



VIBE: A Design Space for Visual Belief Elicitation in Data Journalism

Shambhavi Mahajan¹, Bonnie Chen² , Alireza Karduni³, Yea-Seul Kim⁴, and Emily Wall⁵ 

¹Independent

²Emory University, Atlanta, GA, USA

³Northwestern University, Evanston, IL, USA

⁴University of Wisconsin, Madison, WI, USA

⁵Emory University, Atlanta, GA, USA

Abstract

The process of forming, expressing, and updating beliefs from data plays a critical role in data-driven decision making. Effectively eliciting those beliefs has potential for high impact across a broad set of applications, including increased engagement with data and visualizations, personalizing visualizations, and understanding users' visual reasoning processes, which can inform improved data analysis and decision making strategies (e.g., via bias mitigation). Recently, belief-driven visualizations have been used to elicit and visualize readers' beliefs in a visualization alongside data in narrative media and data journalism platforms such as the New York Times and FiveThirtyEight. However, there is little research on different aspects that constitute designing an effective belief-driven visualization. In this paper, we synthesize a design space for belief-driven visualizations based on formative and summative interviews with designers and visualization experts. The design space includes 7 main design considerations, beginning with an assumed data set, then structured according to: from who, why, when, what, and how the belief is elicited, and the possible feedback about the belief that may be provided to the visualization viewer. The design space covers considerations such as the type of data parameter with optional uncertainty being elicited, interaction techniques, and visual feedback, among others. Finally, we describe how more than 24 existing belief-driven visualizations from popular news media outlets span the design space and discuss trends and opportunities within this space.

CCS Concepts

• **Human-centered computing** → **Visualization theory, concepts and paradigms**;

1. Introduction

Belief-driven visualization has been employed in the wild to solicit and present viewers' beliefs alongside data. One notable example is "You Draw it" from the New York Times, where the authors prompt viewers to draw a line on a blank chart to indicate their beliefs on the correlation between parents' income and the percentage of children who attended college [ACQ15]. After viewers have drawn their beliefs, the actual data appears alongside the viewer's beliefs (Fig. 3), prompting viewers to reflect on their beliefs and the data. Followed by the New York Times, many news outlets have launched articles with visualizations that elicit viewers' beliefs such as CNN Interactives and Guardian Interactives, possibly to further engage viewers with data. Research demonstrates that prompting viewers' beliefs before showing the data promotes viewers to accurately remember data later [KRH17], and promotes more rational belief updating [KWKH19].

However, while there is clear value and further potential in eliciting beliefs in the journalism context, the methods of designing the belief elicitation have not been investigated. For example, there is

no systematic guidance for visualization designers on what design dimensions they should consider. In particular, there are a number of choices that can triangulate and inform the design of belief elicitation, including the parameter that is being elicited and the particular expertise of the intended audience of the visualization.

In this work, we construct a design space for creating a belief elicitation interface to guide the design process. We focus on "graphical" elicitation that matches the modality to the visualized data (c.f., text elicitation where people enter their beliefs using text) in journalism. We first collect examples in the wild to identify design dimensions involved in elicitation interfaces. To iterate upon the design space, we conducted a user study with five participants from varying backgrounds, including data journalism and visualization design. Based on feedback from participants, we re-designed the design space. Using the design space, we demonstrate how examples in the wild can be characterized using our design space and present a usage scenario of how one might use these guidelines in designing a new belief-driven visualization. We con-

clude by outlining current trends and opportunities in the space of belief-driven visualization.

Our contribution is three-fold:

1. We construct a design space for building belief elicitation interfaces based on (i) real-world examples and (ii) a user study.
2. We demonstrate how existing elicitation examples can be analyzed and characterized using the design space.
3. We showcase how the design space can inform the design and re-design of an elicitation interface via a usage scenario.

2. Background & Related Work

Many scholars have described what beliefs are and characterized their qualities. For example, Jerry Fodor, who is an American philosopher, describes beliefs as “representations of ways that the world could be” [Fod83]. While beliefs can be beyond what we could observe depending on their epistemology and ontological views, the working definition we use is the following: **Mental representations related to a phenomenon that people can express using parameters such as numerical or categorical values.** This definition has been used in many behavioral economic and statistics as it provides a way to mathematically formalize one’s beliefs and take into account them in models and theory. For example, prior work elicits people’s beliefs following this definition to study how beliefs impact their decision making [AFZ19, BDKS19] and how their beliefs can be incorporated into mathematical models [OBD*06].

In this work, we study how one’s beliefs can be elicited to benefit visualization viewers and authors. It is natural to assume that people bring in their prior beliefs when interpreting the new information and update their beliefs based on what is shown. This cycle of belief updating has been formalized using Bayes’ theorem where the updated beliefs are the weighted average of prior beliefs and the newly available information weighted by how uncertain they are. Prior work demonstrates its legitimacy of illustrating updated beliefs based on one’s prior and the given information [KWKH19, KKGMMH20]. The belief updating cycle is as follows:

1. **Forming prior beliefs:** Prior beliefs refer to the beliefs that viewers have before seeing a visualization of a phenomenon. A viewer might form their beliefs through various channels, including prior exposure to data, cultural norms, and expectations, etc.
2. **Expressing prior beliefs:** Through an interface, visualization authors can elicit the viewer’s prior beliefs.
3. **Examining evidence:** The viewer examines the presented visualization.
4. **Updating beliefs:** Based on the presented visualization, the viewer can update their beliefs.

Our work focuses on the second stage of this cycle, where a viewer’s prior beliefs are elicited via an interface. We provide a framework for visualization authors to consider when designing belief-driven visualizations.

2.1. Belief, Knowledge, Preference, and Attitude Elicitation

Beliefs can broadly be described as either categorical or probabilistic. Categorical beliefs are when the belief is all or nothing (e.g., belief if housing is affordable or not), whereas probabilistic beliefs measure the strength of the belief (e.g., what is the likelihood of housing becoming affordable in the next 2 years) [Hun96]. Such probabilistic beliefs are observed and *measured* in controlled experiments by behavioral economists. A belief distribution is an assignment of probability to each scenario in the event [Man04]. The belief elicitation process involves the following stages: set-up, elicitation, fitting a distribution and assessing adequacy [GKO05].

In organizational psychology, there is similarly a body of work on elicitation techniques for organizational knowledge and information management. Researchers have proposed a number of taxonomies for direct and indirect elicitation methods like interviews, questionnaires, observations, storytelling, and brainstorming [Gav93] [MBW99] [Wat85] [HSBK95]. An analyst elicits knowledge that is valuable to the organization from an expert [Wat85]. Some work similarly exists examining attitude elicitation, e.g., to assess its impact on provocative topics [HRR20].

Perhaps the closest body of related work is preference elicitation in recommender systems. Designing interfaces to accurately capture user preferences has become the focus of interactive decision support systems [CP04]. HCI researchers evaluate existing preference elicitation techniques in state of the art recommender systems, and establish a set of usability design guidelines [PBW*12] [PCH12]. Preferences are distinct, however, from beliefs – people construct their preferences when necessary rather than having stable and consistent preferences available [PBS*99, SKH04, WJ06]. Belief elicitation, on the other hand, focuses on eliciting prior (already constructed) beliefs from viewers.

2.2. Graphical Belief Elicitation

Eliciting beliefs using graphical techniques has been used in the visualization field. Six years ago, the New York Times launched the “You draw it” interface where articles elicit viewers’ beliefs on a topic by letting them draw a trend on an empty canvas with the annotated x-axis and y-axis. After the user submits their beliefs, the interface shows the actual data alongside their own beliefs [ACQ15]. This was the first interface that implemented graphical elicitation in the visualization context. Since then, a number of data journalism articles have employed belief elicitation techniques to increase engagement with interactive articles.

