Dimensional Reasoning and Research Design Spaces

Stephen MacNeil, Johanna Okerlund, Celine Latulipe
Human-Computer Interaction Lab
University of North Carolina at Charlotte, Charlotte, NC, USA
{smacnei2, jokerlun, clatulip}@uncc.edu

ABSTRACT
We present an exploration into the use of dimensional reasoning and the creation of research design spaces. We hypothesized that researchers, engaged in open-ended creative problem solving, could adapt the methods of design thinking and design spaces to create dimensionalized research design spaces. To investigate how researchers might explicitly engage in such dimensionalization processes, we created and studied (n=5, n=5) an interactive web-based system for the creation of research design spaces. Our results showed that a 'dimensions-first' approach was difficult for researchers to work with. We then created and studied (n=11) a prototype 'examples-first' approach. Our results suggest that the ability of researchers to explicitly dimensionalize their research areas is quite varied and that significant scaffolding is required to help researchers reason in this way.

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords
dimensional reasoning; meta-research; design space; research; ideation; sense-making.

INTRODUCTION
Conducting high-level academic research is a complex, creative, and open-ended task. It can be challenging for researchers to effectively frame new work within the existing cannon because it requires understanding what has been done, how different works relate and the different ways of approaching the topic. Reasoning about the research area iteratively generates insights that inform and position ongoing work.

One way to concretely represent the scope of a research area is to present, through a survey paper, a research design space. A research design space defines an area by categorizing and differentiating various approaches. It gives researchers a common ground on which to position their research or find gaps in the area [25]. An early example of this from the HCI literature is Card et. al's canonical design space of input devices [6], which has been used by many researchers to position their contributions and find new areas to explore. Other examples include a vocabulary and categorization of gestural interactions [56], a taxonomy of 3D tabletop systems [26], and a design space of composite visualizations [30]. What is common among design spaces is that they make design decisions explicit, summarize what is possible, and what is under-explored.

C&C '17, June 27 - 30, 2017, Singapore, Singapore

ISBN 978-1-4503-4403-6/17/06...$15.00
DOI: http://dx.doi.org/10.1145/3059454.3059472

In design domains, a design space defines the bounds of a problem space and solutions in it. Design spaces are often dimensional representations that enumerate the design decisions to be considered, along with potential options for each decision. If externalized as tables, matrices, lists etc, design spaces can be very effective for communicating with collaborators and clients. Representing design spaces with visualizations provides cognitive and perceptual benefits such as pattern recognition [22]. Identifying patterns, gaps, and seeing disparate connections are important for distinguishing between what has and hasn't been tried and may lead to new perspectives and creative leaps [47]. Also, creating external representations of the design space reduces cognitive load [35] and provides opportunities for reflection in practice [52]. The question we ask in this paper is whether design spaces can be appropriated to support researchers. Research tools exist, but, as discussed later, they support only certain types of reasoning and are not appropriate for constructing a research design space.

A research design space (RDS) can be explained through the metaphor of dimensions in physical space. A dimension is an aspect of the research area that encompasses the scope of all possibilities for that aspect. In computing, the algorithmic complexity in terms of time might be one dimension, and algorithmic complexity in terms of storage requirements might be another. One can think of these two dimensions occupying the X and Y axes of a two-dimensional space. The dimensions define the known bounds of the space and existing examples, such as specific algorithms, can be plotted in this space. Of course, a research design space (RDS) is likely to be multi-dimensional and complex. Such spaces are difficult to reason about in the abstract, pushing our cognitive limits, and thus an external representation is very helpful.

What would a tool that supports the creation and exploration of an RDS look like? It would need to represent numerous dimensions, some categorical, some continuous; support adding examples to the space and identifying how each example fits along each dimension; and support ways to visualize and explore the space. Such a tool could be used in conjunction with existing research tools, such as reference management systems, to form an ecology of tools that can support researchers at varying points in their research process. However, it is unclear how researchers reason about RDSs and whether such dimensional reasoning would come naturally. Do they start with examples or dimensions? Is the RDS artifact more valuable than the process of constructing it? Are researchers able to create dimensional representations of their research directly or do they need significant scaffolding and training?

This paper presents three studies that explored how researchers create research design spaces (RDS) and reason dimensionally: two studies adopt a 'dimensions-first' approach and one adopts an 'examples-first' approach. In this paper, we refer to the domain as the "research area"; the artifacts (papers, algorithms, designs, concepts, and solutions) in the research area as "examples"; and the space of all dimensions and examples for that research area as the "research design space" or RDS. The contributions of this paper are: 1) the Design Space Explorer, a tool that affords the creation of a research design space, 2) results from two studies of 'dimensions-first' RDS creation, 3) results from a study of 'examples-first' RDS creation, and 4) insights about how researchers interact with a dimensional representation of their research and how they interact with tools designed to support this process.

BACKGROUND

Many research support tools exist; however, few allow researchers to construct and then reflect on both the details and high-level concepts associated with their research. Design spaces have the potential to allow researchers to balance these two divergent needs. In this section, we describe existing research support tools, the roles of ideation and sense-making in research, review previous work on design spaces, and describe how design spaces might be appropriated for research.

Research Support Tools

A plethora of tools help keep track of and categorize literature. Reference management software [19, 44, 61, 37] supports grouping and sorting PDF documents manually or by their meta-data. Annotations and tags can also be applied to these documents. Docear leverages annotations to provide a visualization of all documents and annotations [2]. Biblioref allows the user to provide their own hierarchical taxonomy of concepts, which they then use to classify papers [39]. Hyperref supports research paper assessment by providing a way to discuss elements of a paper [32]. All of these tools are based on user-defined semantic or relational properties of papers and do not support reasoning about the papers dimensionally.

Mind maps focus on concepts and the relationships between them. The map is a hierarchical directed graph, originating at a central main concept. Other nodes branch off and contain user-added text which represents ideas or thoughts, and edges represent relationships. Tools include MindMeister, XMind, MindManager, FreeMind, and Coggle [46, 59, 45, 23, 9]. Mind maps are useful for some conceptualization and exploration tasks, but do not support dimensional reasoning.

Citation graphs, in which vertices are papers and edges point to cited work, can help researchers find relevant and popular work. Recommender systems that suggest papers based on these graphs are useful, but these recommendations disconnect the researcher from the space as a whole. When visualized, provenance and temporality are preserved, as in the case of Citeology [40]. Analytics help expose temporal and conceptual patterns, as shown in CiteSpace [8]. These visualizations support exploration of the citation graph, but they enforce a relational perspective of the research area that is based on citations and do not help generate a dimensional representation.

Visualizations are also used to explore a paper's content. Docuburst provides a visual summary of the words in a document based on their is-A relationship with other words in the document [10]. Parallel Tag Cloud supports comparisons between documents [11]. Words are listed along each axis, which represents a document, and a path shows where that word appears along other axes. For both techniques, the meaning is preserved as well as the relationships. Both represent the specific document text, which may be too rigid for researchers who want to construct conceptual representations.

Spreadsheets may help reason about examples [21, 53] but were designed for data manipulation. Dix et al. note that the success of spreadsheets for organizing, matching, and grouping is due to their flexibility and familiarity [16]. However, spreadsheets and tables are limited in communicating to the user what they can do with them. Kandogan et al. explore DataBoard, a free form spreadsheet-like interface, and compared it to a plain spreadsheet on simple problem solving tasks [31]. They determined that free form was more flexible and better for iteration, but require more effort. While spreadsheets may be able to represent the dimensions of a design space as a table, they are not particularly interactive and do not support the iterative process of constructing a design space.

