Why Visualize?
Untangling a Large Network of Arguments

Dirk Streeb, Mennatallah El-Assady, *Student Member, IEEE,* Daniel A. Keim,, *Member, IEEE,*
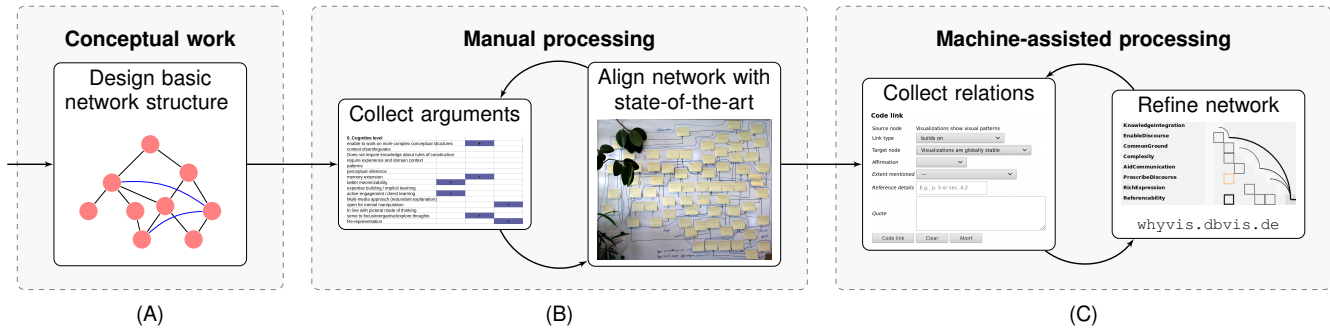and Min Chen, *Member, IEEE.*

Fig. 1: The three major stages in constructing our network of arguments on why visualization works. First, (A) we come up with the idea to look into relations between theoretic arguments, and conceptualize the basic network structure (Sections 2 & 4). Secondly, (B) we collect additional arguments from literature and iteratively align our initial ideas with current research (Section 5). Finally, (C) we implement an interactive system for keeping track of the growing network, make it accessible online for further research, and continue iterative refinement.

Abstract—Visualization has been deemed a useful technique by researchers and practitioners, alike, leaving a trail of arguments behind that reason why visualization works. In addition, examples of misleading usages of visualizations in information communication have occasionally been pointed out. Thus, to contribute to the fundamental understanding of our discipline, we require a comprehensive collection of arguments on "why visualize?" (or "why not?"), untangling the rationale behind positive and negative viewpoints. In this paper, we report a theoretical study to understand the underlying reasons of various arguments; their relationships (e.g., built-on, and conflict); and their respective dependencies on tasks, users, and data. We curated an argumentative network based on a collection of arguments from various fields, including information visualization, cognitive science, psychology, statistics, philosophy, and others. Our work proposes several categorizations for the arguments, and makes their relations explicit. We contribute the first comprehensive and systematic theoretical study of the arguments on visualization. Thereby, we provide a roadmap towards building a foundation for visualization theory and empirical research as well as for practical application in the critique and design of visualizations. In addition, we provide our argumentation network and argument collection online at https://whyvis.dbvis.de, supported by an interactive visualization.

Index Terms—Visualization, Theory, Argument Network, Cognition, Design.

◆

1 INTRODUCTION

VISUALIZATION researchers and practitioners have made attempts to explain to others why visualization works and why it is useful. Benefits of displaying a visualization—e.g., a scatter plot—instead of a table of numeric values may seem obvious to visualization researchers. Articulating precisely which property of the visualization backs up the vaguely-defined benefits is more difficult. The

wide range of positive and negative arguments, such as visual patterns, externalization of memory, and ambiguity of visualizations, poses a fundamental challenge to the field of visualization. Most of us have also heard suggestions that human-centric processes such as visualization have many shortcomings. In the past decades there have been repeated calls for more theoretical advances in the field of visualization, both, by individuals [1], [2] as well as several panels at IEEE VIS conferences [3], [4], [5], [6]. Most recently Chen et al. advocated that "[m]aking significant theoretical advances will lead to significant advances in practical visualization applications." [6, p. 111] This echoes Thomas' assertion that "without fundamental knowledge of what makes certain representations effective, it is not possible to efficiently construct new representations for new classes of information or to know that the new representations will work as designed." [7, p. 70]

While these calls for action have been followed by many

- *Dirk Streeb, Mennatallah El-Assady, and Daniel A. Keim are with the Group of Data Analysis and Visualization at the University of Konstanz, Germany. E-mail: {streeb,elassady,keim}@dbvis.inf.uni-konstanz.de*
- *Min Chen is with the Department of Engineering Science at the University of Oxford, UK. E-mail: min.chen@oerc.ox.ac.uk*

attempts to establish a set of principal arguments, most are of qualitative discourses and a few are mathematical (see Section 3 for details). Investigating aspects one by one increased the comprehension of specific arguments for visualization—most prominently visual patterns. However, connections between individual arguments usually remain implicit. There is no explication of the broad argumentation network. Still, it is necessary to get an overview perspective. Including arguments' relationships is key to gaining a better understanding of why visualization works. For instance, the opportunities that visual patterns offer can only be explained by referring to additional arguments on perception, accessibility of data, and interaction of the user with the visualization (see also Figure 5 on page 9). These relationships that are only implicit in the state-of-the art motivated us to begin with building a network of arguments that entails arguments' dependencies as well as needs for trade-off.

Further, we see a lack of research at a medium level of abstraction. On the one hand highly abstract work looks at information visualization from formal or mathematical perspectives; regularly treating visualizations themselves as black boxes. On the other hand a lot of work is particularly focused on detailed aspects, and often backed up by psychological theories and empirical investigations. Meanwhile, these perspectives are difficult to put into practice. While highly abstract results need to be translated to applicable instructions or guidelines, detailed theoretical and empirical outcomes often lack advice for mitigating conflicts or making trade-offs when conflicts emerge in practical application.

In our network, we connect findings that offer more details than black-box approaches. At the same time, the network abstracts from particular application scenarios, which is, for example, relevant for generalizable empirical research. As a result, the network constitutes a theoratical roadmap allowing for diverse applications (see Sections 7 & 8). While it offers an overview on the conceptual level, it does not entail details on how particular visualizations work or how to exploit the potentials of visualization most effectively. We consider the extraction of guidance on how to design and optimize visualizations for specific demands to be future work. Pictorially speaking, with the network we initiate the mapping of the terrain of visualization theory. Constructing common pathways and routing travelers through the complex landscape is at least one step ahead.

In this paper, we report our theoretical investigations on why visualization works, and present the resultant findings:

- a network of theoretic arguments on why visualization works. We derived the network from theoretical considerations based on the individual arguments' contents. The network connects arguments explicitly and provides a roadmap that aids in structuring discourse (Section 4).
- the detailed collection of arguments underlying the network (see supplementary material). Specifically, we introduce several categorization schemes—based on argumentative standpoint, focal pathway of information flow, as well as dependency on task, user(s) and data—which provide additional structure to the set of arguments (Section 5).

- the presentation of theoretical results on an interactive website. At https://whyvis.dbvis.de/ we feature the argument network in an interactive visualization. We provide relevant references, and quotes facilitate the interpretation of our formulations of arguments. Both, the network and the set of arguments, can be downloaded.
- an outlook on how the network approach and our initial findings (Section 6) open up new opportunities for future research. Moreover, we sketch how theory can be put into practice of critique, automatic recommendation and design of visualizations (Sections 7 & 8).

The remainder of the paper is organized as follows: After defining relevant concepts—like *visualization* and *argument*—in Section 2, we present a brief overview of related work in Section 3. Then, we describe the construction of the argument network based on various relationships among arguments about why visualization works in Section 4. In Section 5, we describe the set of arguments, which form the nodes of our argument network. Right after that, we introduce three categorization schemes that provide additional theoretic structure. First, we categorize arguments along the argumentative standpoint of expression (Section 5.1). Secondly, we look at the focal pathways of information flow through the visualization pipeline that arguments emphasize (Section 5.2). Finally, we consider the dependency of arguments on the visualization task, user(s) and the data at hand (Section 5.3). Having concluded the theoretical parts in Sections 4 and 5, we point out how our work can be used in Section 6. In Section 7, we show in detail how the network may be part in the progress of visualization theory. Additionally, we highlight more potential applications in Section 8. For example, we sketch implications for empirical research (Section 8.1), look at the application of our approach to qualitative and quantitative evaluation of visualization designs (8.2), and finally outline how the network may help at designing visualizations and provide theoretical foundations for design guidelines (8.3). In the end, we sum up in Section 9.

2 TERMINOLOGY AND SCOPE

Before going into details on the argument network, we need to explain two central definitions. These are the interpretations of the terms *visulization* (Section 2.1) and *argument* (2.2). To complete this section we briefly discuss a parallel with an example from biology to show the appropriateness of our definitions within the scope of this paper (2.3).

2.1 What is Visualization?

Of the many distinct interpretations of the term "visualization" two are potential objects to a theory of information visualization: i) the process of visualization, often described by the InfoVis pipeline [8], and ii) the physical object, which is the output of the InfoVis pipeline. While these two interpretations often are considered in combination, we focus solely on the latter. We do not consider other meanings, like mental imagery as quasi-perceptual experience [9].

We focus on visualizations as objects as we try to understand these objects better. While process-driven theoretic work often includes some restrictions on which visual objects are desirable, e.g., by introducing soundness

criteria, they generally lack detail on which of the many sound outputs are preferable, and for what reasons. Detailed inspection of the properties of sound visualizations can provide arguments for the preference of one visualization over another in a given situation. Although the pipeline view has attracted much attention and its outcomes have reached the commercial sector, scientific attention to the output of the pipeline has been comparably low—with the exception of highlighting bad examples.

In this paper, we consider an inclusive definition of visualization. A first working definition might be that any physical—digital or analog—visual object that has been created purposefully for some kind of analysis of data is a visualization. Yet, this definition is too loose in the sense that it also includes printed text, which we explicitly want to exclude. Hence, we demand that some visual feature needs to be meaningful beyond type of sign, and order of objects. For example, a table is a visualization as the position of a cell is meaningful in the sense that it assigns the cell to a column and a row. However, in general text is not, as line breaks are arbitrary, and lines need not be bound to exact horizontal positions as long as the order of symbols is well defined—consider rotated text, text along the border of round logos, and also top-to-bottom systems as in traditional Japanese. For our purpose such loose definition is sufficient as we do not expect to compile a set of properties that all visualizations share. Instead, we collect a network of arguments of which only subsets are put forward for individual visualizations. In fact, several of the arguments could be plausibly put forward for non-visualizations such as printed text. Depending on the task a visualization is intended for, only some of these arguments are relevant.

2.2 What are Arguments?

Arguments are a general and common form of expression used across disciplines involved in the study of visualization, for instance: The fact that X contributes to / reduces the benefit of visualization; in context of Task T, User(s) U and Data D. Practically, we take the conclusion, that visualization works, for granted and put the premises X at the center of our investigations.

As these arguments often lack complete context, we group utterances of similar semantic content. While this procedure may lead to grouping utterances that do not match perfectly, we argue that arguments are rarely put forward with such rigor that it would be possible to read off all details without adding subjective interpretation. Our (grouped) arguments preserve the essence of individual arguments, and abstract from often vague or implicit details. Figure 2 gives an example of how we present the arguments in this paper. A complete list of the arguments we considered can be found in the supplementary material.