Apart from belief elicitation in popular news media, some visualization research has addressed graphical belief elicitation. For instance, Kim et al. [KRH17] investigate the effects of (i) externalizing one’s beliefs and (ii) seeing the data alongside those beliefs on one’s ability to recall the data. The result demonstrates that this technique improves recall, but the effect was not prominent when the elicitation was done in text, compared to being done graphically. Hullman et al. [HKKS18] apply similar techniques in the context of understanding results of scientific experiments and found a similar effect – participants who were prompted to predict the possible outcome of the experiment recalled the data more ac-

curately and were better at predicting other sets of outcomes for new experiments.

More elicitation interfaces have been proposed to record people's probabilistic beliefs regarding proportion parameters [KWKH19], inspired by Bayesian cognition research [VGGT14]. This work demonstrates prompting people to submit a single number multiple times can be beneficial to elicit accurate probabilistic beliefs, compared to directly asking them to construct their subjective probability distribution. A follow-up work devises a new elicitation technique where a participant can set a mode value of their beliefs and define the range where they would not be surprised if the true value falls into the range by using an interactive slider [KKG MH20].

More recently, Karduni et al. [KMWD20] demonstrated a graphical approach could be as effective in eliciting beliefs accurately compared to a more complicated sampling-based procedure of eliciting subjective distributions, such as Markov Chain Monte Carlo with People (MCMC-P) [GKO05]. Nguyen et al. create an authoring tool to support designers to create the interface that elicits beliefs graphically and shows the actual data and other peoples' beliefs consequently [Ngu19]. Most recently, Koonchanok et al. introduced an exploratory visual analytic system where users' expressed their beliefs before seeing new data, which led users to attend to discrepancies between the data and their mental models [KBS*21]. In building on these efforts, our design space complements existing work by providing a systematic framework with which to consider the design of belief-driven visualizations.

3. Initial Design Space & User Evaluation

Based on our literature review, the authors' prior research, and iterative examination of belief-driven articles, we developed an initial version of the design space (Section 3.1). To further refine this initial version of the design space, we conducted a formative user study (Section 3.2).

3.1. Preliminary Design Space

Popular media outlets such as New York Times, Guardian, and FiveThirtyEight regularly publish interactive articles on the web that employ some form of belief elicitation. To formulate our initial design space, we collected and analyzed examples of 14 belief elicitation articles from New York Times, Guardian, CNN, and FiveThirtyEight. Many of the aforementioned media outlets publish summary lists of visual stories and graphics or interactive articles, from which one coauthor manually checked if each article contained a belief elicitation component. Examples were included once they met these criteria: (1) subject's prior belief about the outcome of the assigned task is recorded; (2) the actual data appears to update the subject's prior belief. All of the coauthors iteratively selected two belief elicitation articles at a time and coded them within our design space based on characteristics of the articles that we identified (structured according to the 5Ws + How: *who, what, where, when, why, how* (we ultimately discarded *where* from the design space). The dimensions of the design space were hierarchically defined and varied from iteration to iteration. Dimensions were adjusted when the current iteration revealed characteristics not covered by the design space (e.g., how to characterize the

presence and types of feedback if previous iterations of the articles did not include feedback). Upon each iteration, we applied the current version of the design space to code the subsequent set of two examples for verification, until we reached a saturation where no additional components could be added to the design space, after approximately 7 iterations (included in the supplemental material).

3.2. Formative Study

We next conducted a formative study to incorporate outside feedback into the design space. Additionally, we hoped to understand how experts might use the design space in the creation of a belief-driven visualization. The formative study had several goals:

- assess the utility of the current version of the design space,
- observe and compare the difference when a belief-driven visualization is designed by experts {*with, without*} the use of the design space, and
- refine the space based on user feedback.

Participants. We recruited 5 participants who were familiar with belief elicitation in visual analytics from within our professional networks, including data visualization experts (2), data journalists (2), and UX design practitioners (1). Sessions lasted approximately 60 minutes. We recruited participants who were pursuing or had completed a graduate-level degree in CS, specifically related to HCI or data visualization. There was no compensation for participation.

Tasks. After providing informed consent and going through a briefing on the terminology of belief-driven visualization with the support of a few examples from NYT, participants were first given an unaided belief-driven visualization design scenario (without the use of the design space). Participants were tasked with designing a belief-driven visualization with two problem spaces to choose from to mitigate the possibility that the designer is not comfortable with one of the domains. These design tasks were inspired by existing belief-driven visualizations by one of the venues mentioned earlier (i.e., "Create a system to learn if people know America's geography" based on the NYT example "Where is America's Heartland?" [BQ17] or "Create a system that helps people learn about their food choices impacting climate change" based on the NYT example "Your Questions About Food and Climate Change, Answered" [MPLW19]). During the session, we asked participants to describe their design process by thinking out loud. After they completed the task, we showed them the example belief-driven visualization designed by the publishing news outlet to further inspire them on belief elicitation design.

Next, participants were given a walkthrough of the design space and terminology, followed by a fresh task with a new problem space to avoid carryover from the previous design task. Participants were given a choice between "How do Americans feel about guns?" based on the FiveThirtyEight example "Do You Know Where America Stands On Guns?" [MW18] or "How do socioeconomic factors affect higher education for Americans?" based on the NYT example "You Draw It: How Family Income Predicts Children's College Chances" [ACQ15]. We again chose two different domains in case designers were not comfortable or familiar with

one. In this task, participants were encouraged to reference the design space to think through their design process. After completing both scenarios, we conducted a semi-structured interview with each participant to elicit feedback on the design space.

Qualitative Feedback. Participants expressed a generally positive impression of the design space. It was able to guide the users' design process. For example, P3 said *"I followed the sequence of the design space, and it was similar to my usual design thinking process. It answered my design considerations and helped me organize my thoughts."* While participants suggested organizational changes, they found the content to be comprehensive. P5 said *It is a comprehensive framework. I would have missed out on several design considerations in my usual design process."*

Each participant followed the design space in a systematic way without being suggested what the order should be. P4 said *"The contextual considerations were a good anchor to ask questions about the problem space."* They first identified the contextual considerations (who, why), and then moved on to the interface considerations (when, what, how) once before iterating by going back and forth between the considerations.

However, participants found the number of design considerations to be *visually overwhelming* while simultaneously *lacking sufficient explanation* about each component. Furthermore, participants found it unclear to what extent the problem framing needed to be developed before using the design space. They struggled with scoping a narrower problem statement from the larger problem space. In summary, the formative study led us to identify the primary high-level goal: *the design space, as a resource, needs to be self-contained, including assumptions and guidelines to using it.* This also involved some additional clarifications, described next.

Redesign. Using participant feedback, we made a number of edits to the representation and framing of the design space. First, to resolve the scoping confusion (regarding the level of problem development needed to make use of the design space), we established an assumption that the author starts off with a data set before getting into designing the belief elicitation. We also restructured the context provided in the design space to include (i) belief formation, (ii) belief elicitation, (iii) seeing new data, and (iv) updating belief (rather than the previous context: (i) set up, (ii) elicitation, (iii) fitting a distribution, and (iv) assessing adequacy). Furthermore, we added short descriptors (in gray, e.g., "properties about the user" describing the purpose of the *Who* dimensions, "which stage of belief" describing the purpose of the *When* dimension). We also added annotated usage guidelines (e.g., on the left in pink, to "iterate between steps", and suggestions on "How to elicit truthful beliefs") in the design space.

