

Data Hunches: Incorporating Personal Knowledge into Visualizations

Haihan Lin, Derya Akbaba, Miriah Meyer, and Alexander Lex




Fig. 1: A representation of the process of recording and collaborating with data hunches. When an analyst looks at a data visualization and has a hunch about the data, they are able to externalize the data hunch through visual methods that we introduce in this paper. Following the recording of their data hunch, others can review the hunches and potentially record their own.

Abstract—The trouble with data is that it frequently provides only an imperfect representation of a phenomenon of interest. Experts who are familiar with their datasets will often make implicit, mental corrections when analyzing a dataset, or will be cautious not to be overly confident about their findings if caveats are present. However, personal knowledge about the caveats of a dataset is typically not incorporated in a structured way, which is problematic if others who lack that knowledge interpret the data. In this work, we define such analysts’ knowledge about datasets as *data hunches*. We differentiate data hunches from uncertainty and discuss types of hunches. We then explore ways of recording data hunches, and, based on a prototypical design, develop recommendations for designing visualizations that support data hunches. We conclude by discussing various challenges associated with data hunches, including the potential for harm and challenges for trust and privacy. We envision that data hunches will empower analysts to externalize their knowledge, facilitate collaboration and communication, and support the ability to learn from others’ data hunches.

Index Terms—Data Visualization, Uncertainty, Data Hunches

1 INTRODUCTION

While data-driven decision-making and reasoning is now considered the gold standard across many fields such as business [10], sports [64], and science [22], data workers are also aware of the incompleteness and imperfections of data [8, 50]. In scientific data analysis, for example, caveats about the data are often captured in (digital or paper) lab notebooks, as comments in analysis scripts, and as datasheets recorded alongside a dataset [26], and they are discussed with colleagues in lab meetings, and reported to a general audience in methods sections of scientific papers.

These processes, however, detach a person’s knowledge about caveats to the data from the tools in which the data is analyzed. For example, if a caveat about the data is recorded in a graduate student’s lab notebook, their successor who is asked to re-analyze the data may not have access to the notes and may interpret the data differently. Even if a caveat is noted in the methods section of a paper, a casual reader might skip that section and misinterpret a visualization of the data.

We encountered this challenge in a design study where we collaborated with clinicians who analyze data about blood transfusions to improve patient outcomes and reduce the use of valuable blood-bank resources [43]. When we showed a clinical collaborator a prototype of our visualization tool, he expressed concern about what he saw in the data: that the amount of recycled blood—a patient’s own blood that is reused during surgery—was much lower than he expected. In his

experience, almost all surgeons make extensive use of blood recycling; he had a hunch that the low blood recycling values were due to the data frequently not being recorded during surgeries. A second collaborator focusing on patient blood management was promoting a more data-driven approach in his workplace and worried that when clinicians saw these discrepancies, they would lose trust in the visualization tool. Despite these concerns, we had no way of recording and communicating the clinicians’ hunches in the visualization.

Just like in this example, we frequently find that interactions with a visualization tool trigger experts to express knowledge about problems with data, but that these tools leave few options for them to communicate that knowledge. This knowledge is personal and not available to others, a phenomenon reported in other design studies [47, 51, 54]. However, what if experts could record their hunches directly in a visualization tool in a way that allowed others to interpret the data alongside their hunches? What if a visualization tool supported—and encouraged—the explicit incorporation of expert knowledge during data analysis?

In this paper, we address these questions by proposing a new concept, *data hunches*, to describe the knowledge that analysts bring to their data analysis process that augments and complements their perception of the data, and in turn becomes part of their interpretations as demonstrated in Figure 1. Data hunches are personal knowledge about caveats to the data, such as knowledge about implicit errors [47] or ambiguities [51]. We argue that conceptualizing data hunches provides a new perspective on how to design visual analysis tools for expert, collaborative settings, allowing for the sharing of knowledge about the ways that data is imperfect, partial, and uncertain. This new perspective opens design possibilities for how we might build visualization tools that support both the recording and communicating of hunches.

To this end, our work includes three core contributions:

- Haihan Lin is with the University of Utah. Email: hmlin@sci.utah.edu
- Derya Akbaba is with the University of Utah. Email: derya@cs.utah.edu
- Miriah Meyer is with the Linköping University. Email: miriah.meyer@liu.se
- Alexander Lex is with the University of Utah. Email: alex@sci.utah.edu

- **A conceptualization of data hunches:** We define the term *data hunches* and discuss its relationship to uncertainty, elevating the role of personal knowledge in visual data analysis.
- **An exploration of ways to record and communicate data hunches:** We demonstrate how data hunches can be recorded and communicated within a visualization context using visual techniques to facilitate collaboration.
- **A set of design recommendations for data hunches:** We reflect on design and ethical considerations for supporting data hunches within visualization systems.