2.3 Why investigate Arguments?

As starting with the loose definition of visualization presented in Section 2.1 and the focus on arguments described in Section 2.2 may not be self-explaining, let us argue why it is suitable and productive by looking at a counterpart example from biology. "*Why do species survive?*" may be a question a biologist is interested in. In a sense it is similar to
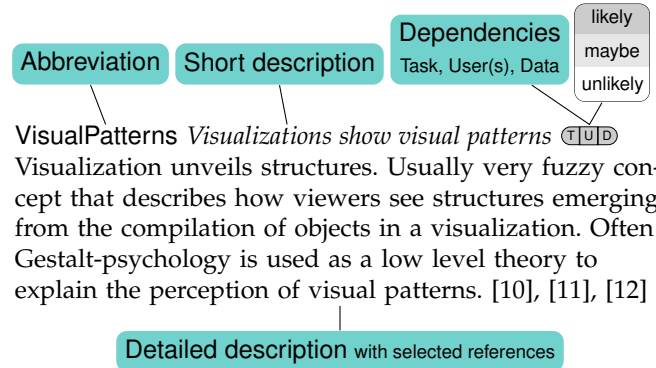

Fig. 2: Argument on VisualPatterns as an example for how we present the arguments in this paper.

our question of "*Why do visualizations work?*". *Visualizations* are comparable to *species* and the overall goal of gaining *insight* is analogous to *survival*. While we expect everyone to agree that describing how one particular species survives requires detailed knowledge of that species, more general findings can be made at a conceptual level (Why?). For example, whales, sharks and penguins survive, in part due to the fact that they are good at swimming. Details on the exact movements they make are not relevant for making this argument. Similarly, bar charts and line charts are reasonable choices when it comes to displaying the change of a numeric dimension over time. Only if more detailed tasks are considered—like comparing values at points in time, or estimating trends; respectively feeding on plankton, squid or fish—low-level properties make a difference.

Very abstract points of view on the other hand make it difficult to distinguish between individual species/visualizations. Stating that living species have some kind of metabolism may be correct and surely is relevant, but at the same time does not provide the means to distinguish between species. Comparably, an information-theoretic view as presented by Chen, Floridi and Borgo [13] can easily distinguish human-centered visualization from machine-centered processes, but it can hardly distinguish between visualizations that represent equivalent information—such as a bar chart and a line chart.

Thus, we choose a medium level of abstraction (i.e., arguments) for investigating why visualization works. Some level of abstraction from Task, User(s) and Data is key to gain generalizable theoretical results. Not abstracting too much is key to distinguish individual visualizations. At the same time, we are well aware that this level of abstraction also provides drawbacks. For instance, it requires a more complex representation and misses distinction on low-level details. Thus, we link our level of abstraction to higher and lower levels. In Section 5.2 we connect arguments to higher level pathways of information flow, and in Section 5.3 we categorize them based on their likelihood to be affected by lower level considerations, namely Task, User(s) and Data.

3 RELATED WORK

Based on the distinctions introduced in Section 2, we introduce related work grouped in four categories: highly

abstract, low-level-focus, pipeline focus, and works with a focus similar to ours.

Highly abstract: A substantial amount of related work is highly abstract. For example, van Wijk [14] looks at visualization as a black box in his process model. More recently Vickers, Faith and Rossiter [15] conceptualized visualization from a semiotic point of view using category theory. Chen, Floridi and Borgo [13] focus on the economic side of visualizations as time saving tools. Chen and Jänicke [16] as well as Chen and Golan [17] take an information-theoretic perspective and abstract visualizations to information transport channels. For visual analytics Sacha et al. [18] present an abstract knowledge generation model that includes the information visualization pipeline [8] next to an abstraction of human interaction.

Low-level focus: Low-level research on single arguments is diverse and spread across disciplines. We highlight few examples that bring out the interdisciplinary nature of visualization research: The fundamentals of color perception are well researched [19], [20]. The InfoVis community focuses on effects of color scales and compositions [21], [22], [23]. Other researchers take their intuitions to the test and, for instance, focus on the effects of different levels of iconicity of symbols [24], [25], [26]. Although some effects of iconicity seem to be intuitive, often researchers were not able to find these effects in their experiments. This clearly shows the need for low-level experimental research on visualization. More practically, effects of visualizations in the context of Bayesian reasoning problems gained interest in the InfoVis community in recent years [27], [28]. Traditionally, mainly psychologists and cognitive scientists investigated this topic [29], [30], [31].

Pipeline focus: Investigating the process how to map data to visual representations has gained repeated attention. The information visualization pipeline by Card, Mackinlay and Shneiderman [8] is well known. Wilkinson published investigations on possible operations in his "The grammar of graphics" [32], on top of which not only Wickham [33] builds. Demiralp et al. [34] present similar theoretic ideas of embedding structures within data in visual representations. In contrast to this paper, the focus of these works is on how to create sound visualizations from data, but not on why the value of some sound visual representation varies between different application contexts. In our network we expect most arguments put forward from this perspective to add detail to the Basic standpoint (see Section 5.1) as well as the second pathway of information flow (5.2). However, our work does not focus on this part of the visualization process, and detailed inclusion of all such arguments would distract from the main points discussed in this paper.

Similar focus: To the best of our knowledge recent investigations on why visualization provides benefits can be dated back to Anscombe [10] in the context of statistics. Larkin and Simon [35] investigate structural diagrams. More recently, Fekete et al. [36] discuss the value of information visualization. Another approach has been taken by Parsons and Sedig [37] who compiled a list of ten properties of visual representations that, at different levels, may explain the appropriateness of a visual representation for a given task. Kindlmann and Scheidegger [38] add three property requirements to be satisfied by effective data visualizations.

While these works present interesting arguments for why visualization is useful in many applications, they largely enumerate arguments independently, and miss to point out the connections in the larger argumentative structure.

4 TOWARDS A NETWORK OF ARGUMENTS

Our main contribution is the argumentation network constructed from arguments we collected across disciplines. As a short disclaimer, we do not consider the network complete. We introduce it as a new way of looking at the arguments from an integrated network perspective, instead of as single entities standing for themselves (see Figure 1 (A)). Further, we hope that the point of view presented in this paper can serve as a common ground for discussing and promoting theoretic approaches within—and maybe even beyond—the information visualization community. Having said that, we want to encourage researchers to contribute to the network. An interactive version of the network, including quotes from and references to those publications presenting the arguments is available online.

4.1 Network Construction & Methodology

We collect arguments by manually extracting them from literature. At first, we consider highly influential publications (A) such as those by Larkin and Simon [35] and Card, Mackinlay and Shneiderman [8]. In the following we continue with papers citing these highly influential papers (B). We then include relevant papers (C) referenced by (B) next to the highly influential ones (A). In this manner we select papers by following references in both directions. We also search for additional works on specific arguments using relevant keywords. Thereby, we focus on aspects of Design and Cognition (see Section 5.1 and Table 1).

While we use a small set of arguments as a starting point (see Figure 1 (A)), we soon begin to collect arguments in a spreadsheet (Figure 1 (B)). Each publication can mention several arguments, and each mentioning can be of one affirmation: ++ Central positive (24 mentions), + Positive (726), +− Ambivalent (44), = Neutral (22), − Negative (28), −− Limiting negative (0), ∗ Mentioned without valuation (25), ? Questioned (7). Though, one has to notice that quite some negative arguments you might have heard about specific visualizations—not visualization in general—are the result of a lack of some positive aspects. Still, many arguments are far from being unanimously seen positive, for example, Animation. Therefore, it can be the case that an argument is described negatively by some publications.

Then, we construct the network in an iterative process. The whiteboard in Figure 1 (B) shows one temporary state of the network. We introduce most of the relationships based on our semantic understanding of the arguments. Only after we established large parts of the network, we begin collecting relationships in parallel to arguments (Figure 1 (C)). Obviously, this procedure leaves us with less mentioned relationships, 49 compared to 933 argument mentions. Taking into account that the latter publications contained about 42% (396) of the node mentions next to mentioning these relationships, we feel rather safe in stating that the connections between arguments have not been as prominent in

the literature in general. This aligns with our impressions gained from reading the first set of publications, and is also reflected in our claimed contribution to provide the first work with a strong focus on the relationships among arguments about why visualization works. We pursue this two-step exploratory approach for three reasons. To begin with, the initial open coding phase allowed us to align our theoretical efforts with previous work. Secondly, we want to derive the network from theoretical considerations based on the arguments' contents, and not as a result of a survey-like effort. Finally, we include relations mentioned in literature for further refinement of the network.

4.2 Argument Relations

Currently, the network is built upon six types of relations between arguments A and B. The most important relation is the *builds on* relation. This *builds on* relation is asymmetric and, if all argumentation is sound, leads to an acyclic network graph. From a logical point of view the relation is weak in the sense that X *builds on* Y does neither imply that Y is sufficient for X nor that X is necessary for Y—nor the opposite directions. We neither make any claims about the sufficiency or necessity of combinations of propositions for others, e.g., X *builds on* Y and Z. Instead, our goal is to compose the loose network in order to lay the groundwork for future discussions. To date, we take the following relations into account:

builds on A can be expressed in a way such that B is a premise of A. This relation is logically weak in the sense that it does not imply necessity nor sufficiency.

conflicts with A and B cannot be the case at the same time, i.e., they are contrary.

is limited by B puts a limit on the benefits of A.

is more specific than A makes a more detailed point than B.

is similar to A and B make points that are close, but include a meaningful difference.

needs to be traded off against A and B can not be the case to full extent at once, but both can apply partially.

For example, there are people who argue that Animation can make visualizations work. Others put forward that an important aspect of visualizations is that they are globally stable (i.e., GlobalStability). Clearly, a visualization cannot be entirely stable and animated at the same time. Hence, Animation *needs to be traded-off against* GlobalStability.

Conflicts and needs for trade-off do not necessarily materialize in practice as we do not expect any visualization to have all properties mentioned as beneficial in any argument. Distinct visualizations may have different properties. The relevance of each single argument for the usefulness of a particular visualization depends on the application scenario.

When inserting relationships, we refrained from adding those implied by the transitivity of the *builds on* relation in most cases (i.e., if A *builds on* B and B *builds on* C, we do not add A *builds on* C). Two exceptions to this general guideline were: i) A relationship was mentioned explicitly in a publication, and ii) we reasoned that the direct relation highlights a central connection. With relations other than *builds on* we did not encounter such cases. In the following section we present some details on the set of arguments we considered when constructing the network.
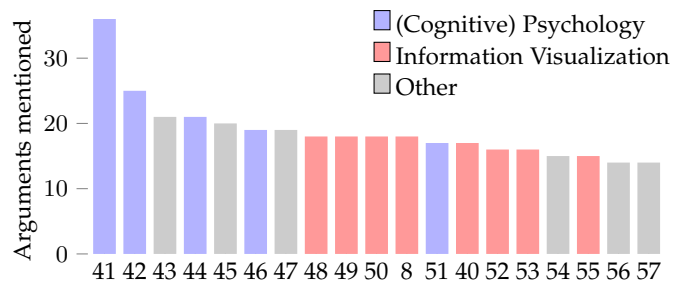

Fig. 3: Nineteen publications mentioning the most arguments. Bars represent publications and are labeled with their identifying numbers according to the list of references of this paper. Notably, many highly diverse publications are not from the InfoVis community.

5 ARGUMENT COLLECTION AND CLASSIFICATION

After coding arguments from an initial set of 57 publications in a spreadsheet in the first phase, we migrated to an online tool, which we specifically developed for the task of constructing an argumentative network. With the help of this tool we were additionally able to track mentioned relationships between arguments. In total, we did 112 codings of 108 publications featuring 933 mentions of arguments and 49 mentions of relationships between arguments.

Overall, we do not claim that our sample is complete. We did not try to provide an exhaustive survey of past work. In fact we prioritized recent work as we expect recent publications to pick up important arguments that have been brought up earlier. Also, we did not undertake any effort to gather a representative sample regarding the positive or negative mentioning of arguments as only the single first mentioning of an argument is of high importance for our goal of building a network featuring as many arguments as possible within reason. Every additional mentioning of an argument can only contribute to the refinement of the formulation of the argument. Our goal is to collect a diverse set of arguments.