4. Design Space

This design space is a guide to assist visualization designers, researchers, and data journalists in their design process of belief-driven visualizations. In constructing the design space, we assume that designers have a data set in mind before designing a belief-driven visualization. Our design space provides a framework to design the two aspects of belief elicitation: **contextual consider-**

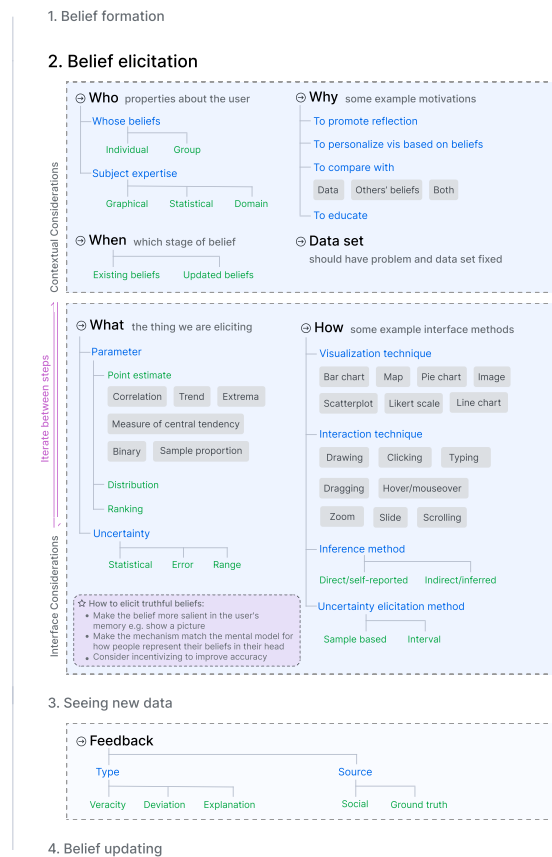


Figure 1: Design space.

ations that guide setting up the context such as target users' characteristics and **interface considerations** that guide setting up the actual elicitation interface.

4.1. Contextual Considerations

Contextual considerations are intended to facilitate narrowing down the problem definition, goals, and intended audience. Triangulating these contextual details will, in turn, inform the specific design choices for the belief-driven visualization.

4.1.1. Who

Whose beliefs are being elicited? The elements of **Who** describe properties about the user that may influence the subsequent design of a belief-driven visualization.

The interface can either elicit a user's *individual* beliefs or a *group's* aggregated beliefs. Eliciting a group's aggregated beliefs introduces additional challenges compared to eliciting an individual's beliefs. Namely, group dynamics affect beliefs, e.g., by tendencies like groupthink [Jan08] or the bandwagon effect [SB15] wherein individuals' behaviors and decisions change as a result of the common opinions of others around them. On the other hand, positive effects like the wisdom of crowds [Sur05] may also

emerge. Higher-order beliefs may also come into play in a group setting (i.e., your beliefs about others' beliefs [MN13]). Collectively, these phenomena make the study of such dynamics in group settings particularly important.

Furthermore, the user's *graphical, statistical, and domain* expertise influence their ability to interact with and correctly interpret the presentation of information intended to elicit their beliefs. As a result, one must consider designs that suit the target audience's graphical literacy [OBD*06] by creating designs that are not overly complex. Statistical literacy may further impact the choice of data presentation (i.e., to minimize elicitation techniques that rely on advanced statistical concepts such as terminology about the user's confidence interval around the belief). Similarly, the interface should ideally provide sufficient background context based on the domain expertise of the target demographic.

4.1.2. Why

Why does the designer want to elicit a user's beliefs to begin with? What are the goals and motivations for eliciting beliefs [Ngu19]? For instance, one may elicit beliefs *to personalize visualizations* based on beliefs. Elicited beliefs can be useful to a researcher or designer to understand and empathize with the user better and provide more targeted information to the individual. Consider a belief-driven visualization related to factors contributing to climate change. Feedback on how to reduce one's carbon footprint could be personalized based on the user's prior beliefs and preferences for food consumption, transportation modalities, etc.

Visualization designers may also wish to elicit beliefs *to compare individuals' data with data about others' beliefs*. It can be insightful to users to learn about others' perspectives. For instance, in the New York Times "Where Is America's Heartland? Pick Your Map" [BQ17] example, the viewer is asked to pick a region of America that they believe is the heartland of the country. Viewers, perhaps surprisingly, see a diverse set of maps that demonstrate other viewers' perspectives. These example motivations are not comprehensive but rather illustrate a small set of possible reasons to elicit beliefs.

4.1.3. When

Figure 1 depicts belief elicitation in the context of (1) belief formation, (2) belief elicitation, (3) seeing new data, and (4) belief updating. **When** refers to whether the elicitation is based on a user's prior belief or based on their updated posterior belief (i.e., after seeing new data). Hence, if one goal of the belief-driven visualization involves understanding how viewers' beliefs *change*, then the visualization designer may opt to elicit both *Existing Beliefs* and *Updated Beliefs*.

4.1.4. Data set

To utilize the design space, we assume that the visualization designer first has a data set reflecting the domain of interest. Some factors that the designer may take into account while choosing a data set include the source and the accuracy (whether it is an objective ground truth or an estimate based on subjective opinions).

4.2. Interface considerations

Interface considerations relate directly to content and design choices for the belief-driven visualization interface. These choices will be influenced by the triangulation of the contextual considerations.

4.2.1. What

What is a measurable parameter of the belief? For example, if you want users to predict median home prices in their neighborhood over the next five years, you can represent it as a trend over time, or if you want to gauge users' beliefs on the top universities for Computer Science, you can elicit it as a ranking. Parametric considerations include the *Parameter Type* and whether or not *Uncertainty* about the belief is also elicited.

Parameter Types could be point estimates, distributions, rankings, etc. A point estimate is a singular value that best represents the data, including measures of central tendency (e.g., median value of home pricing), correlation or trend estimations (e.g., home pricing with relation to median household income), extreme values (e.g., most expensive homes in the county), binary values (e.g., whether a household with a median income in a county can qualify for a mortgage loan or not, or sample proportions (e.g., how many households can afford to buy a house in the county), among others. A distributional parameter could be discrete (e.g., number of semesters taken to graduate from different majors) or continuous (e.g., age of graduating students in a major). A ranking parameter involves defining relative or absolute ordering of some items in a list (e.g., rank order of graduating students by some characteristic such as GPA). These parameters represent a sample of possible types but is not an exhaustive list.

These parameters may be further described by their data type (categorical, numerical, spatial or temporal [S*46]). The choice of which parameter to elicit can be influenced by the motivation (**Why**) for analyzing belief data. For example, if the goal is to promote reflection on inflation, the designer can elicit a *trend* of numerical quantities related to cost of living over time. The choice of which parameter to elicit can also be influenced by the *statistical expertise* of the audience. For instance, eliciting a *point estimate* (such as a numerical value for slope) may not be appropriate for those with less statistical expertise, and a *binary* (such as positive v. negative slope) may be more fitting.

Uncertainty may be desirable to elicit the confidence level about the user's belief. Pang et al. have classified uncertainty visualization into three types: (1) statistical: mean and standard deviation to derive the confidence interval or a distribution (e.g., confidence level indicated on Glassdoor for their salary data [gla21]), (2) error: difference between the correct value and datum (e.g., forecast error is the difference between the actual and the forecasted value for a given period [Swa00]), and (3) range: an interval in which the data exists (e.g., Zillow's Zestimate range [zil21]) [PWL*97], which can inform the design of belief-driven visualizations that elicit uncertainty.

4.2.2. How

How can the belief elicitation be visually represented and what interactive components are most effective? In this dimension, we can

consider the *Visualization Technique*, the *Interaction Technique*, the *Inference Method* and, if eliciting uncertainty, the *Uncertainty Elicitation Method*.