Creativity in Research

Research is an amalgamation of ideation and sense-making. Researchers iteratively curate relevant related work, attempt to make sense of it, identify gaps, and brainstorm solutions that might address the gaps. Sense-making describes the process of moving from raw data to insight. Russell et al. describe the "learning loop complex" in which people create a representation of the data, identify parts that don't fit, and recursively adapt the representation to fit the data [51]. Pirolli and Card build on this idea with the "notional model of sensemaking" which takes into account the preparedness of data [50]. They describe sense-making as having loops that support top-down and bottom-up processing [50]. This defines sense-making as an iterative process; where data, representations, and insights are dynamic. This process involves convergent and divergent thinking, which are important aspects of creativity [12, 27].

As researchers engage in sense-making, they get a better sense for the problem space and develop research questions. Ideating solutions to a research question may involve reviewing related examples, artifacts, or papers. Diverse examples that represent distant parts of a solution space can lead to more creative ideas [7]. Similarly, distant associations of concepts can also lead to creative ideas [42]. On the other hand, examples that are very similar or represent only a subsection of the potential solutions can inhibit creativity and lead to design fixation [29].

Since examples can be both helpful and detrimental for ideation, they should be chosen carefully. By Buchanan's definition, research is a wicked design problem [5] in that it is constantly evolving, incomplete, and may contain conflicting information. Given this definition, choosing relevant examples automatically is very difficult. One solution might be having researchers curate their own examples. Webb et al. suggest that a platform that gives visual structure to objects and authoring capabilities to the user promotes ideation [57]. Popular curation tools, such as Pinterest, allow users to group and save web resources. However, as pointed out by Kerne et al., Pinterest's grid layout is very limiting in terms of spatial organization [34]. Kerne et al. study platforms for curating web artifacts [33, 34] and claim that these platforms should provide room to organize, group, and annotate artifacts.

Design Spaces

Previous work has defined what design spaces are and established methodologies for creating them. In the Question-Option-Criteria (QOC) method, designers use a list of questions to outline the design decisions, the options for each question, and criteria that must be satisfied [38]. Heape [28] presents a similar definition of design space and stresses the iterative nature of design, in terms of the designer's understanding of the solution space, as well as the problem space. Morphological analysis, developed by Zwicky and Wilson, is a matrix-based approach to representing the parameters and conditions in complex problems, through which the problem as a whole can be understood and possibilities can be examined [62, 63]. Dalsgaard et al. and Biskjaer et al. build off the work of Zwicky and Heape, simplifying the process and discussing the benefits of their design space notation for design problems [3, 15]. The design space notion and notation presented by Dalsgaard et al. is targeted specifically for the design of media facades [15], which Biskjaer et al. generalize and discuss in terms of design more broadly [3].

Biskjaer et al.'s framework for documenting and navigating design spaces, defines a design space as "a conceptual space, which encompasses the creativity constraints that govern what the outcome of the design process might (and might not) be." They posit that a design space is "co-constituted, explored and shaped by the designer during the design process" [3]. Their design space schema consists of columns, the headings of which are aspects of the design, rows are potential designs, and the contents of a cell reflect the option of the aspect present in that particular design solution. Biskjaer's work shows that considering a design or research area as a multi-dimensional space is worthwhile because it helps users consider ideas in relation to each other and the problem space. However, to our knowledge, no interactive tools have been developed to facilitate this iterative process of dimensionalizing the space.

Design Processes

While the artifact of a design space is helpful for documentation and communication, perhaps more important is the process one goes through *while* dimensionalizing the space. The process of defining the problem space is a reflective act where the designer or researcher reflects on the problem space, the solution space, and potential opportunities for exploration [52]. Bonnardel and Sumner investigate ways to support reflection in design in the form of critiquing based on certain criteria [4]. Several researchers discuss the use of maps to reflect on different aspects of the design process as a whole [14, 36].

Design processes live in the problem space, the solution space, or both. Beaudouin-Lafon and Mackay present a methodology that focuses on the solution space, where the goal is to generate as many potential design solutions as possible [1]. Schon discusses the necessity for the designer to construct the problem space, given that the complete bounds of the space are not known in advance [52]. Dorst and Cross present design as the iterative co-evolution of problem and solution space [17].

The related work we have reviewed demonstrates that constructing multi-dimensional design spaces is a helpful research task. Methodologies to define design spaces exist, as do platforms for curating information. While these have been shown to be useful for sense-making and in the design ideation process, there appears to be a lack of interactive tools to help researchers externalize a formal, dimensionalized design space.

RESEARCH QUESTIONS

There exists a gap between traditional research tools and the ability to create a dimensional research design space (RDS). However, it is not clear how researchers construct dimensional representations, interact with them, or use them to further their research work. Design spaces are useful for designers, but can researchers use them with the same success designers have? Thus, we have the following research questions:

Q1 Do researchers reason about their research dimensionally?
Q2 Are there other ways researchers reason about research?
Q3 How natural is it for researchers to use a tool that was designed to support dimensional reasoning?
Q4 Can a tool guide researchers to reason about their research dimensionally?
Q5 When constructing a research design space, do researchers start with examples, such as papers, and move to dimensions or do they start with dimensions and fill in examples?
Q6 Is the DSE equally effective for all researchers?
Q7 Is the DSE equally appropriate for all research areas?
Q8 What do researchers gain from dimensional reasoning?

DIMENSIONS FIRST: THE DESIGN-SPACE EXPLORER

To investigate the ways in which people might engage in dimensional reasoning during research we created a 'dimensions-first' prototype, The Design Space Explorer (DSE); a web-based tool that aids researchers in externalizing a set of dimensions that represent a complex, conceptual space. The DSE has three modes:

Create: define the dimensions of the space.
Populate: add examples by selecting values for dimensions.
Explore: interact with visualization of dimensions and points.

Responsive enabling guides users through the RDS creation process, but allows them to iterate between the three tabs.

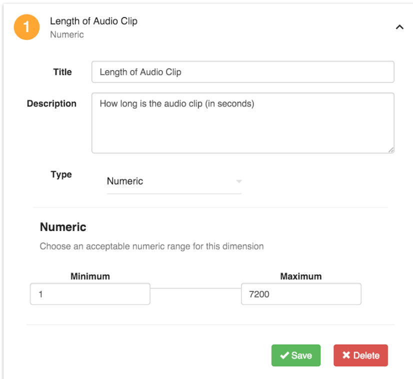

Figure 1. Dimension Card: setting audio length as a numeric dimension.

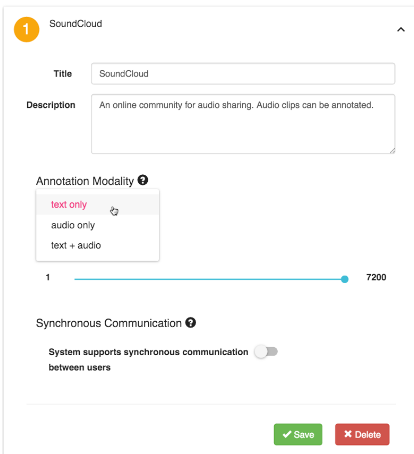

Figure 2. Example Card: adding SoundCloud as an example of an audio annotation system. The user interacts with each dimension widget to specify the values for the current example.