We agree with the views that investigating visualization is an integrated and interdisciplinary matter [39], [40]. Therefore, we considerd publications from several scientific domains. Figure 3 shows those publications that mentioned a diverse set of many arguments. It is easy to see that looking beyond the InfoVis community can broaden our view. In our pursuit to create a network relating arguments these diverse publications are crucial.

Having collected arguments as described above, it soon became clear to us that we need to provide more structure. We propose several categorizations of the arguments. First, we introduce a scheme based on argumentative standpoints in Section 5.1. Secondly, we offer a structuring dependent on the focal pathways of information flow through the InfoVis pipeline an argument focuses on in Section 5.2. Both schemes should be considered as navigational aids in the space of arguments. They do not provide clear cut distinctions in all cases, and arguments regularly span multiple categories. Finally, we categorize arguments by their dependency on Task, User(s) and Data in Section 5.3.

TABLE 1: Overview of the standpoints. n is the number of arguments per argumentative standpoint. Our focus on Cognition and Design becomes obvious.

Standpoint	Short description	n
Teamwork	Arguments about more than one viewer	8
Cognition	Arguments involving a viewer with higher cognitive abilities	31
Interplay	Arguments emerging from the interplay of viewers and visualizations (including perception and interaction)	19
Design	Arguments depending on design choices	32
Semiotics	Arguments on meaningfulness of symbols and space	19
Basic	General mathematical and physical arguments mostly independent of visualization	10

5.1 Argumentative Standpoints

Looking through the collected arguments, we figured that they focus on several aspects of visualization. Overall we found six argumentative standpoints: Basic, Semiotics, Design, Interplay, Cognition, and Teamwork. Especially the borderline between the Cognition and the Interplay standpoints is fluid [1], [58], [59]. Nonetheless, the standpoints can help organize numerous arguments. The structure becomes practically relevant as standpoints hint which options for improvement of visualizations users—designers as well as viewers—might have. Next, we describe the standpoints in more detail (see also Table 1), provide exemplar arguments, and point out how visualizations might be improved.

Basic: The most general standpoint groups mathematical and physical arguments. The distinctive criterion for this standpoint is that arguments are not focused on visualization as such. Instead, their main focus is on general mathematical or logical laws, or on the properties of the physical world. Examples for such arguments are:

TopologyPreservation *Visualizations preserve topology* ⓉⓊⒹ Visualizations can preserve the topological (or geometric) relations of a two-dimensional space that would need to be broken by a one dimensional representation. [35], [38]

StorageInvariance *Visualizations are invariant to data storage formats* ⓉⓊⒹ Some storage formats of data do not represent all features of the abstract structure faithfully. For example, a set stored as a list implies an order that is not present in the abstract set structure. Good visual representations do not reproduce these misrepresentations of underlying data storage formats. [38]

Given the generality of arguments options for improvement are limited. Laws of mathematics and nature can be considered immutable. Technological progress is the main angle for improvement, and can lead to higher quality media and better visualizations. A serious caveat to keep in consideration is that logical and mathematical demands on faithful mappings, in contrast to physical limitations, are not self-enforcing. For example, theoretically there is no addition operation on ordinal scales. Nonetheless, designers are free to add ordinal values represented as numbers, creating visualizations that are not mathematically sound.

Semiotics: Having set the logical and physical preconditions, the Semiotic standpoint pools arguments that build on how meaning is assigned to objects. The main arguments relevant to visualization are that, both, symbols and space are assigned meaning, others include:

FlexibleNotation *Visualizations allow flexible notation* ⓉⓊⒹ Meaning of symbols and space can be changed on demand, it does not need to be defined globally unique. For example, vertical position can have a different meaning at each axis of an parallel coordinate plot, while having no meaning in between the axes. There even can be no explicit definition of the meaning of space. For example, in interactive environments users can sort objects by positioning them in space. With each new object they can refine or change the meaning of positions, while never having an explicit mapping. [41], [60], [61], [62], [63]

Homogeniety *Visual entities are homogeneous* ⓉⓊⒹ Visualizations map entities of different types to visual representations, which are of the same kind and thus easier to put in relation to each other. For example, a bar chart on the speed of light, sound in air, a plane, a horse, and the growth of bamboo shows distinct objects like electromagnetic waves, pressure waves, human artifacts, animals, and size change of a plant as representations of the same kind, i.e., rectangles. [53], [54]

The Semiotic standpoint, much like the Basic standpoint, includes many unchangeable preconditions. However, there still are some options for intervention. For example, cultural communities such as electrical engineers can introduce new conventions for meanings of symbols or colors.

Design: Arguments from the Design standpoint are most interesting for visualization designers as they provide most leverage for improvement in practical application. Properties of visualizations themselves are of central interest from this point of view:

Abstraction *Visualizations abstract structures* ⓉⓊⒹ Visualizations abstract from (irrelevant) details. [42], [64], [65], [66]

Animation *Visualizations are animated* ⓉⓊⒹ An animated visualization can show processes with a natural representation of time. [2], [37], [42], [67]

GlobalStability *Visualizations are globally stable* ⓉⓊⒹ The overall appearance of a visualization does not change without explicit action undertaken by a viewer. Thereby, viewers can on the one hand stay in sync with the external representation and on the other hand have the time to follow their flow of thoughts. [8], [37], [47]

Design decisions provide many options for improving visualizations. However, some of the properties conflict or need to be traded-off against each other. Considering the last two examples, it becomes clear that Animation and GlobalStability are contrary. One option to trade these off would be to offer two views, one animated and one static, side by side. The cost of this trade-off is that each view needs to be smaller and looses detail. The major limiting factor is the task the visualization is designed for. Designers need to decide which properties should be part of the design in order to effectively pursue the objective. Further, interaction capabilities and cognitive abilities of viewers limit the freedom of designers.

Interplay: The Interplay standpoint combines two aspects that are traditionally separated in the InfoVis community: on the one hand arguments on the perception of visualizations, and on the other hand arguments dealing

with actively changing them. The latter aspect becomes more relevant with digital "interactive visualizations" that, for example, allow linking and brushing. Still, most traditional printed or drawn visualizations are interactive to some degree, too [68]. A visualization printed on paper can be folded or annotated with a pen. We combine the two aspects to one standpoint of Interplay as with modern provenance tracking, which may track every move of the eye and cause changes in the visualization, the borderline between the two becomes more and more washed out. Furthermore, for example, Liu, Nersessian and Stasko [58, sec. 6] provide theoretical reasons for assessing interaction bidirectionally. The standpoint of Interplay groups all arguments that directly involve the visualization and the viewer, who does not necessarily employ higher cognitive skills. Examples of arguments are:

AttentionDriving *Visualizations drive attention* ⓉⓊⒹ A visualization guides the attention of viewers to specific parts without the need for conscious search or understanding. [35], [69], [70]

PerceptualInference *Perception infers abstract structures* ⓉⓊⒹ Viewers can assimilate information given in a visualization without a need for heavy conscious processing. Perception substitutes intentional reasoning about abstract structures. [11], [42], [64], [70], [71]

Interaction *Visualizations are open for interaction* ⓉⓊⒹ Visualization enables interactive manipulation. Subjects can physically modify a visualization. This can be annotation, folding the piece of paper the visualization is printed on, or any kind of interaction with a digital visualization on a computer screen. [14], [17], [49], [62], [67], [72], [73], [74]

Improving interplay most likely can come about by training on the human side. Especially with computer-aided interactive visualizations, interaction design can be improved on the side of visualizations. While visual perception can be specifically trained to see certain patterns, manual interaction might be improved by training skills in using input devices such as computer mice.

Cognition: The distinction between the Cognition and Interplay standpoint is the least crisp. There also have been arguments for the non-separability of cognition and interaction [58]. While the Interplay standpoint, as described above, groups arguments that directly involve the visualization and the viewer, this standpoint collects arguments that are mainly on the (complex) cognitive features of viewers. Almost all arguments from the Cognition standpoint are independent of the perceptual channel being vision, audition or some combination. Another distinctive criterion might be that the Cognition standpoint puts a stronger focus on arguments that involve conscious reasoning. Although this criterion is not sufficient, as some parts of perception, and clearly most manual interaction, is conscious. Arguments from the Cognition standpoint include:

ExternalMemory *Visualization supports external memorization* ⓉⓊⒹ Visualizations extend memory. Visualizations can provide storage space for information viewers are not able to or do not want to remember by themselves. [8], [13], [35], [36]

Forgiveness *Visualizations are forgiving* ⓉⓊⒹ Visualizations are not designed to be eternal, but known to be imperfect

and intended to be improved over time. [44], [61], [68], [75]

Rerepresentation *Visualizations re-represent structures* ⓉⓊⒹ Consulting visualizations viewers find other representations of a problem which might make its solution more obvious or accessible. Replacing mental models is facilitated by interacting with alternative external representations. [1], [25], [47], [58], [76]

IdeasByAmbiguity *Ambiguity/inconsistency catalyses ideas* ⓉⓊⒹ A visualization does not provide a single complete and unambiguous interpretation. Inconsistencies within an interpretation or between possible interpretations can push viewers to generate new ideas. [47], [51], [68]

Similar to the Interplay standpoint, with Cognition improvements may primarily come about by training. Educating viewers in analytical reasoning or practicing suitable decision heuristics are two options.

Teamwork: Finally, the Teamwork standpoint groups all arguments that involve more than one viewer. Likewise, the distinction to other standpoints is rather crisp and arguments from this standpoint can be ignored when designing visualizations for exclusively personal use. The main arguments here are on presenting information to others and organizing collective effort. Hence "Collaboration" would have been the most suitable name for this category. In order to have highly distinctive names—compared to "Cognition"—we opted for "Teamwork". Some arguments from this standpoint are:

Referencability *Visualizations are objects to reference* ⓉⓊⒹ Being physical objects visualizations can be used by viewers to reference parts of the structure explicitly. [40], [41], [47], [61]

KnowledgeIntegration *Visualizations support knowledge integration* ⓉⓊⒹ Visualizations can help several viewers to combine their knowledge and draw new conclusions/insights from their combined knowledge. [61]

Improvements can be made by changing the way viewers interact with each other and the visualization. Especially interactive visualizations can enable—or even enforce—interaction protocols that promote positive outcomes.

5.2 Focal Pathways of Information Flow

We also noticed that some arguments focus on the transformation of data to visual representations without mentioning human viewers. Others focus on cognitive processes during understanding visualizations without mentioning data. We conjectured that different arguments might emphasize different pathways in a visualization workflow. This leads to our second categorization scheme based on *focal pathways*. Chen and Golan [17] reviewed many pipelines that have been proposed to describe visualization workflows (e.g., [8], [14], [78]). In this work, we use the pipeline by Chen and Floridi [77] because it encompasses most commonly-seen illustrations of pipelines as well as the seldom-seen aspect of human-human communication. Figure 4 shows six major pathways superimposed on that pipeline:

① Arguments focused on information transformation in individual steps as well as their compositions. Unlike the remaining pathways, ②–⑥, it is more generic but less specific about visualization.
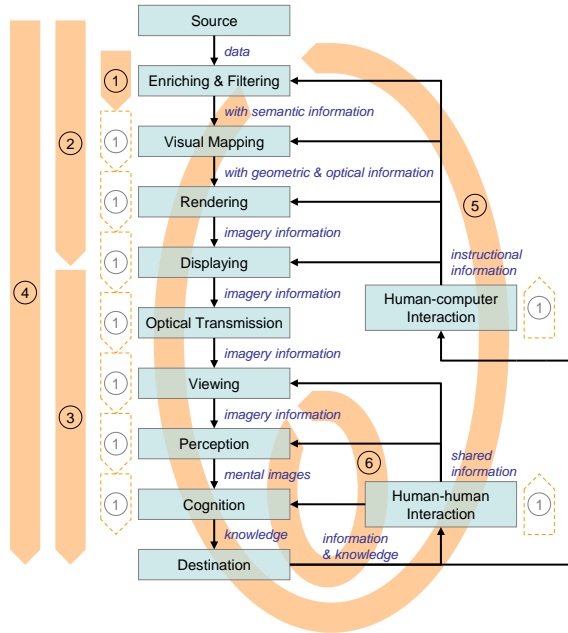
Fig. 4: Six major pathways of information flow on which arguments may have been focused. Pathway ① is an abstraction of single steps. The figure is adapted from Chen and Floridi [77, fig. 1].