For the *Visualization Technique*, some common representations that most user groups are familiar with include bar chart, pie chart, line chart, image, map, scatter plot, list, etc [BVB*13]. Beyond visualization literacy, the designer may consider how well the selected technique matches the user's mental model for the parameter [ZT99]. For instance, if eliciting the user's belief about a trend (e.g., ice cream sales over time), people tend to think about the trends visually and often in the form of a line chart. Hence, it is likely more effective to elicit patterns by having users set the heights of points in a *line chart* than to elicit a series of typed numerical values, or by eliciting via setting heights of bars. Another example is people's beliefs on the credibility of universities - people are likely most naturally familiar with providing relative orderings (e.g., via *ranked lists*) rather than approximating acceptance rates (e.g., via placing universities along a continuous scale such as in a *scatterplot*).

Similarly to visual representation, some *Interaction Techniques* are more familiar to a diverse audience than others. Some techniques in the context of the previous example about eliciting a trend could include drawing (e.g., drawing a trend line), clicking (e.g., plotting individual points on a line graph), typing (e.g., text input for a slope or a series of points), or drag (e.g., adjusting heights of points or bars). Other interaction techniques to consider include hover, zoom, slide, scroll, etc. These techniques can be treated as basic units to design compound and complex interactions as shown in Figure 6. Inclusion and accessibility can be one consideration while selecting one or a combination of these interaction techniques. Explicit instructions on how to interact with the visualization can ensure inclusion for users with lesser graphical expertise.

The way that a system *Infers* the belief may be directly reported by the user or, in the case of beliefs that may require more complex representations or reflection, the belief may be elicited indirectly from other self-reported answers. Indirect inference methods can sometimes be an effective way to gauge a more accurate and/or granular representation of one's beliefs. As an example, you may ask a user's beliefs about Americans' political party affiliations directly (e.g., Democrat, Republican, Independent), or you may elicit it indirectly and with greater detail (e.g., by asking along with a series of relevant social issues, as has been done by Pew Research Center[†]). HypTrails is another example approach for inferring people's beliefs and preferences indirectly through a trail of their digital transitions on the web (e.g., sequences of restaurant reviews) [SHHS15].

In the case of belief elicitation that incorporate *Uncertainty*, the same techniques can be optionally used to elicit uncertainty as well; however, the parameter of interest and uncertainty are typically distinctly elicited. Uncertainty may be elicited as a range of the parameter (e.g., an interval in which the aggregate carbon footprint of a state exists), or the eliciting the parameter over different intervals

rather than a large aggregate value (e.g., carbon footprint in different counties rather than carbon footprint in a whole state).

4.2.3. Feedback Post Belief Elicitation

Showing new data to the user is an optional interface consideration after eliciting the user's belief. How can facts be shared to meet the goals of the elicitation? This dimension is comprised of *type* of feedback and *source* of feedback. The feedback may be *Sourced* from ground truth data or from others' responses (social). The *Type of Feedback* can relate to *Veracity* (showing the correct data), *Deviation* (showing the comparison between the user input and correct data, or showing the comparison between user input and others' inputs), or *Explanation* (supporting the corrected data with details, what-if scenarios, data visualization, etc).

For instance, if one goal is to educate the audience, then you may wish to share ground truth data (*source*) after eliciting a viewer's belief about the trend of carbon emissions over the last decade in the form of the deviation from the true trend (*type*). On the other hand, it may be intriguing for users to learn about others' beliefs from social data (*source*), explained in detail with supporting text and visualizations (*type*).

5. Usage

In this section, we demonstrate how our design space can be used to characterize the existing examples of belief-driven visualizations in the wild and help design (new) or re-design (existing) belief-driven visualizations by walking through an example usage scenario.

5.1. Existing Belief Elicitation Work

News outlets are increasingly making use of belief elicitation in interactive articles as one method to engage readers and promote self reflection [HCHC20]. To understand the space of existing articles that employ belief elicitation and identify the under-explored approaches, we collected examples from a variety of news outlets, including 17 articles from New York Times [Tim20], 14 articles from FiveThirtyEight [Fiv21], 13 articles from CNN [CNN21], and 8 articles from Guardian [Gua21] covering a variety of subjects from journalism, psychology, and politics. We also included some examples of belief elicitation from other sources, including 1 article from Explorable Explanations [Exp21].

Many of the news outlets (and Explorable Explanations) publish summary lists of visual stories and graphics or interactive articles, from which we manually checked if each article contained a belief elicitation component. Some articles were excluded due to (1) the framing of the article was interactive exploration but did not contain any elicitation [ASH14, KS15]; (2) the design of the elicitation was in the form of a quiz or calculator, which was not belief-oriented [AK18, CHJ19]. The complete list of exemplary articles from this analysis is included in Supplemental Materials. After exclusions, we coded 24 examples in total within our design space, including the same 14 examples used in section 3.1.

Two authors independently coded each example. Any disagreements were resolved through discussion and by consulting a third author if necessary. The reliability of the coding depends on the

[†] <https://www.pewresearch.org/politics/quiz/political-typology/>

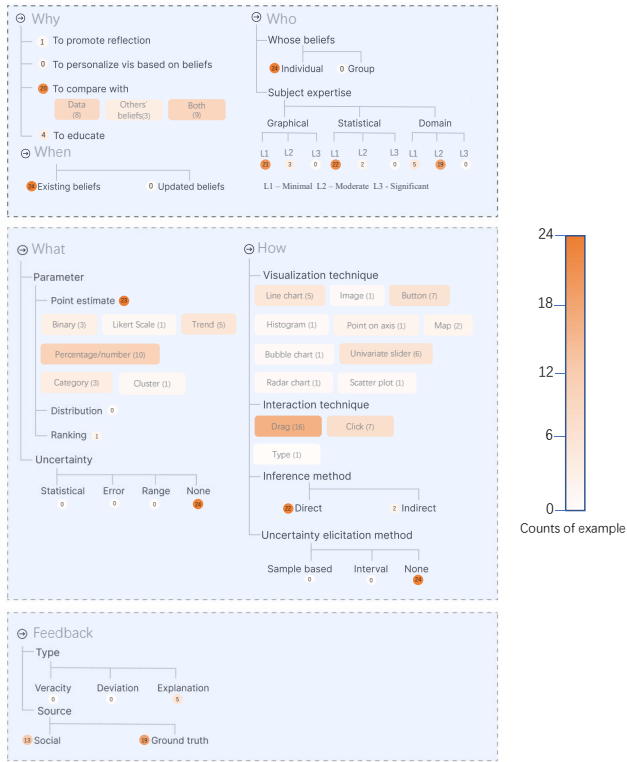


Figure 2: Characterizing belief-driven visualizations within the design space agreement level of both authors. Inter-rater reliability was assessed with Cohen’s KAPPA statistic, which controls for chance agreement [McH12] compared to measuring only percent agreement. A Cohen’s KAPPA of 0.833 was achieved for the codes, indicating high inter-rater reliability. More than half of the disagreements were related to the level of expertise required of the user (e.g., *minimal* v. *moderate* visualization expertise).

Figure 2 shows how existing examples span the design space. The intensity of the orange coloring indicates the number of belief-driven visualizations that were categorized accordingly. That is, darker orange indicates that more examples cover that dimension. Numbers inside circles and bars represent the counts of examples that lie in the corresponding components.

Results. By characterizing the 24 existing examples in our design space, we find that the examples cover a small fraction of our design space. For example, none of these examples from popular news media captured *uncertainty* (**What**) and there is a gap in the directness of techniques used to elicit beliefs (**How** → *Inference Method*). In addition, we observe that all examples elicited beliefs from *individuals* (**Who**) and were designed for people with minimal or moderate expertise (**Who** → *Subject Expertise*). These findings provide several trends and opportunities for belief-driven visualization, which we describe in Section 6.

