour: Style-based exploration of design example galleries. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*. 165–174.
- [56] Eleanor Rosch, Carolyn B Mervis, Wayne D Gray, David M Johnson, and Penny Boyes-Braem. 1976. Basic objects in natural categories. *Cognitive psychology* 8, 3 (1976), 382–439.
- [57] Kathryn Shroyer, Terri Lovins, Jennifer Turns, Monica E Cardella, and Cynthia J Atman. 2018. Timescales and ideospace: An examination of idea generation in design practice. *Design Studies* 57 (2018), 9–36.

- [58] Pao Siangliulue, Kenneth C Arnold, Krzysztof Z Gajos, and Steven P Dow. 2015. Toward collaborative ideation at scale: Leveraging ideas from others to generate more creative and diverse ideas. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. 937–945.
- [59] Pao Siangliulue, Joel Chan, Steven P Dow, and Krzysztof Z Gajos. 2016. 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*. 609–624.
- [60] Pao Siangliulue, Joel Chan, Krzysztof Z Gajos, and Steven P Dow. 2015. Providing timely examples improves the quantity and quality of generated ideas. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*. 83–92.
- [61] Steven M Smith, Thomas B Ward, and Jay S Schumacher. 1993. Constraining effects of examples in a creative generation task. *Memory & cognition* 21, 6 (1993), 837–845.
- [62] Luka Stärk. [n.d.]. Semantic Similarity of Concepts for a Human-Centered Idea Recommendation Feature in the Clustering Application Orchard. ([n. d.]).
- [63] Maximilian Timo Stauss. [n.d.]. Conceptualization and Evaluation of Idea Similarities based on Semantic Enrichment & Knowledge Graphs. ([n. d.]).
- [64] Neil Stewart, Gordon DA Brown, and Nick Chater. 2005. Absolute identification by relative judgment. *Psychological review* 112, 4 (2005), 881.
- [65] Thomas B Ward. 1994. Structured imagination: The role of category structure in exemplar generation. *Cognitive psychology* 27, 1 (1994), 1–40.
- [66] Thomas B Ward. 2001. Creative cognition, conceptual combination, and the creative writing of Stephen R. Donaldson. *American Psychologist* 56, 4 (2001), 350.
- [67] Xiaotong Xu, Judith Fan, and Steven Dow. 2020. Schema and Metadata Guide the Collective Generation of Relevant and Diverse Work. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 8. 178–182.
- [68] Lixiu Yu and Jeffrey V Nickerson. 2011. Cooks or cobblers? Crowd creativity through combination. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1393–1402.

Received October 2020; revised April 2021; accepted July 2021

Table 1. Database categories that were used in idea brainstorming and main interface

Name	Definition
Data	Rapid aggregation, visualization, and dissemination of data about treatments, diagnostics, and up-to-date confirmed cases to inform the public and critical decision-makers.
Economic and Workforce Support	New and continuing job training, innovative financial support methods, employment enabling technologies.
Education	Uninterrupted access to education to those impacted by COVID-19 including distance learning, self-learning without parental or guardian support, and others.
Enabling Community and Connection	Support during and after the crisis for communities and healthcare workers such as from isolation for trauma and depression. Peer-to-peer support, information and other social platforms that allow people to stay engaged and connected with their communities while social distancing.
Humanitarian Assistance	Innovations that support critical delivery of water, food, health, hygiene, and other supplies as well as financial support (i.e. relief funds) during and after a COVID-19 outbreak.
Map volunteers to needs / humanitarian assistance	Innovations that support critical delivery of water, food, health, hygiene, and other supplies as well as financial support (i.e. relief funds) during and after a COVID-19 outbreak.
Medical supply chain / medical equipment	Transportation for COVID-19 impacted supply chains or individuals such as safe transport of specimens and medical equipment or mobility-friendly triaging or treatment units, particularly for rural areas.
New and misinformation	Help people distinguish news from misinformation
Policy-Related	Policies or guidelines to contribute to public well-being and ensure rapid, equitable, inclusive and human-centered responses.
Prevention	Innovations that help prevent the spread of COVID-19, including healthcare worker protection equipment (PPE), healthcare worker training, sterilization, and water, sanitation and hygiene (WASH) innovations. Examples include self-sanitizing PPE or new applications for fast, robust, inexpensive sterilization for reuse.
Testing	Focused efforts for rapid testing and detection with reduced health workforce needs, for both molecular and immunological tests.
Treatment R&D	Innovations that can help treat COVID-19 and related complications, including medical equipment (such as the design or retrofitting of ventilators to enable new units at scale quickly) and emergency treatment units. It also includes healthcare practitioner training and technology assistance to enable maximum use of existing and new systems and approaches (such as virtual reality capabilities for rapid immersion, tele-supervised training or novel approaches to practitioner force multiplication).
Vulnerable Groups	Protection of individuals or marginalized/vulnerable groups at greater risk due to COVID-19 or other factors (e.g. youth, women, refugees, migrants, elderly, etc.).