② Arguments on transformations from data to visual images. This category also includes arguments about visual mapping, a smaller section in the pathway (e.g., Card et al.'s pipeline [8]). Many arguments from standpoint categories Semiotic and Design fall into pathway ②.

③ Arguments focusing on the transformation of visual images to knowledge. This pathway is emphasized by most arguments in the Cognition standpoint category.

④ This pathway is the combination of ② and ③, focusing on the transformation of data to knowledge via visualization. Arguments in this category typically mention something about data, visualization, and mind. The visual exploration path in the pipeline by Keim et al. [78] is such a pathway.

⑤ Arguments on the values of interaction in visualization. The loop in the pipeline by van Wijk [14] is such a pathway.

⑥ Arguments focusing on uses of visualization involving multiple viewers, such as discriminative visualization and collaborative visualization.

As this work revolves around arguments about visualization rather than analytical statistics or algorithms, we do not include the data mining path from Keim et al. [78]. Nevertheless, if necessary we can associate such an argument with pathway ①.

5.3 Dependencies on Task, User(s) and Data

Last but not least, we categorize arguments by their dependencies on *Task*, *User(s)* and *Data*. Here we chose a three point ordinal scale for each dimension, including the values *Likely*, *Maybe*, and *Unlikely* as with most arguments a definite decision is not possible. While the existence of edge cases prevents binary yes/no decisions, the content of arguments usually allows for getting a general impression.

Task ⊤∪D: Obviously, the task being pursued strongly influences the demands put on visualizations. Consequently, the majority of arguments (79) is likely dependent on the task. Nonetheless, there are several aspects, such as the preconditions set by physics and the biology of the human visual system, that are unlikely to be dependent on tasks (33); maybe (7).

User ⊤∪D: While the role of viewers in using visualizations is widely acknowledged, much effort is put into automatizing large parts of the visualization process. Categorizing arguments by user-dependency allows to separate those aspects largely independent of viewers, and to consider these for automatic generation and testing. We classified 68 arguments as likely user-dependent, 49 unlikely, and 2 maybe.

Data ⊤∪D: For information visualization existing data and its structure are central driving forces to the design of visual representations. However, visualizations have properties that are independent of existing data. In total, we consider 63 arguments as likely data-dependent, 39 as unlikely, and 17 maybe.

6 OBSERVATIONS AND PRELIMINARY FINDINGS

To give some examples for considerations that can be undertaken using the argumentation network, first, have a look at Figure 5. There, we present the argument that visualizations show visual patterns (VisualPatterns), which is the most prominent in our collection, next to its supporting background. The partial network is induced by the arguments from the Interplay standpoint (see Section 5.1) and the *builds on* relation. It clearly shows that VisualPatterns is based on several arguments about perception (shown in red). More interestingly, the argument also draws from arguments on the accessibility of data (blue), and the possibility to interact and learn (green). Whereas the importance of a connection between visual patterns and data is obvious for information visualization once stated, interaction and learning only recently gain increasing interest.

One focus of our investigations are the arguments uttered on cognitive aspects (see also "Cognition" in Section 5.1). As expected, many arguments from this standpoint are about thinking (purple in Fig. 6) or related to memory (blue). Also not unexpected are arguments on evaluating data errors and false assumptions (green). The set of arguments building on the forgiving nature of visualizations (red) is larger than we expected. These arguments include some benefits of Ambiguity as well as the possibility to change views and engage in counterfactual reasoning.

Our second focus is on design options (see also "Design" in Section 5.1). Figure 7 depicts the induced graph of arguments based on the *builds on* relation and the arguments from a Design standpoint. Many arguments are on the general composition (green) of visualizations and composition in the context of Task and Data (purple). Another central part of argumentation is on Abstraction and the—often undesired—induced Ambiguity (blue). Recently, anticipating viewers' perception and thinking (red) has gained interest, especially in pursuit of counteracting against presumed [79]
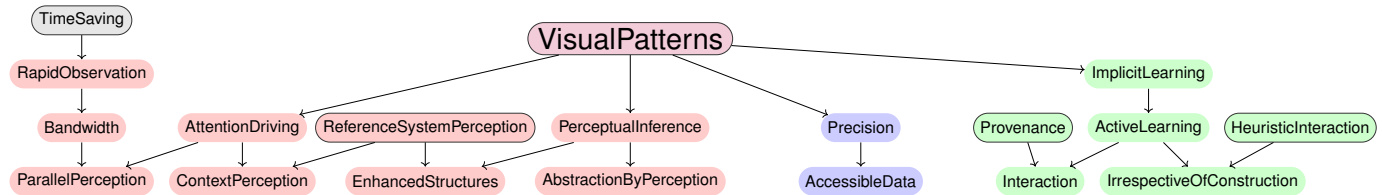
Fig. 5: Induced subgraph of the *depends on* relation for all arguments from the Interplay standpoint. The argument on VisualPatterns builds on aspects from perception, accessibility of data, and interaction; denoted by color. Arguments on which no other arguments from this standpoint build upon are shown with solid borderlines.
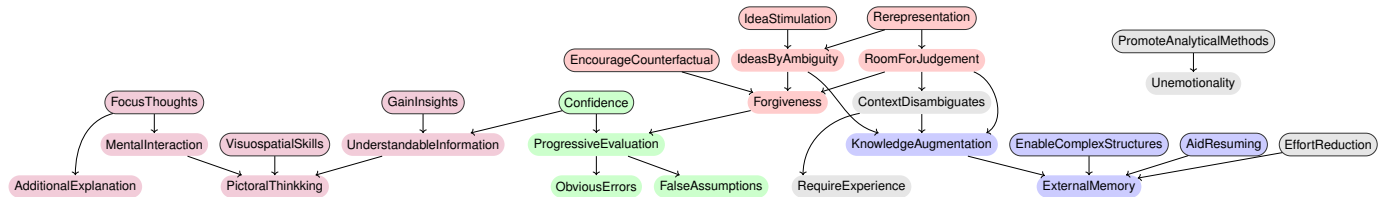

Fig. 6: Induced subgraph of the *depends on* relation for all arguments from the Cognition standpoint. Arguments are grouped by the broader theme they relate to. Purple arguments are on thinking, green ones on evaluation, red arguments build on the forgiving nature of visualizations, the blue ones deal with memory, and gray arguments are on different topics. Arguments on which no other arguments from this standpoint build upon are shown with solid borderlines.

cognitive biases (OvercomeBiases). While many arguments on the design of visualizations are general and independent of viewers, this new trend of designing visualizations for individual viewers puts high demands on designers and their knowledge about future viewers.

As pointed out above, several pairs of arguments go in different directions. In practice, the properties underlying these arguments need to be traded-off. For example, the value of Animation (orange part of Figure 7) continuouesly raised attention. Clearly, Animation *needs to be traded-off against* GlobalStability. The argument network allows us to see both arguments in context, and to estimate some potential effects trade-offs will have by looking at the arguments that *build on* either one. Other examples of relevant trade-offs are between Discreteness and Precision, and between Specificity and Flexibility. We expect that predicting effects of trade-offs in a structured way can inform and guide design decisions in the future.

With regard to dependencies on Task, User(s), and Data, we expected dependencies to accumulate. Figure 8 provides an overview on the dependencies of all arguments. As expected, most general arguments from the Basic standpoint are unlikely user-dependent.

Overall, the current state of the arguments is diverse. There are very prominent arguments as shown in Figure 9. Despite their prominence, arguments usually are not laid out rigorously such that detailed aspects often remain unclear. With our network, we provide a basis for a more precise investigation of arguments, and especially allow to disentangle the dependencies between them. At the same time, there are quite some arguments which are only expressed rarely. For example, Stapleton, Jamnik, and Shimojima [80] are the only ones in our sample who point out that parts of visualizations can be meaningful for themselves (i.e., MeaningfulParts). Kindlmann and Scheidegger [38] argue that visualization works as it does not repeat misrepresentations accepted in data storage formats (i.e., Stor-

ageInvariance). In addition, Patterson [70] highlights that the visual context influences how individual objects are perceived (i.e., ContextPerception). These aspects—like also most relationships—appear to be sparsely researched.

Nonetheless, the network can help progress theoretical research already in its current state. We provide an initial attempt in the following section. More obviously, the gaps we can point out in the network open questions and challenges for future research. In Section 8 we discuss research opportunities.

7 PROGRESSING THEORETICAL RESEARCH

Given a large collection of arguments, many positive and some negative, one cannot help wonder i) if these arguments can be summed up by one or a few more fundamental arguments, and ii) if the conflicts and contradictions can also be explained by these fundamental arguments. In the visualization literature, there are several candidates for such a theoretic investigation, including the data-ink ratio [53], the three algebraic principles [38], and the cost-benefit metric [17]. None of these theoretic propositions has been proved using mathematical means. Like many theoretical propositions in history, the most straightforward way to test them is trying to falsify them by finding counter-examples.

The arguments about "why visualize?" and "why not" represent some wisdoms extracted from positive and negative experience in practice. We can thus assume that each argument is correct at least in some specific conditions. If a theoretic proposition can provide the argument with a fundamental rationale, this is a small contribution to showing its qualities. If the theoretic proposition fails to do so, the argument can serve as a significant counter-example, suggesting that the proposition is likely to be incomplete or subject to unspecified conditions. The large collection of arguments discussed in the previous sections thus offers a valuable opportunity to test theoretic propositions in
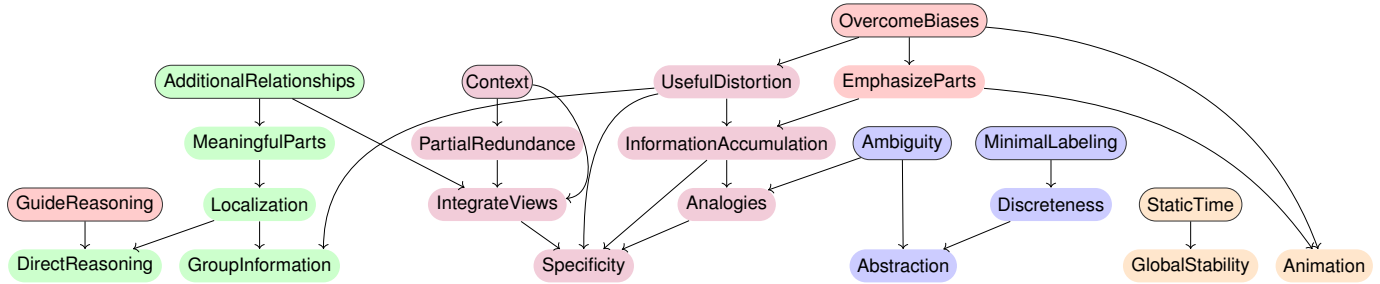
Fig. 7: Induced subgraph of the *depends on* relation for all arguments from the Design standpoint. Arguments are grouped by the broader theme they relate to. Green arguments are about the composition, purple ones about matching visualizations to scenario demands. Red arguments deal with anticipating viewers' thinking, blue arguments are based on abstraction, and orange ones are on stability and animation. Arguments on which no other arguments from this standpoint build upon are shown with solid borderlines.
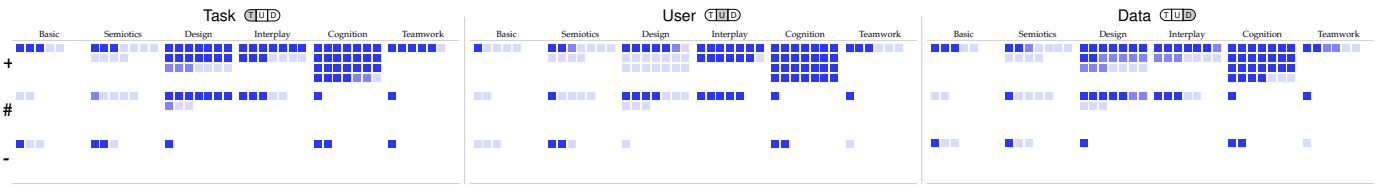

Fig. 8: Task-, User-, and Data-dependencies of arguments aggregated by affirmation (rows) and standpoint (columns). ■ denotes likely dependency; ■ maybe; ■ unlikely. For example, all arguments from the Cognition standpoint are *likely* to be user-dependent across all affirmations.

visualization through falsification. In this section, we first briefly discuss the data-ink ratio [53] and the three algebraic principles [38]. We then focus on the cost-benefit metric [17].