First, relevant dimensions are defined in the RDS, and then the DSE allows researchers to plot examples along the dimensions. Once examples have been plotted, researchers can explore and reflect on the resulting RDS. The DSE builds on Biskjaer's tabular approach to externalizing the design space [3], but by creating an interactive visualization, we leverage human spatial reasoning and enable faster and more fluid exploration of the space than is possible when looking at a static table.

Create: adding dimensions

The first step occurs in the *Spaces* tab, where the user adds a new RDS by giving it a title and a description. They can also load or delete RDSs. Then, responsive enabling guides users to the *Create* tab, where they can add and define new dimensions. Both the *Create* and *Populate* tabs use the metaphor of cards to represent dimensions and examples (Figure 1). To create a dimension, the user provides a name, description, and dimension type (categorical, numeric, or boolean).

Populate: adding example points

After defining at least one dimension, users can move to the Populate tab and add examples, which can be research papers, solutions, devices, systems, or studies. When adding an example, the dimensions that were defined in the Create tab are listed as widgets (Figure 2). The user specifies a value for each dimension using a widget slider (numeric dimensions), drop-down list (categorical dimensions), or a toggle switch (Boolean dimensions). The value for each dimension corresponds its location in the multi-dimensional space.

Explore: Visually reflecting on the whole space

In the Explore tab, users interact with a parallel coordinates visualization of their RDS and have the potential to identify gaps in the research, compare examples, and ideate new solutions. The visualization supports these types of discovery through filtering the examples by brushing along any dimension.

Visualizations leverage human spatial reasoning skills to identify patterns and make sense of the underlying data [22]. We expect that RDS tools would have multiple visualizations, but used one for our study. Scatterdice [18] allows many dimensions to be considered, two at a time, but it is best suited for non-categorical data. Bertin Matrices [48] are compact and support high-dimensional data. However, based on a review of design spaces, many of which were sparsely populated with examples, this compact view is not necessary. Parallel coordinates was chosen because it is flexible and supports continuous and categorical data. The sparseness of the examples suggest that over-plotting and occlusion are unlikely, which can be problematic for parallel coordinates. To improve usability, we implemented the traditional interaction techniques [54] - saturation brushing [58], multiple brushes (one for each dimension), sortable dimensions, and dimensional reordering.

Dimensions in our visualization are similar to facets in faceted navigation. Facets provide multiple filters that each describe one aspect of a item. Facets provide an overview of the relevant aspects of an item that is being searched for [60]. Facets can be either hierarchical or non-hierarchical [20]. They can also be placed into the context of the item itself [43].

DSE Study 1: Pilot Study

In Study 1, computing PhDs and post-doctoral researchers participated in an hour-long task and post-task interview. The session began with an introduction to the concept of a RDS and a walk-through of how to define one with the DSE. Participants then had 60 minutes to interact with the DSE in the context of their own research. They were instructed to input dimensions, populate the space with research examples, and interact with the visualization. After observing their interaction, we conducted a semi-structured interview about how they used the DSE, their research area, and their thinking process.

The goals of the first study were to understand users' approach to using the DSE and dimensionalization, whether they created all dimensions first or were more iterative, what kinds of reflection and ideation the DSE inspired, and their impressions of the value of the DSE.

Results

Four PhD students (2 female) and one post-doc in computing participated in the first study. They were compensated $15 for their time. Two participants used the entire 60 minutes to interact with the DSE, two used nearly the entire 60 minutes, and the fifth participant used 40 minutes. To obtain insights we reviewed the session videos and conducted a thematic analysis.

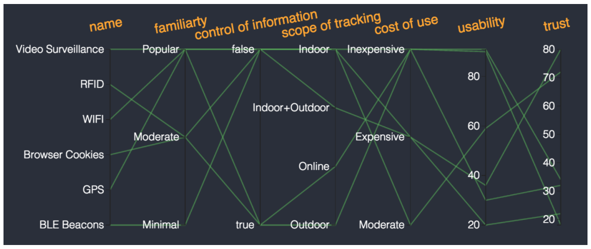

Figure 3. A RDS from the first study, created by P1-5.

Topic Matters

Of the five participants in our study, two chose to use a research topic outside of their current research area because they thought it would be easier to dimensionalize. P1-4 chose an area she had worked on before starting her PhD because, as she said, "I wanted to choose something I could easily populate. I had lots of design spaces I think, like in terms of the main research topic, but I didn't know how to populate them." P1-3 also chose a research area that was easier to dimensionalize. At first, she couldn't decide on any dimensions and asked the facilitator for clarification, to which she responded, "For papers I don't see having too many things like that." She then changed to a different RDS entirely. P1-3's initial RDS was the only RDS that didn't describe a system, but her second RDS, like the others, did describe a system. Both P1-3 and P1-4 expressed an initial concern that it was not obvious to them how to dimensionalize a research area. Dimensionalizing is difficult, but some topics appear to be easier to dimensionalize.

Assigning Dimensions: Metrics vs Features

Most participants created their RDSs from the perspective of their own research and in terms of systems with features and metrics. Three participants (P1-1, P1-4, P1-5) used their own research approach as the first point to populate their RDSs. P1-3 was the only participant who didn't create an RDS based on her own system. Some participants used evaluation metrics and others used potential features of a system when creating their RDSs. For example, P1-1 and P1-3 chose dimensions that reflected design features of existing systems (i.e. "modality", "location"), while others (P1-2, P1-4, P1-5) mostly chose dimensions that could be used as metrics on which to evaluate existing systems (i.e. "usability", "cost of use").

P1-4 referred to her dimensions as "metrics" and chose to use "strength of security" and "strength of usability", which were specific metrics that she had used during her research. P1-5 also used dimensions to represent evaluation metrics but chose them with the goal of investigating the relationship between them. For example, he said he was interested to see if there was a trend between familiarity and trust in his research area of usable privacy and security. Unlike P1-4, whose metrics corresponded to evaluation calculations used in research, P1-5 created metrics based on his own opinions.

Dimensions were also assigned based on features of existing systems. P1-3's RDS consisted mostly of design features relevant to the system that she was designing and one dimension related to the goal of the system (specifically, what type of learning style it corresponds to). P1-1's RDS consisted exclusively of factors to consider when designing a system within his research area (specifically, a system for collaborative video editing). This is consistent with P1-1's definition of design space, which is centered around the idea of parameters and design: "a space with many different parameters, many different dimensions that could be used to create a design or a design could be defined by those different dimensions or parameters." It is not clear from our results whether participants always thought of their research as systems with features and evaluation metrics or whether this was prompted by the DSE.

P1-2's defined a design space as something that is not created or designed, but discovered based on the predefined nature of the research area itself. His RDS of robot motion planning included a combination of mostly metrics for evaluation (i.e., time complexity) and design features of the implementation (i.e., underlying data structure). He noted that using the DSE was natural because "you can think about the field and you think about the different approaches. And the approaches have advantages and disadvantages, which you can highlight here." He admitted to thinking broadly to capture the big picture of the research area, rather than sub-dimensions that would have been more specific.

Iterative and Emergent

All participants iteratively revised, deleted, and added dimensions and examples. P1-4 used the visualization to evaluate her dimensions: "I was continuously going to explore and then I was like, no this doesn't sound right so I had to go back and change maybe the label or delete the dimension or come up with a new dimension." P1-5 noted that he often added an example to his RDS, which made him think of more dimensions to add. He mentioned adding RFID as an example, which he hadn't included in his prior literature review because it was too expensive, but now it led him to add a dimension to reflect that. He said, "I know that adding cost as a dimension was directly tied to remembering RFID. And then video surveillance, I think I went back and added trust and usability.' Here and elsewhere we see examples prompting the addition of dimensions as a way to differentiate examples and to fill out the research space.