We speculate how recording and communicating data hunches could be implemented through a demo tool, available at <https://vdl.sci.utah.edu/data-hunch/>. For transparency and rigor, we include a brief discussion of our methodology in Section 8. We discuss the potential of data hunches, for both good and bad outcomes in Section 10. Despite open questions about how to design for data hunches in ethical and usable ways, we consider this work a step toward a wealth of productive opportunities for valuing and including personal knowledge in visual data analysis.

2 UNCERTAINTY

Measurement errors, modeling assumptions, heterogeneous data recording methods, and missing context are just some of the ways in which values stored in a dataset are neither a perfect nor complete view of a phenomenon in the world. The visualization community has a long history of researching methods to characterize, quantify, and visualize the limitations of data under the heading of *uncertainty*. Within this literature, explicit definitions for uncertainty are sparse, with the definitions that do exist covering a range of interpretations. For example, Hullman refers to uncertainty as “the possibility that the observed data or model predictions could take on a set of possible values” [32], whereas Bonneau et al. state that “uncertainty is the lack of information” [6]. Fields outside of visualization add additional interpretations of the term, such as Walker et al.’s general notion of uncertainty as “any deviation from the unachievable ideal of completely deterministic knowledge of the relevant system” [68], and Covitt et al.’s experiential definition of uncertainty as the “ways in which scientists recognize and analyze limits in their studies and conclusions” [14].

Researchers have openly acknowledged the ill-defined nature of uncertainty, such as Boukhelifa et al.’s statement that “there is no unified single definition of uncertainty across all domains. The general consensus is that there are different meanings and that the term itself encapsulates many concepts” [8]. Brodlić et al. go further and argue that the lack of a clear, consensus definition has held the field back: “the self-referential problem of uncertainty about uncertainty terminology has been a notable stumbling block in this avenue of inquiry” [9].

In this paper, we do not attempt to rectify the uncertainty about uncertainty. Rather, we point out the difficulty of precisely positioning theoretical perspectives in relation to existing uncertainty literature. Despite this challenge, we situate data hunches within the visualization uncertainty literature because this body of work has focused on understanding and characterizing how people, visualizations, and imperfect and partial data come together. More specifically, the uncertainty literature focuses on the ways that data is limited representations of the world, and how people (can) become aware of these limitations. Such perspective is in contrast to other visualization research threads that focus on expert knowledge more broadly, such as work on insight that considers different types of expert knowledge that impacts insight generation [38]. In this section, we briefly summarize the literature on quantitative and qualitative uncertainty, and argue that data hunches complement the existing work in this area.

2.1 Quantitative Uncertainty

Researchers provide a variety of characterizations of uncertainty through descriptions of the many ways that data can be uncertain. Potter et al. [58] use a characterization from computational sciences that describes *epistemic* versus *aleatoric* uncertainty, with the former describing the ways in which a lack of knowledge about and from the data induce a computationally unknowable uncertainty, and the latter encompassing data limitations that can be assumed and modeled

statistically. Padilla et al. compare *direct quantitative* uncertainties and *indirect qualitative* uncertainties [53]. In this framing, direct uncertainties are quantifiable expressions such as confidence intervals and probability distributions, whereas indirect uncertainties can be expressed only qualitatively. Direct, quantifiable expressions of uncertainty are typically computed from sources of imperfections and partialities, including limited data collection resources, such as not being able to sample every person in a population of interest; limited measurement capabilities, such as the precision of an instrument; or limited knowledge about the future, such as the unpredictability of forecasting weather [63].

The visualization community has historically focused on developing and testing methods for visualizing quantifiable uncertainty [6, 52, 58]. Some approaches attempt to intuitively encode uncertainty through modifications of a data item’s graphical mark using quantile dot plots [53], glyphs [37, 52, 73], or value-suppressed color schemes [13]. Other approaches have instead explored visual representations that directly display summary statistics [12, 56, 57] or use animations showing hypothetical outcomes [34]. Researchers have also developed uncertainty-specific evaluation techniques, such as eliciting users’ internal models of probability distributions, recording the effects of uncertainty on decision-making, and assessing participants’ sense of confidence after viewing uncertainty visualizations [33].