Tufte proposed that visual designs can be optimized by improving the ratio between the amount of data being depicted and the amount of ink used [53]. This theoretic proposition can explain a number of arguments such as MultipleMeanings, MinimalLabeling, Compactness, and SpatialConstraints, but has difficulties in supporting many arguments that are User-dependent or Task-dependent and focus on pathways ③ or ④ in Figure 4. If this theoretic proposition was complete, it would suggest that the better data-ink ratio the better a visualization enables viewers' cognitive capabilities. One can find counter-examples where redundant ink reduces cognitive load.

For example, PartialRedundancy includes both positive and negative views on redundancy. Rheingans and Landreth first detected the benefits of redundancy in visualization using an empirical study [81]. Chen and Jänicke later explain such benefits using the information-theoretic concept of error detection and correction [16]. Arguments VisualPatterns, EnableComplexStructures, and Rerepresentation center around the merit that visualizations reveal structures in data. Many structures are depicted explicitly, such as connections and spatial partitions. In practice visual representations are often not economic in using inks. For example, typically a connection line in graph or tree visualizations uses a fair amount of ink for 1 bit of information. This leads to the mixed argument of LimitedAbstraction.

Kindlmann and Scheidegger proposed three algebraic principles to ensure good visualization designs, namely *representation invariance*, *unambiguous data depiction*, and *visual-*

data correspondence [38]. The theoretic proposition suggests that a visualization process is useful if it follows the three principles, and it is corroborated by several arguments such as StorageInvariance, GlobalStability, and PrescribeDiscourse. However, this proposition inadvertently supports negative arguments such as InformationLoss since many visual representations used in practice would feature down-sampling from higher-resolution data (e.g., overview visualization) and projection from higher-dimensional space (e.g., volume visualization) making ambiguity inevitable.

Meanwhile, with Misleading many examples of deceptive visualizations [82] do follow the three principles. Hence, some arguments for or against visualization cannot be explained by whether or not the visual designs follow the three principles. In general, data-ink ratio [53] and the three algebraic principles [38] are meaningful abstractions, especially of arguments less sensitive to variations of Users and Tasks. However, both are incomplete as counter-examples can be identified by following some arguments that are Task-dependent (e.g., InformationLoss) or User-dependent (e.g., Misleading).

Chen and Golan proposed an information-theoretic metric for analyzing the cost-benefit ratio of visualization processes [17]. It considers a visualization process as a sequence of transformations, $P_1, P_2, \ldots, P_i, \ldots, P_n$, where P_i may be running an algorithm for visual-mapping, viewing a visualization image, performing an interaction, and so on. Iterative processes are accommodated through sequentialization. The metric consists of three measures:

Alphabet Compression measures the entropy reduction ($\in \mathbb{R}$, in *bits*) of a transformation by comparing the input alphabet with the output alphabet. Most transformations in
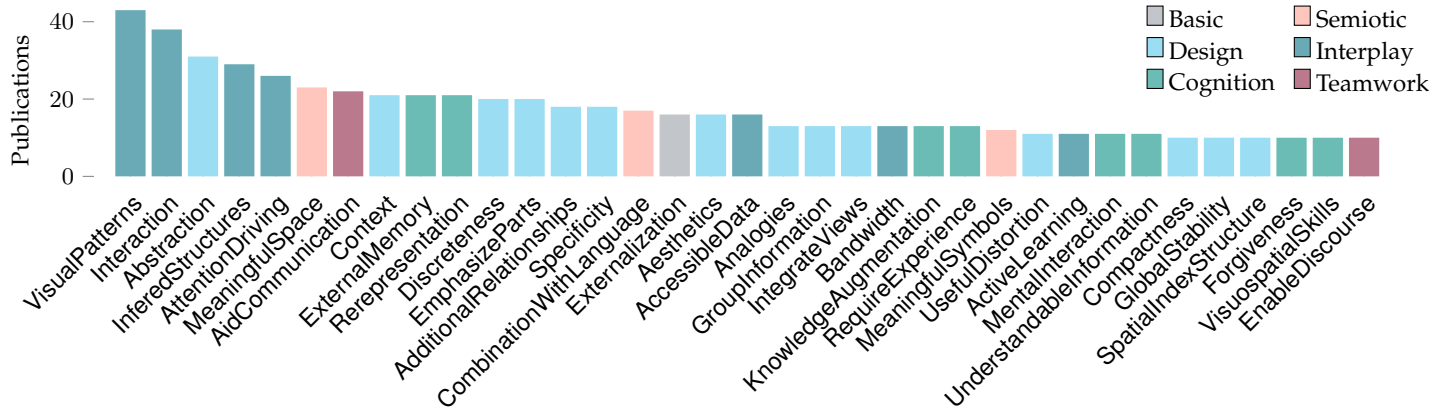
Fig. 9: Arguments mentioned in at least ten publications, colored by standpoint (see Section 5.1). While we focused the sample of publications on the Design and Cognition standpoints (as described in Section 5), arguments like VisualPatterns, Interaction, MeaningfulSpace, or AidCommunication are still prominent. They may be central arguments for visualization. For instance, in our network we can see the central position of VisualPatterns within the Interplay standpoint (see Figure 5).

a visualization process, such as filtering, visual mapping, projecting, selecting, navigating, decision-making, etc., result in positive measurements of Alphabet Compression.

Potential Distortion measures the potential distortion ($\in \mathbb{R}_{\geq 0}$, in *bits*) if one attempts to reconstruct the input alphabet based on the output alphabet. A viewer's soft knowledge is implicitly featured here as the viewer can exploit extra knowledge to infer the input alphabet. Potential distortion is mostly positive.

Cost measures the cost ($\in \mathbb{R}_{>0}$) for the transformation and reconstruction. Ideally, it is measured in an energy unit, but may be approximated by time or monetary cost.

The trade-off among the three measures is expressed as:

$$\frac{Benefit}{Cost} = \frac{Alphabet\ Compression - Potential\ Distortion}{Cost} \quad (1)$$

For its detailed mathematical definition, we refer you to Chen and Golan [17]. The metric can be applied to individual transformations, a series of transformations, and the whole process. It suggests that visualization is useful whenever it offers more cost-benefit than viewing data directly or applying statistical algorithmic analysis, and that a visual design is better than another if the former offers a better cost-benefit ratio than the latter. The metric gives rise to several assertions concerning visualization processes:

- Losing information (i.e., alphabet compression) is a ubiquitous phenomenon in data intelligence processes, and has a positive impact on *Benefit*. Hence, losing information (e.g., InformationLoss) should not lead to a conclusion "not to visualize".
- A viewer's soft knowledge can be used to reduce the potential distortion. For example, when a time series is displayed in conjunction with a y-axis shifted away from the zero base, it may cause less potential distortion to an expert than a novice. So argument Misleading is User-dependent and Task-dependent. For some viewers and with some tasks, a so-called "vislie" may not lead to distortion.

- Visualization tasks are implicitly encoded in the output alphabets of some transformations, usually towards the end of the sequence $P_1, P_2, \ldots, P_{n-1}, P_n$. For example, if P_n is a transformation from observation to decision, the output alphabet consists of a set of possible choices. The task of P_n is typically to choose from a decision alphabet, i.e., making a decision.
- Modifying a transformation P_i may change its alphabet compression, potential distortion and cost, and may also change the three measures in subsequent transformations $P_{i+1} \ldots, P_n$. Hence, optimizing a visualization process must be done holistically.

In general, the metric in Equation 1 exhibits a trade-off among three measures. This suggests that ambivalent arguments are likely derived from cost-beneficial trade-offs and negative arguments from less successful trade-offs. The need to consider the three measures holistically in optimizing a visualization aligns with the complexity in designing visual representations and visualization systems.

Broadly, the positive impact of *alphabet compression* can explain many positive arguments related to abstraction, highlighting and interaction, and can counter the negative arguments related to information loss. The role of knowledge in reducing *potential distortion* can explain the conditional dependence on Users or Tasks of many arguments since different users may have, and different tasks may require, different levels of knowledge. It also supports many positive arguments in comparison with analytical statistics and algorithms, since these usually offer more alphabet compression but much less help to the reconstruction of input alphabets. The need to consider *cost* as part of the metric supports many positive arguments related to time saving, external memorization, and the utilization of cognitive capabilities such as visual search, pattern recognition, and knowledge acquisition and deployment. In the following subsections, we discuss several arguments in detail.

7.1 Arguments about Representation and Meanings

Arguments Homogeniety and StorageInvariance state that the same visual representation can be used for different data objects or the same data in different formats, while FlexibleNotation, DifferentMappings, and the failure of ConsistentNotation point out that the same data object can be depicted using different visualizations. In addition, when considering the meanings conveyed by visualizations, MultipleMeanings, IdeasByAmbiguity, and Ambiguity assert that visual representations can carry multiple meanings, in some cases independently and in other cases ambiguously.

Information-theoretically, we can consider visual mappings from data to visualizations as *transformations* from data alphabets to visualization alphabets, and the semantic mapping from visualization objects to meaning as transformations from visualization alphabets to alphabets of semantic meanings. Let us denote the two types of transformations as P_{D2V} and P_{V2M} respectively. Transformation from different data objects to the same visual representation suggests the presence of many-to-one mapping in P_{D2V}, which is a form of alphabet compression in Equation 1. The alphabet compression in P_{D2V} is one of the causes of the ambiguity in P_{V2M}, i.e., the potential distortion.

Some may suggest that mapping different data objects to the same visual representation is an improper way to conduct visualization. In fact, this is ubiquitous in visualization. For example, given data objects, 1.99, 2.00, and 2.01, the chance of them being mapped to the same visual representation is very high. When a visualization process works, alphabet compression in P_{D2V} results in more beneficial consequences, as for example suggested by MultipleMeanings, than causing adverse problems such as ambiguous meaning (Ambiguity).

We can juxtapose P_{D2V} and P_{V2M} with the transformation from objects and events to some entities of languages (e.g., words, phrases, and sentences), and from languages to meanings. In the context of languages, the role of human knowledge is absolutely critical in delivering the many-to-one mappings in the first transformation and untangle the one-to-many mappings in the second transformation. It is not difficult to extrapolate from languages back to visualization to realize the critical role of viewers in visualization. With some *a priori* knowledge about the data and the knowledge of visual representations, a better trade-off between alphabet compression and potential distortion can be achieved than in situations without such knowledge.

Some arguments agree with the major role of human knowledge in visualization. MeaningfulSpace suggests that visualizations convey information with empty space. Similarly, speech can convey information via intonation, and even silence. Without context and spoken parts of the conversation, silence would hardly be understandable. NaturalMeaning expostulates that visual representations convey meaning "naturally" using metaphors of real world objects. In comparison to languages, which are costly to learn, this naturalness makes knowledge of visualization easier to acquire and more readily available. Without such knowledge, the trade-off between alphabet compression and potential distortion in Equation 1 will unlikely favor visualization.