As participants iterated, their mental models of their RDS evolved. In particular, P1-3 showed a change in thinking from narrow to broad throughout the study. Her RDS began with very specific dimensions about the system she was developing, most of which resulted in a category called "irrelevant", which was selected for most examples and resulted in a less interesting visualization. Towards the end of her session, she added a category that was more relevant, broader, and helped to differentiate the systems in the RDS. She had previously been thinking of ways the approaches were similar, but when she stepped back to consider differences, it lead to a more interesting set of dimensions.

Missed Opportunities for Reflection

Participants used the DSE to compare different examples in their RDSs. P1-4 noted, "I could see for example how is a particular CAPTCHA doing. I could also see if I needed to compare for example two or three at a time, did they have the same strength for example." P1-2 noticed that there was a lot of crossover among his examples and commented that they were essentially doing the same thing. P1-5 said, "I was curious to see if there were any trends" and explained in the interview where he might have seen a relationship between trust and familiarity, though the examples he had populated the RDS with did not show evidence of this relationship. P1-1 commented that the visualization didn't make it easy to see the empty spaces (gaps). He often only had one or two items for each dimension, which may explain why he saw few gaps. Despite being shown how to filter (brush) the visualization along each dimension, which may have helped identify gaps, none of the participants used this feature.

Value

Some participants used DSE to identifying new dimensions and organize their thoughts. They also offered ways they might use it more in the future. P1-4 thought the DSE could be helpful for brainstorming, "to maybe explore your research topic. To try to find more angles where you can look at that research topic and find problems and see how things are related." So even though P1-4 used the DSE during the study as an evaluation tool, she would use it in the future for ideation. P1-5 would also use it in the beginning of research: "I think I would use it early on to really get a grasp of what I'm looking for in my research. If only to understand what exists right now, that's something I would want to do early on to know what can I borrow from or generate that's novel based on that."

Some participants, such as P1-1, thought it might be helpful for organizing and exploring new research areas, "[The DSE] would certainly give you an idea of what is existing out there and what are the properties." P1-3 said she would use the DSE early in her research process, but for curation and organizing as she found more sources, rather than ideation. P1-2 thought the DSE would be helpful for developing high-level conceptual understanding of their research area: "You probably think you know all these things but until you actually sit down and lay them all out you realize you don't know as much as you thought you did ... So it definitely helps your understanding."

P1-3 was the only participant to use another tool to externalize her thought process. She opened a text editor and typed a list of relevant factors. She said that she wanted to have another tab in the DSE where she could do whatever she wanted, such as make a collage or flow chart. She thought such a feature would provide value. She self-reported that because she has a background in art, she thinks of design as more open-ended and abstract and referred to the DSE as being "so specific."

DSE Study 2: Longitudinal Study

In our first study, we received mixed results from participants. Some struggled to define dimensions and others were starting to use it to compare and evaluate examples. If dimensional reasoning emerges over time, then using the DSE for a longer period of time might allow researchers to reflect differently.

For our second study, we performed a six week case study with graduate research students. Participants came into our lab for 3 in-person interviews, they were sent weekly surveys, and they were asked to use the DSE for 2-3 hours weekly on their own time. During the first interview, we asked about their research areas and practices, gave an introduction to the

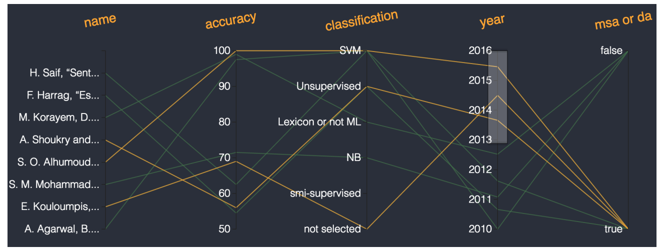

Figure 4. A RDS for P2-3 shows different approaches for classifying tweets for sentiment analysis. The year dimension is brushed to highlight only most recent papers (in orange).

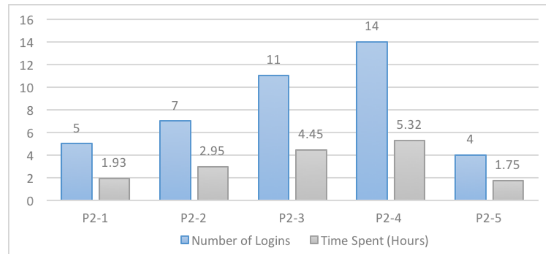

Figure 5. Participant usage of the DSE during the 6 week study.

concept of design spaces and dimensionalization, walked them through using the DSE, and they performed a practice task of dimensionalizing a list of fruit to get experience using it.

The goals of Study 2 were to continue to study the DSE in terms of how people use it and ways it supports or doesn't support dimensional reasoning, but with respect to larger sense-making and reflection loops than in Study 1 since participants were able to use the DSE longitudinally.

Results

Five graduate research students (4 female) participated in the longitudinal study. Participants were recruited via email and were compensated for their time, receiving a $10 gift-card at the initial training and a $20 gift-card at both the mid-way and final interviews. Three of the five students were involved with design research such as formal design methods, architectural design, and design patterns. The other two students were designing new algorithms to use in their research.

Contexts and Intentions of Use

Each participants had different purposes and use cases for the DSE. Several saw the DSE as a tool for knowledge management. In the beginning, P2-1 said, "I'm terrible at taking notes and this would organize my thoughts." P2-3 was writing a survey paper and wanted to use the DSE to evaluate whether she had organized her sources correctly. She said, "I need it in my research" and also, "I can arrange the names of the papers and the names of the authors and the names of the algorithms that they have used." P2-4 was using the DSE to curate research papers on topics that she was considering working with such as the acoustics of physical space. P2-4 also had a very unique space that investigated the American Southern Literary tradition for inspiration to design a museum that celebrates southern culture. P2-5 was also using the DSE to curate papers as he was reading them, but used the DSE to compare these papers. P2-2 was trying to characterize and classify the bounds of her set of Design Patterns for CS education.

Name	*Area of study*	*Topic of research*	*Year of study*	*Usage of tool*
P2-1	Computing	Crowdsourced design in citizen science	PhD, year 3	Knowledge management, compare approaches
P2-2	Computing	CS education research	PhD, year 2	Categorize design patterns for CSE
P2-3	Computing	Sentiment analysis in social media	PhD, year 4	To write a survey paper
P2-4	Architecture	Uncertain	MSc, year 2	Explore thesis topics, collect design inspiration
P2-5	Bioinformatics	Computational methods to study how plants influence health	PhD, year 4	Compare metrics for path ranking, evaluate sections in papers

Table 1. Participants' areas of research, their academic backgrounds, and how they used the Design Space Explorer in Study 2.

Successes

The most common observation from this study was that the DSE helped participants think about their RDS systematically and develop a holistic view of their research. P2-4 noted the DSE helped broaden her perspective: "It really helps to dimensionalize them because I start to think of them less as articles and more as articles in like different categories that they fit into. Which is something I haven't really done before." P2-1 echoed the sentiment that considering the space holistically was important, "when I was writing I was only focusing on one aspect, but when I was looking at the chart [visualization] I wasn't sure if I was fair enough in my writing."