5.2. Usage Scenario

Here, we will demonstrate a hypothetical scenario in which the design space can be used to iterate on the design of a variety of belief-

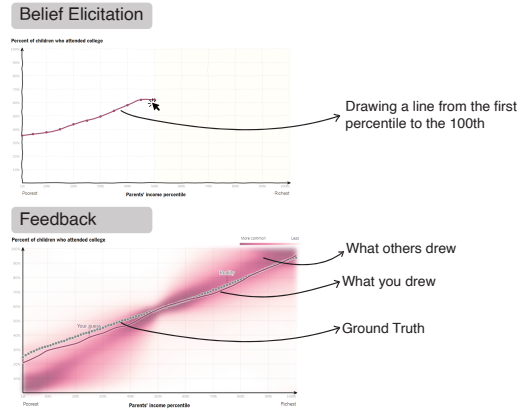


Figure 3: Example belief-driven visualization by New York Times.

driven visualizations, each with potential benefits and drawbacks. In this section, we will build upon an existing New York Times article with a belief-driven visualization starting with a dataset about the relationship between parents’ income percentile and children’s educational attainment. The article is titled “*You Draw It: How Family Income Predicts Children’s College Chances*” [ACQ15].

5.2.1. Background

The two first sentences of the article define the overarching questions behind the belief-driven article: “*How likely is it that children who grow up in very poor families go to college? How about children who grow up in very rich families?*” Inferring from these questions, we formalize that the **Why** behind this article is to educate individuals about the relationship between parents’ income and children’s educational attainment (**Contextual Considerations** → *Data set*). Since this is a NYT article, we can infer that this article is aimed at the NYT readership; that is, the **Who** is the general public with minimal to moderate *graphical* and *statistical* expertise. However, since the question is related to US demographics, engaging with this article would require moderate knowledge in the *domain* of US demographics. Furthermore, the article elicits beliefs before showing the data to gauge readers’ *existing beliefs* (**When**).

Given these questions, the article continues to provide a brief paragraph explaining how different relationships would be visually represented. Small sparklines in the text are used to describe the relationship between income and educational attainment (e.g., positive relationship, meaning higher income leads to higher educational attainment, is a positive sloped line). Additionally, the article defines more complicated relationships such as “situations in which chances level off after a certain income threshold” with a small sparkline that starts with a steep slope in the left and becomes horizontal in the right.

Following the tutorial, readers see an empty chart with two axes (Figure 3, top). The horizontal axis is labeled parents’ income percentile (from poorest to richest) and the vertical axis is labeled percent of children who attended college (**What** → *parameter*). Readers can draw their belief (**How** → *interaction technique*) about the relationship between these two parameters as a line (**How** → *visualization technique*). The line is editable and can be modified by

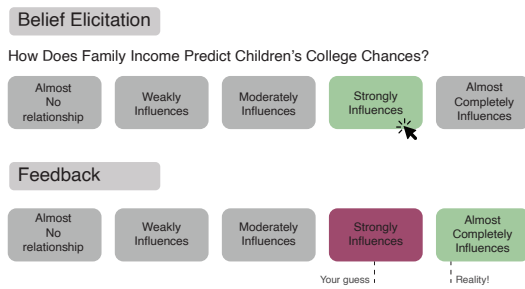


Figure 4: Midge's first sketch, a simple button-based design

clicking and dragging at fixed points along the line. There is also a point in the chart provided as a hint for users: "Free tip: Your line should go through this point".

After the line is drawn, the user can click on the *I'm Done* button to signify the conclusion of the belief elicitation. The article then updates with two new charts and several paragraphs of descriptive text. The first chart provides (**feedback** → **source**) about the differences between reader's expressed belief and the ground truth. The second chart shows "aggregate choices of [...] other New York Times readers" (**feedback** → **source**) through a heat map overlaid on the same chart (Figure 3, bottom). The remainder of the article includes additional details about the findings from the article and their interpretation. Additional sentences such as "Your line was not steep enough, there's more inequality than you guessed" and "Your line was relatively straight, reflecting one of the more striking findings of this research: the relationship between college enrollment and parental-income rank is linear" reinforce the provided visual feedback.

5.2.2. Leveraging the Design Space

Midge is a designer and journalist with a background in psychology. She is working for a marvelous and super hip online magazine. She pitches a belief-driven article about new research showing striking results about the relationship between parents' income levels and their children's educational attainment (**Data set**). She wants her designs to communicate the gravity of the inequality evident from this trend (**Why**). She starts sketching a few design options using a new design space for belief-driven visualizations. She knows that her designs will be seen by a large and diverse audience. So she assumes that her audience (**Who**) has very minimal *graphical, statistical, and domain* knowledge. She wants to gauge their beliefs *before* showing the data (**When**) to engage them early on in the article.

In her first sketch (Figure 4), she decides to use an ordinal parameter (**What** → **Parameter**) as a proxy for the *correlation* between the two variables. She uses a 5-button Likert scale design. She hopes her designs to be as easy to understand for the layperson as possible. Therefore after deliberating with different terms to verbally communicate the relationships such as "predicts" or "influences", she arrives at 5 choices ranging from *almost no relationship* to *almost completely influences*. She uses buttons as the *visualization technique* (**How**) and elicits users' belief through a *click* (**interaction technique**). Her design incorporates **feedback**




Figure 5: Midge's second sketch, eliciting beliefs and visualizing feedback using a continuous scale

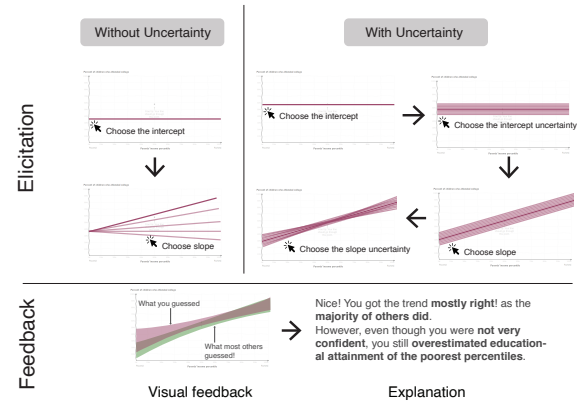


Figure 6: Midge's third sketch, drawing a line by choosing intercept and slope, receiving social and ground truth feedback.

by highlighting the difference (*type: deviation*) between users' choices and reality (*source: ground truth*) by highlighting the correct button.

After evaluating her designs, she realizes simplifying the relationship between these two variables might not communicate the striking findings of the research since with this design, it is not feasible to easily communicate the magnitude of the difference between users' beliefs and ground truth. Although this could provide *social* (**feedback** → **source**) through a bar chart encoding count of responses in each category (**How** → *visualization technique*), it would not capture the subtleties of users' beliefs. Additionally, the language around prediction would be unlikely to resonate with the *statistical expertise* of her readership (**Who**).

In her second iteration (Figure 5), she elects to *parameterize* users' beliefs (**What**) as a continuous number between 0 and 1 (*point estimate*). However, since her audience has minimal *statistical expertise* (**Who**), she uses the same wording used in her previous design. Users *interact* (**How**) with this elicitation technique by *dragging* a point on an axis that represents the continuous scale. *Ground truth* (**feedback** → **source**) can be shown to users as a point on the same scale, and *social feedback* can be shown to users as a range where the majority of users' choices lie. After evaluating this design, she finds that her current designs, omit an important part of the story she aimed to communicate with users. Although her second design elicited users' beliefs about the strength of the relationship (*correlation*), it did not reflect the massive difference between the poorest and wealthiest percentiles.