7.2 Arguments about Information Loss and Distortion

Undeniably, visualizations—intentionally or unintentionally—may cause information loss and distortion as suggested by negative arguments, such as InformationLoss and Misleading. Even many positive arguments imply that visualizations may not work, if some information is lost or distorted, e.g., losing topological information in TopologyPreservation, failing to follow mathematical laws in LogicMathematic, using inconsistent notations in ConsistentNotation, and failing to maintain global stability in GlobalStability. The central assumption of these arguments is that information is useful. So, it should not be lost or distorted.

There are numerous definitions of "information" (e.g., [83]). Although "useful information" is a commonly-used phrase, the definition of usefulness depends on viewers, tasks, and contexts. In information theory definitions of information [84] mostly are semantically neutral. Some have a slightly negative implication. For example, one well-known definition of information is Shannon entropy [85]. It measures the average uncertainty of an alphabet, hinting less may be better in many situations. Another definition of information is Kullback-Leibler Divergence [86], which measures the difference between two alphabets, hinting that the difference may be bad when measuring distortion and may be good when measuring learning.

Equation 1 uses these information-theoretic measures to define two different types of information. It defines alphabet compression as the reduction of Shannon entropy after a transformation, hinting losing information in terms of uncertainty may be beneficial. It defines potential distortion as the difference between the original input alphabet and the reconstructed alphabet, hinting the original information in the input alphabet is beneficial. Thus, benefit is expressed by subtracting potential distortion from alphabet compression, which is the trade-off between reducing one measure of information and maintaining another.

Recall our discussion about languages in Section 7.1. In comparison to the alphabet of all possible real world objects and events, the alphabet of language entities is much smaller. So, the transformation from objects and events to language entities exhibits a huge amount of alphabet compression. Potential distortion is thus inevitable from time to time depending on the complexity of objects and events, the purposes of describing them, and the people who speak (write) and listen (read). It would be insane to suggest that spoken or written words must capture every piece of information about an object or an event. Although it might not be helpful, losing some topological information about an object or an event, failing to follow mathematical laws, and having a bit of inconsistency would usually be tolerated by the listeners and readers. Even distortion is allowed in languages. For example, irony and sarcasm are forms of distortion; successful deployment depends on listeners.

This trade-off between alphabet compression and potential distortion can explain why it is not difficult to find visualizations that break the preconditions of TopologyPreservation, LogicMathematic, ConsistentNotation, and GlobalStability but still are very useful. Often visualizations regarded as "vislies" by Misleading are widely used in practice. For example, most finance sites (e.g., Yahoo Finance) display

time series of stocks and exchange rates without following the guideline of showing the zero baseline [53], [82]. Likely, most of their users have adequate financial knowledge to alleviate the potential distortion due to the misleading visualizations. Most metro maps distort the geographical locations of stations (i.e., geometry) in order to enable better presentation of topological connections. Most city maps show the accurate location of train stations without their topological connections in order to reduce cluttering.

Nevertheless, arguments TopologyPreservation, Logic-Mathematic, ConsistentNotation, and GlobalStability express desirable properties of visualizations in many applications, where omitting such preconditions may lead to distortions with grave consequences. Here the third element of Equation 1 comes into play. Visualizations need to balance the benefits and the costs of displaying different aspects of data, and the costs of not displaying them.

7.3 Arguments about Patterns and Structures

Many arguments articulate that a major reason for using visualizations is that they reveal patterns and structures. Some attribute this capability to visual representations (e.g., Rerepresentation, Abstraction, VisualPatterns, EnableComplexStructures), while others attribute it to human cognitive capabilities (e.g., PerceptualInference, EnhancedStructures, ContextPerception). The two groups of arguments focus on different transformations, P_{D2V} and P_{V2M}, in Section 7.1.

In fact, these arguments are mostly made against scenarios of viewing data without visualization. For the first group of arguments, there are obvious merits over viewing data when P_{D2V} can correctly extract patterns and structures and explicitly depict (e.g., connection lines between dots) or highlight them (e.g., showing clusters of dots in different colors). The second group of arguments focuses on phenomena where patterns and structures are unknown to the first transformation P_{D2V} and it is the viewers' cognitive capabilities that enable the identification of patterns and structures in the second transformation P_{V2M}.

Although we are yet to have a full understanding about such cognitive capabilities, it is generally believed that this is due to a combination of human capabilities such as visual search, selective attention, Gestalt-grouping, heuristics, memory, knowledge, reasoning, hypothesizing, and so on [87]. In the literature, there also is evidence suggesting shortcomings of these capabilities, such as inattentional blindness, illusion, biases, forgetting, and so on [88], [89].

This leads to the question, what if visualization is compared against automated techniques for spotting patterns and structures using techniques such as classification, clustering, association analysis, principle component analysis, dimensionality reduction, neural networks, decision trees, etc. Would such comparison lead to arguments against using visualizations, or for automated techniques to tell us what to see? Hence, with the shortcomings of human cognitive capabilities and the improvement of machine intelligence, the arguments at the beginning of this subsection could potentially be used as counter-arguments.

The cost-benefit metric [17] can be applied to machine-centric as well as human-centric transformations in any data intelligence workflow. One major shortcoming of a machine-centric process is the potential distortion during the reverse mapping from an output alphabet to an input alphabet. As discussed earlier, humans' knowledge can alleviate the potential distortion by taking into account some additional variables that are not in the data and using additional heuristics to reason about the data. Recently Tam et al. [90] reported two observational studies, each of which compared a fully-automated workflow with a visual analytics workflow. In both cases they found that the visual analytics processes achieved better results. They quantitatively estimated the amount of knowledge of human analysts that can be useful to the visual analytics workflow, and found that there was more information in the knowledge than in the data used in these two case studies. Both, the data space and the knowledge space are measured using Shannon entropy. The use of human knowledge in the reverse mapping from visualization to data was confirmed in a recent empirical study by Kijmongkolchai, Abdul-Rahman and Chen [91].

While automated processes can achieve higher rates of alphabet compression and lower running cost, they will certainly be less competitive in alleviating potential distortion. When human analysts are in a workflow, $P_1, P_2, \ldots, P_n$, using visualizations to observe different processes, visualizations facilitate external memorization and provide the workflow as well as the human minds with provenance of transformations within the workflow, as suggested by ExternalMemory and Provenance. The externalized knowledge further helps in reducing potential distortion.

Stasko [92] made a related argument that visualization helps the confidence of analysts (i.e., Confidence). Imagine a machine-centric workflow $P_1, P_2, \ldots, P_n$. Further, imagine a human-centric secondary workflow $C_1, C_2, \ldots, C_n$ running parallel to $P_1, P_2, \ldots, P_n$. Let us define confidence as an alphabet with five levels $[1, 2, 3, 4, 5]$ from very doubtful to very confident. Assume that the analyst initially does not know which value of the 5-scale should be, so the confidence alphabet has the maximum entropy (about 2.3 bits) for each P_i. At every stage C_i, the analyst observes some visualizations related to P_i, pondering confidence. If the analyst is certain about a particular confidence value (e.g., 2) due to some errors of P_i, the entropy of the confidence alphabet becomes 0 bits. At the end of the process, this secondary workflow can achieve alphabet compression of up to $2.3n$ bits. Hence, using visualization to determine if one is confident or not about a machine-centric workflow brings about benefit.

Summary: Due to space limitations, we can only discuss a small number of arguments here. We considered all arguments in our collection. For most arguments, we find that individually each typically focuses on one of the three information-theoretic measures. When we consider several related or conflicting arguments together, the notion of some trade-off reveals itself. In general, the cost-benefit metric presents an abstract argument about the trade-off among three information-theoretic measures. In the literature, the optimization of this trade-off is often phrased as an argument that visualization facilitates effective and efficient cognitive processes in performing tasks involving data [13], [55] (i.e., TaskEfficiency). Here, the word "effective" captures the sense of high alphabet compression and low potential distortion, "efficient" captures the sense of high alphabet compression (for subsequent processes) and low cost.

8 RESEARCH OPPORTUNITIES AND APPLICATIONS

Having presented an example of how to use the network in visualization theory, in this section we present more potential applications of our network approach to empiric visualization research (Section 8.1) as well as the practice of qualitative and quantitative evaluation (8.2), and the design of visualizations (8.3).

8.1 Implications for Empirical Research

The theoretical discourse in the previous section uses the network for putting an existing theory to the test. In empirical research the network may reveal potential hypotheses about aspects that have not been investigated empirically in a rigorous manner, thus pointing out new research opportunities. For example, empirical research on ActiveLearning and FlexibleUsage of visualizations is limited. Especially when it comes to HeuristicInteraction there remain plenty research opportunities. For example, what heuristics do visualization experts, domain experts, and novices apply? How effective are these? And, why do they differ in efficiency? Given the network we further expect that some visualizations are usable without knowledge about their rules of construction (IrrespectiveOfConstruction). In which cases this theoretical assertion holds needs to be investigated.

Secondly, the network aids in formulating hypotheses as it untangles arguments, and hence properties of visualizations. As an example, please reconsider the VisualPatterns argument in Figure 5. The subgraph clearly depicts factors argued to be the foundation of visual patterns, namely AttentionDriving, PerceptualInference, Precision, and ImplicitLearning. While aspects of visual patterns, in general, have been investigated extensively, future research may look more closely at (implicit) learning.

Furthermore, the current state of the network is far from being complete. Likely there are missing connections, and relationships lacking empirical support. Relationships including at least one argument from the Cognition standpoint are good starting points for potentially interesting research questions. For instance, how ambiguity in a representation or different representations of the same problem can help in solving problems (i.e., IdeasByAmbiguity, and Rerepresentation) is only partially understood, and not in the focus of the InfoVis community today. Clarifying the relation between not only these arguments, but also their prerequisites is a next step towards better understanding. Establishing the network structure in this area will connect arguments expressed as early as by Feynman [93], and Larkin and Simon [35], with those by Krish [47], and Tversky [51]. Additionally, new arguments focused on these particular aspects may be added.

Another way to use the network is to research disagreements. These can be conflicts between arguments encoded in the network, or different affirmations of single arguments and relationships. For example, Animation is discussed broadly with diverging affirmations. Categorizing empirical research on the topic by connecting the findings of experiments with related arguments in the network can explicate differences between experiments and help integrate findings. Subsequently, the provided structure facilitates retracing lines of thought, and uncovering needs for targeted replication. Last but not least, the references embodied in the network provide a source of relevant publications.

Taken together, the network can provide empiric researchers with a model at a useful level of abstraction. It offers enough detail to relate individual effects, and sufficient abstraction to embed them in a bigger picture. Resulting research designs allow for generalizable empiric research that may lead to widely applicable insights.

8.2 Qualitative and Quantitative Evaluation

Next to inspiring novel empiric investigations, the network may also be used in the critique of existing visualizations. As a practical example we consider the network graph on the whiteboard used during the construction of our network (see Figure 1 (B)) and the interactive adjacency matrix view in the online tool we developed (Figure 1 (C)) used for later refinement. The network helps us to define demands as a set of arguments. For the case of constructing an argument network such set could consist of the need to repeatedly interrupt and resume working on the network (AidResuming), and setting the expectation that the current network is only an intermediate result (Forgiveness).

In order to engage in a solid critique, we must ensure that we do not miss related or otherwise relevant aspects. For example, in the network we find that AidResuming *builds on* GlobalStability. The whiteboard diagram is an example of a visualization that is very stable. Any change is a result of a viewer's active engagement as long as some very basic preconditions are met. For instance, there is no wind blowing away sticky notes. The matrix view, on the other hand, is not necessarily stable from a global perspective when adding a new relationship. As arguments are sorted such that all *builds on* relationships lay in the lower triangle the sorting is automatically updated once a new relationship is added. Arguments with sort order undetermined by the *builds on* relation may flip when a new order is calculated. Such flips undermine the demand for GlobalStability.