The DSE helped some researchers approach research and writing differently. P2-4 used the DSE for a literature review and said, "[I'm] becoming more aware of how to identify gaps in research." She said that the visualization, "shows you which ones [dimensions] I've found more [examples] of so far and which I need to focus on seeing if there are other types of examples for." P2-5 used the DSE to establish the novelty of his research compared with others and to structure papers, "I'm trying to see what kinds of things are emphasized in each of the different sections - your results, intro, and methods" He explained that this was supported by the DSE, "very clearly saw where things were actually lacking so that was very easy for me to see where I should be writing or where I should make more of a point to talk about in each of the different sections where everybody else was kind of missing it."

Challenges

The challenges faced by participants were related to either the visualization or a mis-match between their mental models of their research and the dimensional representation with which they interacted. P2-3 described the visualization as difficult for her to use and that she would have preferred a table. P2-2 liked the visualization, saying "this style of visualization, I haven't seen it before and this is really a good view of the space," and "This kind of visualization makes me perform a better gap analysis on my patterns." However, she was ultimately unable to use it and hoped for a more compact representation.

The mismatch between participants' ways of reasoning and the intended dimensional reasoning was also an issue. Here are some ways of reasoning that participants expected:

Semantic Reasoning All participants expressed a desire for support of some sort of semantic reasoning and the ability to assign meaning to examples with tags. For example, P2-2 described: "[I] wanted to mark multiple options for this category ... This [example] could be random or self-selected so I actually have two of them so this is not so practical."

Relational Reasoning Many participants visually or conceptually related two examples or dimensions. For example, P2-1 explained that the visualization gave her a false sense of relatedness, though she explained that in her mind some dimensions actually were related, "I could see some inter-relations between the two [dimensions in the visualization]." P2-4 thought the DSE supported relational thinking, "you are looking more for the connections between them and the similarities that you wouldn't have thought about on your own"

Hierarchical Reasoning Most participants expressed a hierarchical understanding of their research area. P2-2 created a new space to represent meta-dimensions corresponding to her original RDS, "I decided to go to a higher level to see at least a glance." P2-4 noted that some dimensions were more important than others and should be ordered accordingly.

Blended Multiple participants made comments about the need to support multiple types of reasoning. P2-1 noted that, "I wanted to see inter-relations between aspects," and "If I had the checkbox I could do that." P2-1 explained: "when I think of my dimensions some of them are inter-related - it's not that they're always a hierarchy."

Discussion for Study 1 and Study 2

In our first study, we observed the entirety of the participants' interaction. They iterated frequently as they shaped their RDSs and defined and refined dimensions based on examples that they added. This is consistent with the view that sense-making is an iterative process [50] and that designers iterate through periods of 'problem structuring' and 'distinct problem-solving phases' [24]. This also suggests that dimensions emerging from examples might be a more natural way of creating design spaces (Q5). Researchers approached the DSE from the perspective of their own research and were most likely to encode their research as systems and in terms of metrics, features, or both (Q6, Q7). It is likely that design spaces, which were intended for designers, forced researchers to encode their research as systems that could be designed. Even though it might have been more beneficial to build a RDS for their current research topic, some participants deliberately chose topics that they perceived as easier to dimensionalize. This suggests that dimensionalizing is cognitively challenging (Q3), and that certain topics lend themselves more obviously to dimensionalization (Q6, Q7). In particular, research involving the design of a system appears to be easier than abstract topics to dimensionalize, as all participants chose research areas consisting of artifacts or systems for dimensionalization (Q6, Q7). The DSE guided researchers to dimensional representations (Q3); however, it is not clear that participants always reflected on them dimensionally (Q1). Participants described a variety of ways that the DSE changed their perspective about their own research, such as thinking more holistically (Q8). Dimensions allowed them to compare and contrast papers so that papers once thought of as solitary were actually related to others (Q8).

In the second study, researchers had more time to become comfortable with the DSE and were more vocal about mismatches between their mental models and the dimensional way in which they were being asked to interact (Q1). There is evidence of researchers reasoning dimensionally but there were also other modes of reasoning, such as semantic, relational, and hierarchical (Q2). Researchers didn't consistently receive the benefits associated with design spaces such as identifying gaps in their research. Participants created multiple small fragmented versions of the same RDS rather than elaborating a single detailed space. In search, information foraging describes how people maximize results and minimize effort [49]. For example, search engines encourage searchers to quickly scan results for information rather than delving deeply into individual websites. It may be similar here, simple fragmented spaces are good enough and participants satisfice [55] rather than iteratively refine, which is cognitively taxing.

In both studies, participants opted away from complex and detailed design spaces. In the first study, we see that participants chose spaces that were easier to dimensionalize. In the second study, researchers created multiple, simple fragmented versions of the same design space rather than creating one complex and detailed space. This may suggest that some parts of research are better suited to dimensions (Q7), that some dimensions are more obvious than others, or that the process is more important to researchers than the artifact (Q3).

EXAMPLES FIRST: PAPER PROTOTYPE

Given that so many design spaces exist in literature, we assumed it would be straight-forward to create a tool to support the construction of RDSs. But, given the challenges faced by participants in the first study, we employed a paper prototype in our third study. We did this to reduce the assumptions we were making and to provide more flexibility.

In the first two studies, participants experienced some difficulties in constructing their design spaces, especially in terms of a mismatch between their own mental models of their research and the affordances provided for dimensionalizing. Since several participants were interested in semantic encodings, our third study allowed users to employ an examples-first approach that supports tagging. We expected that this might mimic researchers' traditional process of tagging papers, but could still result in a dimensional representation. To investigate the examples-first, semantic encoding approach, we devised a paper-prototype and wizard-of-oz study [13] described next. Given the results from our first study, this low-fidelity approach gave us more flexibility.

Study 3: Tag-Based Card Study

In Study 3, participants were asked to arrive with a list of papers, artifacts, or examples that were relevant to their current research and then they completed the following tasks: 1) create examples by writing paper names on index cards, 2) tag these examples by writing an attribute that describes the example

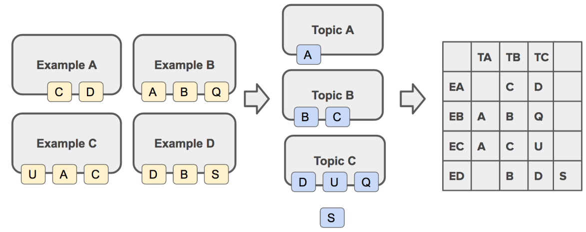

Figure 6. Participants labeled index cards with the example name and tagged them with attributes (in yellow). Participants grouped tags (in blue) by putting them on a 'topic' card and labeling it. Participants looked at the dimensional representation formed by plotting examples and corresponding tags along grouped dimensions in a table.

Figure 7. A participant from Study 3 constructs their RDS.

on a post-it note and attaching it to the index card, 3) group related tags from the previous step by placing them on index cards and labeling the index cards, and 4) review a resulting RDS spreadsheet created by a wizard. The RDS had examples as rows and groupings (dimensions) as columns. The value of each cell was the tag from that group for that example. We chose a spreadsheet because participants were familiar with them and creating a table in real-time was more feasible in a Wizard of Oz study than creating an interactive visualization. Figure 6 shows an overview of the process.

To help participants understand the process, we walked them through an example in which we created a design space of cars. We used 4 cars, and tagged them with attributes (such as 'red', '30MPG', '3 doors', etc.) which were then grouped into categories such as 'color', 'MPG' and 'number of doors'.