In her third iteration (See Figure 6), Midge elects to *parameterize (What)* and *visualize (How)* users' beliefs as the intercept and slope of a line in a chart with horizontal values encoding parents' income percentile and the vertical values encoding percent of children attending college. For her new design, she employs a *hover + click interaction technique (How)*. Users would first hover the mouse on the chart to select the intercept which is the percentage of children attending college for the lowest parents' income percentile (Figure 6 top-left). Users would then hover on the chart again to decide on the slope of the line, which highlights the rate in which educational attainment increases as we move up in each income percentile (Figure 6 top-left). In this design, *ground truth (feedback → source)* would be shown as a line of a different color, and *social (feedback → source)* is provided as a heatmap (Figure 6 middle and right).

As Midge iterated through the designs, she became interested in writing an article that provides different types of feedback based on different levels of uncertainty. One benefit of the current iteration of Midge's design, is that she can easily create a new elicitation technique to capture users' beliefs and *uncertainties (What)* around their beliefs. Using uncertainties, Midge designed *social (feedback → source)* that provides visual and textual *explanations (→ Type)* for users who are highly certain in their beliefs v. ones who are very uncertain (Figure 6 bottom). Uncertainties can be elicited by choosing a *range (What → Uncertainty)* around each selected parameter. The "with uncertainty" version of Midge's design adds two steps to the elicitation technique: (1) users first choose the intercept, then (2) the uncertainty around the intercept, then (3) the slope, and finally (4) the uncertainty around the slope (Figure 6 top-right).

Midge and the editorial team find the third iteration to be the most suitable for their goals; however, it does require greater *statistical expertise* of users (**Who**). Midge and the editorial team resolve to work on simplifying the surrounding phrasing and creating efficient tutorials. They test out their techniques with a small number of users to gauge the effectiveness of their final design. They hope to evaluate the usability and usefulness of the technique, and the extent to which this technique represents users' true beliefs.

6. Discussion

6.1. Applications Beyond Journalism

Given a framework or design space for belief-driven visualizations, there are numerous direct applications and extensions. As seen through our qualitative analysis of existing belief-driven journalistic articles, using such visual belief elicitation techniques allows for customizing the content shown to viewers. Given that users' belief congruence to the data they observe has an impact on how they update their beliefs [KMWD20, KRH17], policy makers, in public health for instance, can develop strategies to directly interact with individuals who have incongruent beliefs related to the information at hand (e.g., related to vaccination policies).

Furthermore, in many practices where decision makers need individual or aggregated input from stakeholders, such belief-driven visualizations can streamline the process of gathering direct input from stakeholders. For example, city planners can gauge the public's beliefs about an optimal location for a new train station by

eliciting individual beliefs about the impacts of the proximity to train stations to other aspects they find important.

We also envision that using belief-driven visualizations as part of visual analytics has the potential to improve users' decision-making in light of uncertain data. For instance, knowing an individual's beliefs about data prior to their analysis can enable a new wave of interfaces that can intelligently detect [WBF17] and mitigate [WSE19] suboptimal analysis behaviors consistent with e.g., confirmation bias [Nic98] (where an individual's search and interpretation of information prefers information consistent with pre-existing beliefs and discounts inconsistent information).

6.2. Trends and Opportunities

Based on our findings from coding existing belief-driven visualization examples from popular news media (Section 5.1), we identify several trends and opportunities for belief-driven visualization.

Eliciting from Broader Audiences. From a designer's perspective, there are a number of factors from the design space that have not been covered by examples from popular news media. For instance, *none* of the examples we identified captured *uncertainty (What)*, though some work in academic research does. Hence, there is an opportunity to design uncertainty elicitation techniques that are usable and understandable to broad audiences. Furthermore, we observe a gap in techniques that elicit beliefs *indirectly (How → Inference Method)*. Indirect inference methods may be able to decompose more complex beliefs in ways that cater to broader audiences; hence, this may be a promising direction for future work.

Eliciting from Specific Audiences. We further observe that all examples elicited beliefs from *individuals (Who)* (as is to be expected based on the typical context of individuals reading news articles). How might elicitation techniques engage multiple people in co-located or asynchronous settings, accounting for dynamics and biases of *groups*, to elicit truthful collective beliefs? Furthermore, the vast majority of examples required only *minimal* or *moderate* expertise. Hence, there is further opportunity to leverage the advanced expertise of some groups to explore potentially more accurate elicitation techniques that leverage advanced visualization or statistical knowledge of viewers.

Exploring Social Beliefs. While a few exceptions [ACQ15], most of the journalism examples do *not* incorporate social aspects of beliefs. However, notable research in the visualization field indicates that seeing other people's beliefs can influence your beliefs in interpreting data when they are not aligned with your own [KRH18]. Evidence beyond the visualization domain indicates that being exposed to opinions that are not aligned with your own can prompt you to think about the credibility of the information [Kim15], suggesting that social interventions can be useful to correct misinformation. Future work can further explore how a belief-driven visualization can leverage the function of other people's beliefs to maximize its benefit.

Validating Techniques. Finally, while some work is dedicated to evaluating the elicitation technique in other fields [GRI4], very little work exists that validates any of these techniques within the

context of visualizations (a notable exception being the line+cone technique for eliciting beliefs about correlations by Karduni et al. [KMWD20]). There is an opportunity to design experiments that can produce empirical results about the efficacy of different elicitation techniques for eliciting beliefs quickly, simply, and truthfully. Elicitation methods can be evaluated on the measurement properties of validity, reliability, responsiveness and feasibility, as suggested by Johnson et al [JTH*10].

This paper provides a literature- and data journalism-based approach to creating a design space for belief-driven visualizations. Here we discuss the limitations of our approach, applications of belief-driven visualization, and trends and opportunities.

6.3. Accurate Elicitation

One guiding consideration when choosing how to design the belief elicitation is how to ensure an accurate elicitation of someone's beliefs. Akin to limitations in accurate verbalization of thoughts in think-aloud protocols in usability studies [FST20], how to accurately capture cognitive phenomena is not trivial. However, there are a number of prior research efforts that can inform the design of elicitation interfaces that accurately capture beliefs.

For instance, prior work by behavioral economics on *scoring rules* may inform accurate elicitation. Scoring rules are used to incentivize people to quantify their opinions accurately and truthfully. Scoring rules measure the accuracy of probabilistic forecasting (e.g., the chance of rain fall in meteorology). A reward system designed around a proper scoring rule will incentivize the forecaster to report probabilities equal to their personal beliefs [Bic07]. Accuracy may be improved in the elicitation by considering incentives on accuracy [Wan11] that may promote more conscious reflection.

Furthermore, Kahneman and Tversky have described three heuristics people often use to make judgments of probability: availability; representativeness; and anchoring and adjustment [TK74]. These heuristics can be leveraged to increase the accuracy of the belief elicitation. For instance, it may result in a more accurate belief to guide users by giving one initial data point as an anchor for their assessment and then asking users to adjust remaining values with respect to this data point (reducing noise).

Similarly, consider framing the task by matching the user's expected mental model to the representation used in the interface. For instance, if people most often consider a trend visually as a line, then eliciting a series of numbers may be less accurate. In general, it may be difficult to accurately assess the probability of an event depending upon the numerical, statistical or domain expertise of the expert/user. Garthwaite et al. have reviewed and summarized psychological research on people's ability to estimate statistical summaries [GKO05], which may inform when to use a particular numerical parameter. In particular, people tend to estimate well binary data and sample proportions [Shu61], as well as measures of central tendency when distributions are symmetric [BS66]. Designers of belief-driven visualizations can consider these tradeoffs in people's competencies when developing elicitation interfaces.

6.4. Limitations

Our approach for deriving the design space for belief-driven visualizations provides a framework for describing existing and designing new belief-driven visualizations; however, our approach has at least 4 primary limitations. First, our design space spans a particular level of abstraction – we do not provide more precise guidance related to e.g., how to weigh and decide between alternative visual encoding choices.