One has to note though that our network is not intended to entail all relevant aspects to be considered when deciding which of the two to use. Being theory-driven and focused on the properties of visualization objects, for example, it does not include the effort needed to implement visualizations. In our case, you can clean a whiteboard, and pick some sticky notes and pens in a minute, but you may need days or weeks to implement an interactive visualization. Applying the network to critique visualizations should be done mindfully and non-mechanically. Knowing demands and considering criteria external to the argument network is key to making high-quality choices. Nonetheless, the network can add structure to a large part of the evaluation process.

Another research opportunity is the collection and development of quantitative evaluation techniques / quality measures. There are several measures available, many more than we can feature here [94]. To give only one prominent example, Tufte's data-ink ratio [53] can be used to measure Discreteness. Nonetheless, there are arguments for which to the best of our knowledge no measures exist. One such example is measuring how well two or more views are integrated with each other (IntegrationOfViews).

The network may be used to bridge the gap between visualization theory, quantitative measures and to date often

implicit or vague demands. At the same time the network offers lines of thought why a single quality measure, such as Tufte's data-ink ratio, often is not enough to capture visualization quality. In practice, providing automated assistance by proposing suitable visualizations is limited by the quality of this measure-to-demand link.

8.3 Design of Visualizations

Finally, the network may also be applied in the design of novel visualizations. As with traditional visualization design, the first step is to specify initial demands. To date designers continue with finding a design that meets these initial demands. Using the network introduces an intermediate step. Before entering the design phase, further implicit demands can be inferred from the network. Picking up the VisualPatterns example again (see Figure 5), one can infer that VisualPatterns *builds on* driving viewer's attention to relevant parts (AttentionDriving), but also enabling the viewer to learn how to interact with a visualization effectively (ImplicitLearning). Quite obviously, learning can only happen efficiently if the attention driving effect is not too strong. Otherwise, viewers need to put constant effort in overriding exaggerated distracting forces. Driving attention to visual patterns is a good start for visual analysis of data. However, once analysts proceed beyond spotting a pattern to investigating its origin, reliability and context, they must be able to free themselves from being permanently attracted to the pattern. Applying the network in the design of novel visualizations leads to a richer picture of requirements, even before the first prototype has been sketched.

Besides being directly applied in the design of visualizations, the network can be used to underpin design guidelines with theory. In line with its application in empiric research (Section 8.1) the network provides a theoretic structure, which matches the level of abstraction of many guidelines. One common example is to avoid visual decorations unrelated to the data / chartjunk [53]. Why to avoid decorations might be founded on several theoretic arguments. First, decorations change the visual context and hence perception (ContextPerception). Secondly, decorations likely drive attention (AttentionDriving) away from the relevant visual patterns. Thirdly, decorations may add Ambiguity in cases calling for clarity. A detailed investigation could continue from here, but exceeds the scope of this paper.

9 CONCLUSION

In this work, we curate a network of arguments based on a sizeable collection of mostly positive and some negative arguments as to "why visualization works". With the network we offer a first roadmap of theoretical arguments on the aspects underlying benefits provided by visualization. We categorize these arguments based on several schemes, allowing the examination of arguments in groups. The categorizations aid in navigating within the network and offer connections to coarser and more detailed levels of abstraction. Perhaps most interestingly, we observe a number of needs for trade-off among arguments, and present numerous opportunities for future research and practical application. Finally, we make the network and the collection

of arguments available online as an open dataset about visualization research, resulting from manual coding of 108 publications.

Our work suggests a new scope in developing a theoretical foundation of visualization [6]. Arguments about "why visualization works" are usually related to "how does visualization work", and "how to make visualizations work", especially at a more detailed level regarding individual visual representations and visualization systems. It will require a huge effort to examine conflicts and needs for trade-offs as well as to bridge between the abstract concepts discussed in this work and concrete mechanisms. Without doubt, these efforts will lead to profound impact on practice.

REFERENCES

[1] M. Scaife and Y. Rogers, "External cognition: How do graphical representations work?" *Int. J. Human-Computer Studies*, vol. 45, no. 2, pp. 185–213, 1996.

[2] H. C. Purchase, N. V. Andrienko, T. J. J. Kelly, and M. O. Ward, "Theoretical foundations of information visualization," in *Information Visualization - Human-Centered Issues and Perspectives*. Springer, 2008.

[3] T. J. Jankun-Kelly, R. Kosara, G. Kindlmann, C. North, C. Ware, and E. W. Bethel, "Is there science in visualization?" in *IEEE Visualization Panel*, Baltimore, MA, 2006.

[4] C. Ziemkiewicz, P. Kinnaird, R. Kosara, J. Mackinlay, B. Rogowitz, and J. S. Yi, "Visualization theory: Putting the pieces together," in *IEEE VisWeek Panel*, Salt Lake City, UT, 2010.

[5] Ç. Demiralp, D. Laidlaw, J. V. Wijk, and C. Ware, "Theories of visualization - are there any?" in *IEEE VisWeek Panel*, Providence, RI, 2011.

[6] M. Chen, G. Grinstein, C. R. Johnson, J. Kennedy, and M. Tory, "Pathways for theoretical advances in visualization," *IEEE Computer Graphics and Applications*, vol. 37, no. 4, pp. 103–112, 2017.

[7] J. J. Thomas, *Illuminating the Path: The Research and Development Agenda for Visual Analytics*, J. J. Thomas and K. A. Cook, Eds. IEEE Computer Society, 2005.

[8] S. K. Card, J. D. Mackinlay, and B. Shneiderman, *Readings in information visualization - using vision to think*. Academic Press, 1999.

[9] N. J. Thomas, "Mental imagery," in *The Stanford Encyclopedia of Philosophy*, spring 2018 ed., E. N. Zalta, Ed. Metaphysics Research Lab, Stanford University, 2018.

[10] F. J. Anscombe, "Graphs in statistical analysis," *The American Statistician*, vol. 27, no. 1, pp. 17–21, 1973.

[11] S. M. Kosslyn, *Elements of graph design*. W. H. Freeman and Company, 1994.

[12] H. Wickham, D. Cook, H. Hofmann, and A. Buja, "Graphical inference for infovis," *IEEE Trans. Visualization and Computer Graphics*, vol. 16, no. 6, pp. 973–979, 2010.

[13] M. Chen, L. Floridi, and R. Borgo, "What is visualization really for?" in *The Philosophy of Information Quality*. Springer, 2014.

[14] J. J. van Wijk, "The value of visualization," in *Proc. IEEE Conf. Visualization*, 2005.

[15] P. Vickers, J. Faith, and N. Rossiter, "Understanding visualization: A formal approach using category theory and semiotics," *IEEE Trans. Visualization and Computer Graphics*, vol. 19, no. 6, pp. 1048–1061, 2013.

[16] M. Chen and H. Jaenicke, "An information-theoretic framework for visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 16, no. 6, pp. 1206–1215, 2010.

[17] M. Chen and A. Golan, "What may visualization processes optimize?" *IEEE Trans. Visualization and Computer Graphics*, vol. 22, no. 12, pp. 2619–2632, 2016.

[18] D. Sacha, A. Stoffel, F. Stoffel, B. C. Kwon, G. Ellis, and D. A. Keim, "Knowledge generation model for visual analytics," *IEEE Trans. Visualization and Computer Graphics*, vol. 20, no. 12, pp. 1604–1613, 2014.

[19] S. K. Shevell, Ed., *The science of color*, 2nd ed. Elsevier Science, 2003.

[20] P. Kay, B. Berlin, L. Maffi, W. A. Merrifield, and C. Richard, *The world color survey*. CSLI Publications, 2010.

[21] B. Rogowitz, L. A. Treinish, and S. Bryson, "How not to lie with visualization," *Computers in Physics*, vol. 10, no. 3, pp. 268–273, 1996.

[22] S. Mittelstäädt and D. A. Keim, "Efficient contrast effect compensation with personalized perception models," *Computer Graphics Forum*, vol. 34, no. 3, pp. 211–220, 2015.

[23] M. Steiger, J. Bernard, S. Mittelstädt, S. Thum, M. Hutter, D. A. Keim, and J. Kohlhammer, "Explorative analysis of 2d color maps," in *Int. Conf. in Central Europe on Computer Graphics, Visualization and Computer Vision*, 2015, pp. 151–160.

[24] W. Gaissmaier, O. Wegwarth, D. Skopec, A.-S. Müller, S. Broschinski, and M. C. Politi, "Numbers can be worth a thousand pictures: Individual differences in understanding graphical and numerical representations of health-related information," *Health Psychology*, vol. 31, no. 3, pp. 286–296, 2012.

[25] M. Sirota, L. Kostovičová, and M. Juanchich, "The effect of iconicity of visual displays on statistical reasoning: evidence in favor of the null hypothesis," *Psychonomic Bulletin & Review*, vol. 21, no. 4, pp. 961–968, 2014.

[26] J. Boy, A. V. Pandey, J. Emerson, M. Satterthwaite, O. Nov, and E. Bertini, "Showing people behind data: Does anthropomorphizing visualizations elicit more empathy for human rights data?" in *Proc. Conf. Human Factors in Computing Systems*, 2017.

[27] L. Micallef, P. Dragicevic, and J.-D. Fekete, "Assessing the effect of visualizations on Bayesian reasoning through crowdsourcing," *IEEE Trans. Visualization and Computer Graphics*, vol. 18, no. 12, pp. 2536–2545, 2012.

[28] A. Khan, S. Breslav, M. Glueck, and K. Hornbæk, "Benefits of visualization in the mammography problem," *Int. J. Human-Computer Studies*, vol. 83, pp. 94–113, 2015.

[29] K. Yamagishi, "Facilitating normative judgments of conditional probability: Frequency or nested sets?" *Experimental Psychology*, vol. 50, pp. 97–106, 2003.

[30] G. L. Brase, "Pictorial representations in statistical reasoning," *Applied Cognitive Psychology*, vol. 23, no. 3, pp. 369–381, 2008.

[31] ——, "The power of representation and interpretation: Doubling statistical reasoning performance with icons and frequentist interpretations of ambiguous numbers," *J. Cognitive Psychology*, vol. 26, no. 1, pp. 81–97, 2014.

[32] L. Wilkinson, *The grammar of graphics*. Springer, 1999.

[33] H. Wickham, "A layered grammar of graphics," *J. Computational and Graphical Statistics*, vol. 19, no. 1, pp. 3–28, 2010.

[34] Ç. Demiralp, C. E. Scheidegger, G. L. Kindlmann, D. H. Laidlaw, and J. Heer, "Visual embedding: A model for visualization," *IEEE Computer Graphics and Applications*, vol. 34, no. 1, pp. 10–15, 2014.

[35] J. H. Larkin and H. A. Simon, "Why a diagram is (sometimes) worth ten thousand words," *Cognitive Science*, vol. 11, no. 1, pp. 65–100, 1987.

[36] J.-D. Fekete, J. J. van Wijk, J. T. Stasko, and C. North, "The value of information visualization," in *Information Visualization - Human-Centered Issues and Perspectives*. Springer, 2008.

[37] P. Parsons and K. Sedig, "Adjustable properties of visual representations: Improving the quality of human-information interaction," *J. Association for Information Science and Technology*, vol. 65, no. 3, pp. 455–482, 2014.

[38] G. L. Kindlmann and C. E. Scheidegger, "An algebraic process for visualization design," *IEEE Trans. Visualization and Computer Graphics*, vol. 20, no. 12, pp. 2181–2190, 2014.

[39] J. Meyer, J. Thomas, S. Diehl, B. Fisher, D. A. Keim, D. H. Laidlaw, S. Miksch, K. Mueller, W. Ribarsky, B. Preim, and A. Ynnerman, "From visualization to visually enabled reasoning," in *Scientific Visualization: Advanced Concepts*, ser. Dagstuhl Follow-Ups, H. Hagen, Ed. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2010, vol. 1, pp. 227–245.