Results

11 graduate student participants (recruited through graduate school email lists) participated in a 90 minute session and were compensated with $15 gift cards. Participants were primarily computing majors: 3 Engineers (1MS, 2PhD), 8 CS students (5MS, 3PhD). Topics ranged from recommendation systems, to interaction log analysis, to road maintenance theory.

Existing practices Understanding participants' existing practices and how they consider their research area was important because it is possible that these things influence the way in which they perform the card tagging task. Most participants do not use organizational tools or techniques in their daily research life. Most take notes on a paper and then move on. P3-8 uses folders of papers and P3-9 has a hierarchy of folders and subfolders. P3-2's notes on each paper were highly organized, with categories, subcategories, and related bullet points. P3-11 had the most complex organization method, using a series of documents to organize, tag, relate, and index the papers.

When asked to describe how they decide if examples are relevant and the ways in which they use them to inform their research, there was a range of answers. Some use the literature to find promising approaches or solutions to a problem (P3-2, P3-6), to explain behavior they observe (P3-5), to inform the steps of their implementation (P3-3, P3-7), to explain what the problem is (P3-6), and most commonly to extract aspects that can be applied to another problem (P3-2, P3-8, P3-9, P3-10, P3-11). Unsurprisingly, none of the participants say that their primary use of related work is for the purpose of defining the space of their research.

Process Most participants easily tagged the examples, grouped the tags, and successfully completed the task. Two participants completed the task in an unexpected way: P3-2 and P3-9 tagged categories of papers rather than papers themselves. P3-2 didn't run into any trouble as the papers in each category were similar, but P3-9 had trouble applying tags consistently to papers. We discussed the intended process with P3-9 at the end, saying that the intention was to tag papers, but he felt the way he did it made sense, noting that "most of the times I read the papers to make this kind of categorizing in my brain".

Most participants reported that the process felt natural. This was partially due to the familiarity of tags in general. P3-6 explained, "It was natural for me because I am trained to assign index terms for my papers. I always do that." P3-10 said, "This follows the same mental process we follow when we write the related work section in our papers". Some reported even having the categories in mind as they were doing the initial tagging (P3-1, P3-5, P3-6, P3-8). Only two reported that it was unnatural: one because he was not accustomed to using pen and paper and the other because she was considering it as a way to keep track of the steps of her current implementation, which are always changing. P3-6 saw the value of the process for younger researchers working to construct an understanding of a new space, but felt he did not need such scaffolding: "Your way is like going from details to big picture. [..] For me the way I think is going from big picture to small picture or details". For P3-6, the DSE might have been a more natural experience. However, this also points to the fact that tools, such as the DSE or our Examples-First prototype, may be better suited to different stages of research.

Final artifact The final spreadsheet resulting from P3-1 contained a very well-formed RDS: it had columns that were clear and reasonably independent, rows that represented a variety of previous works, and most cells filled in, indicating how each work fit along each dimension. This was impressive given that the RDS was constructed in 90 minutes. In the other participant spreadsheets, there were the beginnings of RDSs, but they were incomplete. Some RDSs had cells containing lists of tags, meaning that the example lived at multiple points along that dimension. Most of the RDSs had at least 1 empty cell, but 4 of them had a significant number of empty cells that the participant did not feel they could fill in. While a limited number of empty cells and lists of tags can still be present in an RDS, if there are too many, it becomes difficult to see the bounds of the space.

Perceived benefits of RDS as Artifact When asked what benefits participants thought the RDS table would provide them, most mentioned the ability to easily see which papers discuss particular topics (P3-3, P3-5, P3-6, P3-8, P3-11). For example, P3-8 said, "it would make it easier for me to go back and look at these papers using the particular information in a table format." Others would use it to plan out their steps of implementation (P3-3, P3-7). P3-2 wanted to find the frequency of each tag to see the distribution across examples. P3-10 thought she could use the table to compare approaches, but she said that any display of the papers and the tags would allow her to do that. P3-11 pointed out that some papers that he knew were similar due to the fact that they had similar tags in a given category. Some participants said they would refer to it as they were writing a survey paper or related works section (P3-8, P3-11), but none of them would share the actual table in the paper as a representation of their RDS.

When asked if there were additional ways in which participants would like to interact with the table, they listed the ability to add notes (P3-1), link to relevant information or papers (P3-3, P3-10), color code tags (P3-4, P3-10), or drill down to see more specific tags (P3-6). P3-11 had an interesting suggestion, which was to use the table as the basis of a network graph, where papers would be displayed nearer or farther based on how similar they were along a given dimension.

Study 3 Discussion

Every participant completed the task and most found the process natural. However, based on their resulting RDSs and what they said about the table representation, it is clear that most were not able to create a useful RDS in the time given, nor did they understand the paper prototype as a way of working towards a dimensional representation. Thinking about research and trying to develop an understanding of how a set of papers related to one another is a cognitively demanding task, and getting to a useful dimensional representation by tagging papers likely requires more time than the 90 minutes in our study.

A few mentions of dimensions show promise for a tool that supports dimensional reasoning (Q4). P3-10 mentioned comparing examples, though not in a multi-dimensional way as she said the table was unnecessary. P3-2's interest in tag frequency can be thought of as dimensional in a way: he was interested in the dense areas of the design space along one dimension at a time. P3-11 was the only participant to consistently use the word "dimension" to describe the columns in the table and desired a network graph representation that displayed the relatedness of papers along one dimension at a time.

Participants typically used fragments of related work to apply to their current work or inform their approach (Q2). Perhaps because of this, they were not able to see the paper prototype or its benefits through the lens of dimensions. The participants who claimed that they had never considered their papers at one time and who didn't use any organizational tools, commented on the table as a way of tracking their papers. Experienced researchers described the importance of understanding the relationships between papers, but they had this understanding already and discussed the paper prototype in relation to what they already knew (Q6).

DISCUSSION

Given the prevalence of design spaces in research, we expected that student researchers would be better able to construct them. We observed a few challenges faced by participants when trying to reason dimensionally: (1) research papers represented fragments of information, (2) starting with dimensions was not natural, (3) multiple ways of reasoning, (4) topic matters when dimensionalizing, (5) focusing on metrics instead of features, and (6) starting with examples didn't lead to dimensional representations.

Our results suggest that most researchers think about papers in terms of how they use them, rather than what they are. Researchers may take a technique from one paper and apply it to a problem discussed in another paper. Instead of thinking holistically about these papers as points in a dimensional space, researchers appeared to think of them as pieces of a puzzle they are trying to solve. Thinking of papers as fragments in this way is problematic, because it results in empty design spaces. The "dimensions-first" approach appeared to force participants to think more holistically about the examples and their features. This resulted in richer RDSs, but the process felt less natural to them. Participants were more interested in these RDSs, but didn't discover many gaps in their research or novel ideas. It may be that they were too familiar with their work, but maybe they only enumerated the obvious features. A study by McCaffrey and Spector found that designers overlooked 56% of the potential features of a candle [41] and that the overlooked, obscure features led to more innovative solutions. Therefore, it may be beneficial to suggest dimensions automatically based a paper's keywords or content or otherwise guide researchers to reflect more fully.