Second, some aspects were scoped out of our design space, including specifics of incentive structure for eliciting truthful beliefs. Since most of these methods are self-reported, further comprehensive investigation on how to elicit truthful beliefs is needed. Another possibility could be to further explore indirect methods, where a belief can be observed or captured through behavior (e.g., Hyprails [SHHS15]) rather than self-reported.

Third, our design space was derived by iterating on examples from data journalism. However, there are a number of other adjacent areas that may be informative for belief-driven visualization design. We described the role of topics such as preference elicitation and knowledge elicitation, but these were not exhaustively surveyed nor explicitly included in the iterative coding and derivation of the design space. For instance, we explicitly scoped out elicitation articles that were not belief-driven (e.g., preference elicitation that uses buttons ask whether it is important to the user if the US wins Olympics gold medals: shorturl.at/kJQT0). The academic literature on graphical belief elicitation as well as both academic and adjacent journalism techniques (such as preference elicitation) can be informative nonetheless, and can inform design spaces that span a broader scope.

Finally and perhaps most glaringly, our focus is on *graphical* methods of belief elicitation using data visualization. However, there is further necessity to explore other modalities from the perspective of optimal elicitation techniques depending on the elicited parameter, as well as considering alternative modalities from an accessibility perspective.

7. Conclusion

In this paper, we synthesized a design space for creating belief-driven visualizations. Through iterative coding, revision and formative feedback from visualization experts, we developed a design space that addresses questions of *who, why, when, what, and how* to provide *feedback* to users' in belief-driven visualizations. We characterize 24 existing belief-driven visualizations collected from popular news media outlets along the dimensions of the derived design space and describe observed trends and opportunities in research on belief-driven visualization. Finally, we demonstrate how the design space might be used as a supplemental tool to designers through a usage scenario. This design space contributes to literature on belief-driven visualization by providing a formal framework with which to characterize existing belief-driven visualizations and structure the design of future belief-driven visualizations.

References

[ACQ15] AISCH G., COX A., QUEALY K.: You draw it: How family income predicts children's college chances - new york times.

- <https://www.nytimes.com/interactive/2015/05/28/upshot/you-draw-it-how-family-income-affects-childrens-college-chances.html>, May 2015. (Accessed on 11/20/2021). 1, 2, 3, 7, 9
- [AFZ19] ARMONA L., FUSTER A., ZAFAR B.: Home price expectations and behaviour: Evidence from a randomized information experiment. *The Review of Economic Studies* 86, 4 (2019), 1371–1410. 2
- [AK18] ASCHWANDEN C., KOEZE E.: How fast would you run a marathon? - fivethirtyeight. <https://projects.fivethirtyeight.com/marathon-calculator/>, November 2018. Accessed on 07/27/2021. 6
- [ASH14] ASHKENAS J.: Can you live on the minimum wage? - new york times. <https://www.nytimes.com/interactive/2014/02/09/opinion/minimum-wage.html>, February 2014. Accessed on 07/27/2021. 6
- [BDKS19] BAILEY M., DÁVILA E., KUCHLER T., STROEBEL J.: House price beliefs and mortgage leverage choice. *The Review of Economic Studies* 86, 6 (2019), 2403–2452. 2
- [Bic07] BICKEL J. E.: Some comparisons among quadratic, spherical, and logarithmic scoring rules. *Decision Analysis* 4, 2 (2007), 49–65. 10
- [BQ17] BADGER E., QUEALY K.: Where is america's heartland? pick your map - new york times. <https://www.nytimes.com/interactive/2017/01/03/upshot/where-is-americas-heartland-pick-your-map.html>, January 2017. (Accessed on 11/20/2021). 3, 5
- [BS66] BEACH L. R., SWENSON R. G.: Intuitive estimation of means. *Psychonomic Science* 5, 4 (1966), 161–162. 10
- [BVB*13] BORKIN M. A., VO A. A., BYLINSKII Z., ISOLA P., SUNKAVALLI S., OLIVA A., PFISTER H.: What makes a visualization memorable? *IEEE transactions on visualization and computer graphics* 19, 12 (2013), 2306–2315. 6
- [CHJ19] CLARKE S., HULLEY-JONES F.: Find out which brexit deal is right for you - guardian. <https://www.theguardian.com/politics/ng-interactive/2019/jan/18/find-out-which-brexite-deal-is-right-for-you>, January 2019. Accessed on 07/27/2021. 6
- [CNN21] CNN: Cnn interactives. <https://www.cnn.com/specials/multimedia/cnn-interactives>, 2021. Accessed on 07/27/2021. 6
- [CP04] CHEN L., PU P.: *Survey of preference elicitation methods*. Tech. rep., 2004. 2
- [Exp21] EXPLORABLE: Explorable explanations. <https://explorabl.es/>, 2021. Accessed on 07/27/2021. 6
- [Fiv21] FIVETHIRTYEIGHT: Fivethirtyeight interactives. <https://projects.fivethirtyeight.com/>, 2021. Accessed on 07/27/2021. 6
- [Fod83] FODOR J. A.: *Representations: Philosophical essays on the foundations of cognitive science*. Mit Press, 1983. 2
- [FST20] FAN M., SHI S., TRUONG K. N.: Practices and challenges of using think-aloud protocols in industry: An international survey. *Journal of Usability Studies* 15, 2 (2020). 10
- [Gav93] GAVRILOVA T.: Choice of knowledge elicitation technique: the psychological aspect. *International Journal on Information Theory and Applications* 1, 8 (1993), 20–26. 2
- [GKO05] GARTHWAITE P. H., KADANE J. B., O'HAGAN A.: Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association* 100, 470 (2005), 680–701. 2, 3, 10
- [gla21] Company salaries. *Glassdoor* <https://www.glassdoor.com/Salaries/index.htm>, 2008–2021. 5
- [GR14] GOLDSTEIN D. G., ROTHSCILD D.: Lay understanding of probability distributions. *Judgment & Decision Making* 9, 1 (2014). 9
- [Gua21] GUARDIAN T.: Guardian interactives. <https://www.theguardian.com/interactive>, 2021. Accessed on 07/27/2021. 6
- [HCHC20] HOHMAN F., CONLEN M., HEER J., CHAU D. H.: Communicating with interactive articles - distill. <https://distill.pub/2020/communicating-with-interactive-articles/>, 2020. Accessed on 07/27/2021. 6
- [HKKS18] HULLMAN J., KAY M., KIM Y.-S., SHRESTHA S.: Imagining replications: Graphical prediction & discrete visualizations improve recall & estimation of effect uncertainty. *IEEE transactions on visualization and computer graphics* 24, 1 (2018), 446–456. 2
- [HRR20] HEYER J., RAVEENDRANATH N. K., REDA K.: Pushing the (visual) narrative: the effects of prior knowledge elicitation in provocative topics. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020), pp. 1–14. 2
- [HSBK95] HOFFMAN R. R., SHADBOLT N. R., BURTON A. M., KLEIN G.: Eliciting knowledge from experts: A methodological analysis. *Organizational behavior and human decision processes* 62, 2 (1995), 129–158. 2
- [Hun96] HUNTER D.: On the relation between categorical and probabilistic belief. *Noûs* 30, 1 (1996), 75–98. 2
- [Jan08] JANIS I. L.: Groupthink. *IEEE Engineering Management Review* 36, 1 (2008), 36. 4
- [JTH*10] JOHNSON S. R., TOMLINSON G. A., HAWKER G. A., GRANTON J. T., FELDMAN B. M.: Methods to elicit beliefs for bayesian priors: a systematic review. *Journal of clinical epidemiology* 63, 4 (2010), 355–369. 10
- [KBS*21] KOONCHANOK R., BASER P., SIKHARAM A., RAVEENDRANATH N. K., REDA K.: Data prophecy: Exploring the effects of belief elicitation in visual analytics. 3
- [Kim15] KIM Y.: Exploring the effects of source credibility and others' comments on online news evaluation. *Electronic News* 9, 3 (2015), 160–176. 9
- [KKGMMH20] KIM Y.-S., KAYONGO P., GRUNDE-MCLAUGHLIN M., HULLMAN J.: Bayesian-assisted inference from visualized data. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 989–999. 2, 3
- [KMW20] KARDUNI A., MARKANT D., WESSLEN R., DOU W.: A bayesian cognition approach for belief updating of correlation judgement through uncertainty visualizations. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 978–988. 3, 9, 10
- [KRH17] KIM Y.-S., REINECKE K., HULLMAN J.: Explaining the gap: Visualizing one's predictions improves recall and comprehension of data. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017), ACM, pp. 1375–1386. 1, 2, 9
- [KRH18] KIM Y.-S., REINECKE K., HULLMAN J.: Data through others' eyes: The impact of visualizing others' expectations on visualization interpretation. *IEEE transactions on visualization and computer graphics* 24, 1 (2018), 760–769. 9
- [KS15] KING R., SILVER N.: Which flight will get you there fastest? - fivethirtyeight. <https://projects.fivethirtyeight.com/flights/>, November 2015. Accessed on 07/27/2021. 6
- [KWKH19] KIM Y.-S., WALLS L. A., KRAFFT P., HULLMAN J.: A bayesian cognition approach to improve data visualization. In *Proceedings of the 2019 CHI conference on human factors in computing systems* (2019), pp. 1–14. 1, 2, 3
- [Man04] MANSKI C. F.: Measuring expectations. *Econometrica* 72, 5 (2004), 1329–1376. 2
- [MBW99] MOODY J. W., BLANTON J. E., WILL R. P.: Capturing expertise from experts: The need to match knowledge elicitation techniques with expert system types. *Journal of Computer Information Systems* 39, 2 (1999), 89–95. 2
- [McH12] MCHUGH M. L.: Interrater reliability: the kappa statistic. *Biochemia medica* 22, 3 (2012), 276–282. 7
- [MN13] MANSKI C. F., NERI C.: First-and second-order subjective expectations in strategic decision-making: Experimental evidence. *Games and Economic Behavior* 81 (2013), 232–254. 5