[40] D. A. Keim, J. Kohlhammer, G. Ellis, and F. Mansmann, *Mastering the Information Age Solving Problems with Visual Analytics*, D. Keim, J. Kohlhammer, G. Ellis, and F. Mansmann, Eds. Eurographics Association, 2010.

[41] B. Tversky, "Visualizing thought," *Topics in Cognitive Science*, vol. 3, no. 3, pp. 499–535, 2011.

[42] M. Hegarty, "The cognitive science of visual-spatial displays: Implications for design," *Topics in Cognitive Science*, vol. 3, no. 3, pp. 446–474, 2011.

[43] A. Gelman and A. Unwin, "Infovis and statistical graphics: Different goals, different looks," *J. Computational and Graphical Statistics*, vol. 22, no. 1, pp. 2–28, 2013.

[44] D. C. Gooding, "Visualizing scientific inference," *Topics in Cognitive Science*, vol. 2, no. 1, pp. 15–35, 2010.

[45] A. Groß, *Die Bildpädagogik Otto Neuraths: Methodische Prinzipien der Darstellung von Wissen*, F. Stadler, Ed. Springer, 2015.

[46] M. T. McCrudden and D. N. Rapp, "How visual displays affect cognitive processing," *Educational Psychology Review*, pp. 1–17, 2015.

[47] D. Kirsh, "Thinking with external representations," *AI & Society*, vol. 25, no. 4, 2010.

[48] R. E. Patterson, L. M. Blaha, G. G. Grinstein, K. K. Liggett, D. E. Kaveney, K. C. Sheldon, P. R. Havig, and J. A. Moore, "A human cognition framework for information visualization," *Computers & Graphics*, vol. 42, pp. 42–58, 2014.

[49] W. Dou, C. Ziemkiewicz, L. Harrison, D. H. Jeong, W. Ribarsky, X. Wang, and R. Chang, "Toward a deeper understanding of the relationship between interaction constraints and visual isomorphs," *Information Visualization*, vol. 11, no. 3, pp. 222–236, 2012.

[50] S. Bertschi and N. Bubenhofer, "Linguistic learning: A new conceptual focus in knowledge visualization," in *Proc. 9th Int. Conf. Information Visualisation*, 2005.

[51] B. Tversky, "On abstraction and ambiguity," in *Studying Visual and Spatial Reasoning for Design Creativity*. Springer, 2015.

[52] H. C. Purchase, "Twelve years of diagrams research," *J. Visual Languages and Computing*, vol. 25, pp. 57–75, 2014.

[53] E. R. Tufte, *The visual display of quantitative information*. Graphics Press, 2001[1983].

[54] S. Krämer, "Zur Grammatik der Diagrammatik," *Zeitschrift für Literaturwissenschaft und Linguistik*, vol. 44, no. 4, pp. 11–30, 2014.

[55] T. Munzner, *Visualization Analysis & Design*. CRC Press, 2014.

[56] P. C.-H. Cheng, "What constitutes an effective representation?" in *Proc. 9th Int. Conf. Diagrammatic Representation and Inference*, 2016.

[57] M. W. Johansen, "What's in a diagram? on the classification of symbols, figures and diagrams," in *Model-Based Reasoning in Science and Technology*. Springer, 2014.

[58] Z. Liu, N. J. Nersessian, and J. T. Stasko, "Distributed cognition as a theoretical framework for information visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 14, no. 6, pp. 1173–1180, 2008.

[59] W. A. Pike, J. T. Stasko, R. Chang, and T. A. O'Connell, "The science of interaction," *Information Visualization*, vol. 8, no. 4, pp. 263–274, 2009.

[60] H. Wainer, "How to display data badly," *The American Statistician*, vol. 38, no. 2, pp. 137–147, 1984.

[61] M. J. Eppler, "What is an effective knowledge visualization? Insights from a review of seminal concepts," in *Knowledge visualization currents: from text to art to culture*. Springer, 2013.

[62] A. Endert, M. S. Hossain, N. Ramakrishnan, C. North, P. Fiaux, and C. Andrews, "The human is the loop: new directions for visual analytics," *J. Intelligent Information Systems*, vol. 43, no. 3, pp. 411–435, 2014.

[63] K. Cook, N. Cramer, D. Israel, M. Wolverton, J. Bruce, R. Burtner, and A. Endert, "Mixed-initiative visual analytics using task-driven recommendations," *IEEE Conf. Visual Analytics Science and Technology*, pp. 9–16, 2015.

[64] W. S. Cleveland and R. McGill, "Graphical perception: Theory, experimentation, and application to the development of graphical methods," *J. American Statistical Association*, vol. 79, no. 387, pp. 531–554, 1984.

[65] E. R. Tufte, *Visual explanations*. Graphics Press, 1997.

[66] R. Goebel, "A sketch of a theory of visualization," in *Proc. Int. Conf. Information Visualization Theory and Applications*, 2014.

[67] L. J. Trevena, B. J. Zikmund-Fisher, A. Edwards, W. Gaissmaier, M. Galesic, P. K. J. Han, J. King, M. L. Lawson, S. K. Linder, I. Lipkus, E. Ozanne, E. Peters, D. R. M. Timmermans, and S. Woloshin, "Presenting quantitative information about decision outcomes: a risk communication primer for patient decision aid developers," *BMC Medical Informatics and Decision Making*, vol. 13, no. S-2, p. S7, 2013.

[68] J. Walny, S. Carpendale, N. H. Riche, G. Venolia, and P. Fawcett, "Visual thinking in action: Visualizations as used on whiteboards," *IEEE Trans. Visualization and Computer Graphics*, vol. 17, pp. 2508–2517, 2011.

[69] S. M. Kosslyn, *Graph design for the eye and mind*. Oxford University Press, 2006.

[70] R. E. Patterson, "Cognitive engineering, cognitive augmentation, and information display," *J. Society for Information Display*, vol. 20, no. 4, pp. 208–213, 2012.

[71] J. Norman, "Differentiating diagrams: A new approach," in *Proc. 1st Int. Conf. Theory and Application of Diagrams*, 2000.

[72] J. S. Yi, Y. a. Kang, J. T. Stasko, and J. A. Jacko, "Toward a deeper understanding of the role of interaction in information visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1224–1231, 2007.

[73] W. Aigner, "Understanding the role and value of interaction: First steps," in *Proc. Int. Workshop on Visual Analytics (EuroVA)*, 2011.

[74] P. Parsons, K. Sedig, A. Didandeh, and A. Khosravi, "Interactivity in visual analytics: Use of conceptual frameworks to support human-centered design of a decision-support tool," in *Proc. 48th Hawaii Int. Conf. System Sciences*, 2015.

[75] D. Spiegelhalter, M. Pearson, and I. Short, "Visualizing uncertainty about the future," *Science*, vol. 333, no. 6048, pp. 1393–1400, 2011.

[76] R. Arias-Hernández, T. M. Green, and B. D. Fisher, "From cognitive amplifiers to cognitive prostheses: Understandings of the material basis of cognition in visual analytics," *Interdisciplinary Science Reviews*, vol. 37, no. 1, pp. 4–18, 2012.

[77] M. Chen and L. Floridi, "An analysis of information visualization," *Synthese*, vol. 190, no. 16, pp. 3421–3438, 2013.

[78] D. A. Keim, G. Andrienko, J.-D. Fekete, C. Görg, J. Kohlhammer, and G. Melançon, *Visual Analytics: Definition, Process, and Challenges*. Berlin, Heidelberg: Springer, 2008, pp. 154–175.

[79] G. Gigerenzer, "The bias bias in behavioral economics," *Review of Behavioral Economics*, vol. 5, no. 3–4, pp. 303–336, 2018.

[80] G. Stapleton, M. Jamnik, and A. Shimojima, "What makes an effective representation of information: A formal account of observational advantages," *J. Logic, Language and Information*, vol. 26, no. 2, pp. 143–177, 2017.

[81] P. Rheingans and C. Landreth, "Perceptual principles for effective visualizations," in *Perceptual Issues in Visualization*, G. Grinstein and H. Levkowitz, Eds. Springer-Verlag, 1995, pp. 59–74.

[82] A. V. Pandey, K. Rall, M. L. Satterthwaite, O. Nov, and E. Bertini, "How deceptive are deceptive visualizations?: An empirical analysis of common distortion techniques," in *Proc. 33rd ACM Conf. Human Factors in Computing Systems*, 2015.

[83] E. H. McKinney Jr. and C. J. Yoos II, "Information about information: A taxonomy of views," *MIS Quarterly*, vol. 32, no. 2, pp. 329–344, 2010.

[84] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. John Wiley & Sons, 2006.

[85] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, 1948.

[86] S. Kullback, *Information Theory and Statistics*. Dover Publications, 1978.

[87] C. Ware, *Information Visualization: Perception for Design*, 2nd ed. Morgan Kaufmann, 2004.

[88] E. E. Smith, S. Nolen-Hoeksema, B. Fredrickson, and G. R. Loftus, *Atkinson & Hilgard's Introduction to Psychology*, 14th ed. Cengage Learning Emea, 2003.

[89] H. Gleitman, A. J. Fridlund, and D. Reisberg, *Psychology*, 6th ed. W. W. Norton & Co., 2003.

[90] G. K. L. Tam, V. Kothari, and M. Chen, "An analysis of machine- and human-analytics in classification," *IEEE Trans. Visualization and Computer Graphics*, vol. 23, no. 1, 2017.

[91] N. Kijmongkolchai, A. Abdul-Rahman, and M. Chen, "Empirically measuring soft knowledge in visualization," *Computer Graphics Forum*, vol. 36, no. 3, pp. 73–85, 2017.

[92] J. T. Stasko, "Value-driven evaluation of visualizations," in *Proc. 5th Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*, 2014.

[93] R. Feynman, *The Character of Physical Law*. Modern Library, 1965.

[94] M. Behrisch, M. Blumenschein, N. W. Kim, L. Shao, M. El-Assady, J. Fuchs, D. Seebacher, A. Diehl, U. Brandes, H. Pfister, T. Schreck, D. Weiskopf, and D. A. Keim, "Quality metrics for information visualization," *Computer Graphics Forum*, 2018.

Dirk Streeb is Ph.D. candidate at the Graduate School of Decision Sciences and member of the Data Analysis and Visualization Research Group at the University of Konstanz, Germany. He received his M.Sc. in Social and Economic Data Analysis from the University of Konstanz in 2016. His research interests include theoretical aspects of information visualization and visual analytics as aids in human decision-making processes.

Mennatallah El-Assady is a research associate in the Data Analysis and Visualization Research Group at the University of Konstanz, Germany, and in the Visualization for Information Analysis lab at the University of Ontario Institute of Technology, Canada. She received her M.Sc. degree in Information Engineering from the University of Konstanz in 2015. Her research interests include visual text analytics, user-steerable topic modeling, and discourse/conversational text analysis.

Daniel A. Keim is professor and head of the Data Analysis and Visualization Research Group in the Computer Science Department at the University of Konstanz, Germany. He has been actively involved in data analysis and information visualization research for more than 25 years and developed novel visual analysis techniques for large data sets. Dr. Keim got his Ph.D. degree in Computer Science from the University of Munich, Germany. Before joining the University of Konstanz, Dr. Keim was associate professor at the University of Halle, Germany, and Technology Consultant at AT&T Shannon Research Labs, NJ, USA.

Min Chen received his Ph.D. degree from University of Wales in 1991. He is currently a professor of scientific visualization at Oxford University and a fellow of Pembroke College. Before joining Oxford, he held research and faculty positions at Swansea University, UK. His research interests include visualization, computer graphics and human-computer interaction. His services to the research community include papers co-chair of IEEE Visualization 2007 and 2008, IEEE VAST 2014 and 2015, and Eurographics 2011; co-chair of Volume Graphics 1999 and 2006, and EuroVis 2014; associate editor-in-chief of IEEE TVCG; and co-director of Wales Research Institute of Visual Computing. He is a fellow of BCS, EG and LSW.