Most participants in the "dimensions-first" study expressed a desire for the DSE to support other representations of their research that we have described as being semantic, hierarchical, or relational. For these participants the process did not feel natural and they often chose a topic for their RDS based on what they perceived to be easiest to dimensionalize. This suggests that dimensional reasoning may be easier for some topics than others. Participants in Study 3 had an easier time with the "examples-first" process and expressed less of a desire to integrate other modes of reasoning. However, based on the ways in which these participants discussed the benefits of the paper prototype and resulting table, it was clear that they did not internalize the dimensional aspect or reason dimensionally. Many participants chose to use metrics that evaluate the research rather than features that define the research. Participants also didn't appear to understand that there was a difference between metrics and features when asked.

While participants in the "dimensions-first" study struggled with dimensions and suggested easier ways to approach the task, they were forced to reason dimensionally. This appeared to place a high cognitive load on participants, but it also allowed them to gain more insights and to internalize the dimensionality of the end result. The tagging approach in Study 3 appeared to be easier and induce less cognitive load, but it also generated fewer insights for participants. The most useful RDSs resulted from researchers doing the hard cognitive work of thinking about the dimensions first, rather than letting them emerge from tags. More research is needed to understand these tradeoffs and what kind of systems can support researchers in thinking dimensionally.

LIMITATIONS

We had small numbers of participants and so our results are suggestive rather than conclusive. Future work should investigate a larger and more diverse sample of researchers. Research topics varied from CS education to southern literary tradition to road maintenance theory. The first and third studies took place in a lab and did not replicate the environment in which researchers typically work. Short-term studies (1-2 hours) were balanced with one longitudinal study (6 weeks), but future work could investigate longer timelines. Research is collaborative, but in our studies we removed the potential confound of multiple investigators to isolate researchers' thought processes. Future work might have researchers co-construct RDSs. Participants received minimal guidance so that their natural process could be observed, but explicit training with dimensional reasoning might lead to better outcomes. Finally, participants constructed their spaces without any automated assistance. Mining keywords from papers, suggesting related keywords, or searching for papers in the DSE would have likely made the process faster and easier for participants.

CONCLUSION

We have presented three studies that explore the process of dimensional reasoning in researchers. Our investigation contributes the first understanding of how researchers engage in dimensional reasoning and how digital tools might support and scaffold such reasoning. We developed a top-down approach, the DSE, which allows researchers to create dimensions and fit examples to those dimensions. Researchers iteratively refined these dimensions based on examples added. This resulted in a dimensional representation but appeared to be a cognitively demanding process for participants. A paper-prototype study followed in which researchers tagged examples and then dimensions organically emerged as researchers grouped the tags together semantically. This process fit the way in which researchers typically interact with papers; however, it often didn't always result in a dimensional representation.

Our findings about dimensional reasoning suggest that it is possible to scaffold researchers to create dimensional representations; however, it appears that dimensional reasoning comes more naturally to some researchers than to others, and also suits certain types of inquiry more than others. Dimensional representations did help researchers think more holistically about their research and afforded easier comparisons. However, we observed that in addition to reasoning dimensionally, researchers also engaged in semantic, relational, and hierarchical reasoning. These results suggest that supporting dimensional reasoning can offer benefits, but needs to be incorporated into tools or systems that also support other types of reasoning. Thus, to fully support the creative, open-ended task of research, rich systems that offer multiple tools for organizing, representing and exploring are important, and different modes of reasoning need to be accommodated.

REFERENCES

1. Michel Beaudouin-Lafon and Wendy E Mackay. 2003. Prototyping tools and techniques. *Human Computer Interaction – Development Process* (2003), 122–142.
2. Joeran Beel, Bela Gipp, Stefan Langer, and Marcel Genzmehr. 2011. Docear: An Academic Literature Suite for Searching, Organizing and Creating Academic Literature. In *Proceedings of the 11th Annual International ACM/IEEE Joint Conference on Digital Libraries (JCDL '11)*. ACM, New York, NY, USA, 465–466. DOI: http://dx.doi.org/10.1145/1998076.1998188
3. Michael Mose Biskjaer, Peter Dalsgaard, and Kim Halskov. 2014. A Constraint-based Understanding of Design Spaces. In *Proceedings of the 2014 Conference on Designing Interactive Systems (DIS '14)*. ACM, New York, NY, USA, 453–462. DOI: http://dx.doi.org/10.1145/2598510.2598533
4. Nathalie Bonnardel and Tamara Sumner. 1996. Supporting evaluation in design. *Acta Psychologica* 91, 3 (1996), 221–244.
5. Richard Buchanan. 1992. Wicked problems in design thinking. *Design issues* 8, 2 (1992), 5–21.
6. Stuart K Card, Jock D Mackinlay, and George G Robertson. 1990. The design space of input devices. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 117–124.
7. 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. *Journal of mechanical design* 133, 8 (2011), 081004.
8. Chaomei Chen. 2006. CiteSpace II: Detecting and visualizing emerging trends and transient patterns in scientific literature. *Journal of the American Society for information Science and Technology* 57, 3 (2006), 359–377.
9. Coggle. 2016. https://coggle.it/. (2016).
10. Christopher Collins, Sheelagh Carpendale, and Gerald Penn. 2009a. Docuburst: Visualizing document content using language structure. In *Computer graphics forum*, Vol. 28. Wiley Online Library, 1039–1046.
11. Christopher Collins, Fernanda B Viegas, and Martin Wattenberg. 2009b. Parallel tag clouds to explore and analyze faceted text corpora. In *Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on*. IEEE, 91–98.
12. Arthur Cropley. 2006. In praise of convergent thinking. *Creativity research journal* 18, 3 (2006), 391–404.
13. Nils Dahlbäck, Arne Jönsson, and Lars Ahrenberg. 1993. Wizard of Oz studies: why and how. In *Proceedings of the 1st international conference on Intelligent user interfaces*. ACM, 193–200.
14. Peter Dalsgaard, Kim Halskov, and Rune Nielsen. 2008a. Maps for design reflection. *Artifact* 2, 3-4 (2008), 176–189.
15. Peter Dalsgaard, Kim Halskov, and Rune Nielsen. 2008b. Towards a design space explorer for media facades. In *Proceedings of the 20th Australasian Conference on Computer-Human Interaction: Designing for Habitus and Habitat*. ACM, 219–226.
16. Alan Dix, Rachel Cowgill, Christina Bashford, Simon McVeigh, and Rupert Ridgewell. 2016. Spreadsheets As User Interfaces. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '16)*. ACM, New York, NY, USA, 192–195. DOI: http://dx.doi.org/10.1145/2909132.2909271
17. Kees Dorst and Nigel Cross. 2001. Creativity in the design process: co-evolution of problem–solution. *Design studies* 22, 5 (2001), 425–437.
18. Niklas Elmqvist, Pierre Dragicevic, and Jean-Daniel Fekete. 2008. Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. *IEEE transactions on Visualization and Computer Graphics* 14, 6 (2008), 1539–1148.
19. EndNote. 2016. http://endnote.com/. (2016).
20. Jennifer English, Marti Hearst, Rashmi Sinha, Kirsten Swearingen, and Ka-Ping Yee. 2002. Hierarchical faceted metadata in site search interfaces. In *CHI'02 Extended Abstracts on Human Factors in Computing Systems*. ACM, 628–639.
21. Microsoft Excel. 2016. https://office.live.com/start/Excel.aspx. (2016).
22. Jean-Daniel Fekete, Jarke J Van Wijk, John T Stasko, and Chris North. 2008. The value of information visualization. In *Information visualization*. Springer, 1–18.
23. FreeMind. 2016. http://freemind.en.softonic.com/. (2016).
24. Vinod Goel and Peter Pirolli. 1992. The structure of design problem spaces. *Cognitive science* 16, 3 (1992), 395–429.
25. TC Graham, Leon A Watts, Gaëlle Calvary, Joëlle Coutaz, Emmanuel Dubois, and Laurence Nigay. 2000. A dimension space for the design of interactive systems within their physical environments. In *Proceedings of the 3rd conference on Designing interactive systems: processes, practices, methods, and techniques*. ACM, 406–416.
26. Tovi Grossman and Daniel Wigdor. 2007. Going Deeper: a Taxonomy of 3D on the Tabletop. In *Horizontal Interactive Human-Computer Systems, 2007. TABLETOP '07. Second Annual IEEE International Workshop on*. 137–144. DOI: http://dx.doi.org/10.1109/TABLETOP.2007.18
27. Joy Paul Guilford. 1956. The structure of intellect. *Psychological bulletin* 53, 4 (1956), 267.