- [MPLW19] MOSKIN J., PLUMER B., LIEBERMAN R., WEINGART E.: Your questions about food and climate change, answered - new york times. <https://www.nytimes.com/interactive/2019/04/30/dining/climate-change-food-eating-habits.html>, April 2019. (Accessed on 11/20/2021). 3
- [MW18] MEHTA D., WOLFE J.: Do you know where america stands on guns? - fivethirtyeight. <https://projects.fivethirtyeight.com/guns-parkland-polling-quiz/>, March 2018. (Accessed on 11/20/2021). 3
- [Ngu19] NGUYEN F.: Belief-driven data journalism. *Computation+ Journalism* (2019). 3, 5
- [Nic98] NICKERSON R. S.: Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology* 2, 2 (1998), 175–220. 9
- [OBD*06] O'HAGAN A., BUCK C. E., DANESHKHAH A., EISER J. R., GARTHWAITE P. H., JENKINSON D. J., OAKLEY J. E., RAKOW T.: *Uncertain judgements: eliciting experts' probabilities*. John Wiley & Sons, 2006. 2, 5
- [PBS*99] PAYNE J. W., BETTMAN J. R., SCHKADE D. A., SCHWARZ N., GREGORY R.: Measuring constructed preferences: Towards a building code. In *Elicitation of preferences*. Springer, 1999, pp. 243–275. 2
- [PBW*12] POMMERANZ A., BROEKENS J., WIGGERS P., BRINKMAN W.-P., JONKER C. M.: Designing interfaces for explicit preference elicitation: a user-centered investigation of preference representation and elicitation process. *User Modeling and User-Adapted Interaction* 22, 4 (2012), 357–397. 2
- [PCH12] PU P., CHEN L., HU R.: Evaluating recommender systems from the user's perspective: survey of the state of the art. *User Modeling and User-Adapted Interaction* 22, 4-5 (2012), 317–355. 2
- [PWL*97] PANG A. T., WITTENBRINK C. M., LODHA S. K., ET AL.: Approaches to uncertainty visualization. *The Visual Computer* 13, 8 (1997), 370–390. 5
- [S*46] STEVENS S. S., ET AL.: On the theory of scales of measurement. 5
- [SB15] SCHMITT-BECK R.: Bandwagon effect. *The international encyclopedia of political communication* (2015), 1–5. 4
- [SHHS15] SINGER P., HELIC D., HOTHO A., STROHMAIER M.: Hyp-trails: A bayesian approach for comparing hypotheses about human trails on the web. In *Proceedings of the 24th International Conference on World Wide Web* (2015), pp. 1003–1013. 6, 10
- [Shu61] SHUFORD E. H.: Percentage estimation of proportion as a function of element type, exposure time, and task. *Journal of Experimental Psychology* 61, 5 (1961), 430. 10
- [SKH04] SIMON D., KRAWCZYK D. C., HOLYOAK K. J.: Construction of preferences by constraint satisfaction. *Psychological Science* 15, 5 (2004), 331–336. 2
- [Sur05] SUROWIECKI J.: *The wisdom of crowds*. Anchor, 2005. 4
- [Swa00] SWAMIDASS P. M.: *Forecast Errors: Encyclopedia of production and manufacturing management*. Springer Science & Business Media, 2000. 5
- [Tim20] TIMES N. Y.: 2020: The year in visual stories and graphics. <https://www.nytimes.com/interactive/2020/12/30/us/2020-year-in-graphics.html>, December 2020. Accessed on 07/27/2021. 6
- [TK74] TVERSKY A., KAHNEMAN D.: Judgment under uncertainty: Heuristics and biases. *science* 185, 4157 (1974), 1124–1131. 10
- [VGGT14] VUL E., GOODMAN N., GRIFFITHS T. L., TENENBAUM J. B.: One and done? optimal decisions from very few samples. *Cognitive science* 38, 4 (2014), 599–637. 3
- [Wan11] WANG S. W.: Incentive effects: The case of belief elicitation from individuals in groups. *Economics Letters* 111, 1 (2011), 30–33. 10
- [Wat85] WATERMAN D. A.: *A guide to expert systems*. Addison-Wesley Longman Publishing Co., Inc., 1985. 2
- [WBF17] WALL E., BLAHA L. M., FRANKLIN L., ENDERT A.: Warning, bias may occur: A proposed approach to detecting cognitive bias in interactive visual analytics. In *2017 IEEE Conference on Visual Analytics Science and Technology (VAST)* (2017), IEEE, pp. 104–115. 9
- [WJ06] WEBER E. U., JOHNSON E. J.: Constructing preferences from memory. *The Construction of Preference, Lichtenstein, S. & Slovic, P., (eds.)* (2006), 397–410. 2
- [WSE19] WALL E., STASKO J., ENDERT A.: Toward a design space for mitigating cognitive bias in vis. In *2019 IEEE Visualization Conference (VIS)* (2019), IEEE, pp. 111–115. 9
- [zil21] What is a zestimate? zillow's zestimate accuracy. *Zillow*, May 27, 2021, <https://www.zillow.com/zestimate/>, 2021. 5
- [ZT99] ZACKS J., TVERSKY B.: Bars and lines: A study of graphic communication. *Memory & cognition* 27, 6 (1999), 1073–1079. 6