28. Chris Heape. 2007. *The Design Space: the design process as the construction, exploration and expansion of a conceptual space*. Ph.D. Dissertation. SDUSDU, Det Tekniske FakultetFaculty of Engineering, Mads Clausen InstituttetThe Mads Clausen Institute.
29. David G Jansson and Steven M Smith. 1991. Design fixation. *Design studies* 12, 1 (1991), 3–11.
30. Waqas Javed and Niklas Elmqvist. 2012. Exploring the design space of composite visualization. In *Visualization Symposium (PacificVis), 2012 IEEE Pacific*. IEEE, 1–8.
31. Eser Kandogan, Juho Kim, Thomas P. Moran, and Pablo Pedemonte. 2011. How a Freeform Spatial Interface Supports Simple Problem Solving Tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 925–934. DOI: http://dx.doi.org/10.1145/1978942.1979079
32. Rudolf K Keller, Anurag Garg, and Tao Tao. 1993. HyperRef: online support for research literature assessment and documentation. In *Proceedings of the 11th annual international conference on Systems documentation*. ACM, 163–175.
33. Andruid Kerne, Eunyee Koh, Steven M Smith, Andrew Webb, and Blake Dworaczyk. 2008. combinFormation: Mixed-initiative composition of image and text surrogates promotes information discovery. *ACM Transactions on Information Systems (TOIS)* 27, 1 (2008), 5.
34. Andruid Kerne, Andrew M. Webb, Steven M. Smith, Rhema Linder, Nic Lupfer, Yin Qu, Jon Moeller, and Sashikanth Damaraju. 2014. Using Metrics of Curation to Evaluate Information-Based Ideation. *ACM Trans. Comput.-Hum. Interact.* 21, 3, Article 14 (June 2014), 48 pages. DOI:http://dx.doi.org/10.1145/2591677
35. David Kirsh. 2010. Thinking with external representations. *Ai & Society* 25, 4 (2010), 441–454.
36. Giovan Francesco Lanzara and Lars Mathiassen. 1985. Mapping situations within a system development project. *Information & Management* 8, 1 (1985), 3–20.
37. LitAssist. 2016. http://www.litassist.com/. (2016).
38. Allan MacLean, Richard M Young, Victoria ME Bellotti, and Thomas P Moran. 1991. Questions, options, and criteria: Elements of design space analysis. *Human–computer interaction* 6, 3-4 (1991), 201–250.
39. Renan Manola, Renata SS Guizzardi, and Roberta Lima Gomes. 2008. Biblioref: a semantic bibliographic reference management system. In *Companion Proceedings of the XIV Brazilian Symposium on Multimedia and the Web*. ACM, 149–151.
40. Justin Matejka, Tovi Grossman, and George Fitzmaurice. 2012. Citeology: visualizing paper genealogy. In *CHI'12 Extended Abstracts on Human Factors in Computing Systems*. ACM, 181–190.
41. Tony McCaffrey and Lee Spector. 2011. Innovation is built on the obscure: Innovation-enhancing software for uncovering the obscure. In *Proceedings of the 8th ACM Conference on Creativity and Cognition*. ACM, 371–372.
42. Sarnoff Mednick. 1962. The associative basis of the creative process. *Psychological review* 69, 3 (1962), 220.
43. Yevgeniy Medynskiy, Mira Dontcheva, and Steven M Drucker. 2009. Exploring websites through contextual facets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2013–2022.
44. Mendeley. 2016. https://www.mendeley.com/. (2016).
45. MindManager. 2016. https://www.mindjet.com/mindmanager/. (2016).
46. MindMeister. 2016. https://www.mindmeister.com/. (2016).
47. Lior Noy, Yuval Hart, Natalie Andrew, Omer Ramote, Avi Mayo, and Uri Alon. 2012. A quantitative study of creative leaps. In *Proc. Third Int. Conf. Comput. Creativity, Dublin*. 72–76.
48. Charles Perin, Pierre Dragicevic, and Jean-Daniel Fekete. 2014. Revisiting bertin matrices: New interactions for crafting tabular visualizations. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 2082–2091.
49. Peter Pirolli and Stuart Card. 1995. Information foraging in information access environments. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM Press/Addison-Wesley Publishing Co., 51–58.
50. Peter Pirolli and Stuart Card. 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis*, Vol. 5. 2–4.
51. Daniel M Russell, Mark J Stefik, Peter Pirolli, and Stuart K Card. 1993. The cost structure of sensemaking. In *Proceedings of the INTERACT'93 and CHI'93 conference on Human factors in computing systems*. ACM, 269–276.
52. Donald A Schon. 1992. Designing as reflective conversation with the materials of a design situation. *Research in Engineering Design* 3, 3 (1992), 131–147.
53. Google Sheets. 2016. https://www.google.com/sheets/about/. (2016).
54. Harri Siirtola and Kari-Jouko Räihä. 2006. Interacting with parallel coordinates. *Interacting with Computers* 18, 6 (2006), 1278–1309.
55. Herbert A Simon. 1955. A behavioral model of rational choice. *The quarterly journal of economics* (1955), 99–118.

56. Consuelo Valdes, Diana Eastman, Casey Grote, Shantanu Thatte, Orit Shaer, Ali Mazalek, Brygg Ullmer, and Miriam K. Konkel. 2014. Exploring the Design Space of Gestural Interaction with Active Tokens Through User-defined Gestures. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 4107–4116. DOI: http://dx.doi.org/10.1145/2556288.2557373
57. Andrew M Webb and Andruid Kerne. 2011. Integrating implicit structure visualization with authoring promotes ideation. In *Proceedings of the 11th annual international ACM/IEEE joint conference on Digital libraries*. ACM, 203–212.
58. Edward J Wegman and Qiang Luo. 1997. High dimensional clustering using parallel coordinates and the grand tour. *Computing Science and Statistics* 28 (1997), 93–101.
59. XMind. 2016. http://www.xmind.net/. (2016).
60. Ka-Ping Yee, Kirsten Swearingen, Kevin Li, and Marti Hearst. 2003. Faceted metadata for image search and browsing. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 401–408.
61. Zotero. 2016. https://www.zotero.org/. (2016).
62. Fritz Zwicky. 1967. *New methods of thought and procedure: contributions to the Symposium on Methodologies; Pasadena, California, May 22-24, 1967*. Springer.
63. Fritz Zwicky. 1969. *Discovery, Invention, Research: through the morphological approach*. Macmillan.