The community’s characterization of what types of uncertainty are quantifiable would seemingly place our blood-reuse example in Section 1 as something that could be quantified, a limitation of the measurement capabilities. In principle, we could attempt to model and quantify this limitation, but any metric is likely to be grossly inaccurate because of the abstracted nature of the knowledge about the imperfections, a point raised by Thomson et al.: “In addition to uncertain measures, analysts are concerned with abstract uncertainties such as the credibility of a particular source or the completeness of a set of information. As the uncertainty becomes more abstract, it is more difficult to quantify, represent, and understand” [63]. Instead, researchers adopt different perspectives for abstract sources of uncertainty—qualitative perspectives—that focus on the knowledge people have about the limitations of data.

2.2 Qualitative Uncertainty

Qualitative uncertainty—also referred to as indirect or epistemic uncertainty—has been described as “the quality of knowledge concerning how effectively facts, numbers, or hypotheses represent reality” [53]. Definitions of qualitative uncertainty make explicit references to knowledge, shifting the emphasis from exploitable information about the data, to inaccessible subjective knowledge: “Epistemic uncertainty generally, but not always, concerns past or present phenomena that we currently don’t know but could, at least in theory, know or establish” [65]. In contrast to quantitative uncertainty, qualitative uncertainty is not easily quantified, and is generally conveyed through “caveats about data” [65].

The sources of qualitative uncertainty stem from the same imperfections and partialities that metrics for quantifying uncertainty pertain to. Boukhelifa et al. [8] classify these sources as imperfect, messy, and missing *data*; imperfect and limited *models*; (approximate) digital representations of data in a visualization *interface*; and *cognitive* differences between interpretations from individual analysts. Most relevant for the work in this paper are sources of qualitative uncertainty from data, defined by McCurdy et al. as *implicit error*: “a type of measurement error that is inherent to a dataset but not explicitly recorded, yet is accounted for qualitatively by experts during analysis, based on their implicit domain knowledge” [47].

The predominate way that visualization designers encode qualitative expressions of uncertainty is through text-based annotations. In her interview study with visualization practitioners, Hullman reports that “uncertainty as a qualitative expression of a gap in knowledge came up in most interviews with interviewees as well as several survey responses. 62% of survey respondents had used text to warn their viewers of the potential for uncertainty in results” [32]. Other approaches use visual approaches to communicate qualitative uncertainty, such as the use of perceptually imprecise visual encoding channels like sketchiness [7] or glyphs [51]. A different approach taken in both the visualization

and machine learning communities is to explicitly expose information about the data collection process, providing analysts with contextual information that allows them to incorporate personal knowledge about potential shortcomings of the data during their interpretations [2,26,54].

Several visualization systems explore ways for recording expert knowledge about qualitative sources of uncertainty. For example, in their tool for supporting public health experts, McCurdy et al. [47] designed a template with structured questions that enabled the experts to record what they knew about implicit errors in the data. The recorded results were marked on the visualizations with glyphs that provided annotations when clicked. Similarly, Franke et al. [23] collected confidence about data sources from historians through a web interface. These varying levels of confidence were then presented alongside other data in a hierarchical tree view, representing the distribution of confidence along different dimensions of the data source in question.

These two examples demonstrate that experts *know* about limitations of their data; that the sources of uncertainty can be known. This, however, is in contrast to existing definitions of qualitative uncertainty that position it as something unknowable: “epistemic uncertainty describes uncertainties due to lack of knowledge and limited data which could, in principle, be known, but in practice are not” [58]. This focus on the unknowable also appears in more general uncertainty characterizations [6]. This contradiction leads us to ask the question: unknown by whom?

The work we present in this paper is a significant shift in how we think about designing for qualitative uncertainty. Data hunches frame the knowledge about sources of uncertainty away from the visualization tool designers, to the experts who conduct the analysis. The concept of data hunches is an explicit acknowledgment that many sources of qualitative uncertainty are in fact known—known to the experts who can articulate the knowledge when triggered by their interactions with a visualization tool. This shift complements the visualization community’s perspectives on uncertainty by focusing on knowledge about sources of uncertainty, and thus opening up opportunities to design new ways that visualizations can support recording and communicating this knowledge during data analysis.

3 DATA HUNCHES

Our community’s current framing of uncertainty implies that data workers and tool builders are responsible for identifying, characterizing, and quantifying sources of uncertainty from data. Data hunches instead acknowledge and incorporate experts who come to data analysis with deep knowledge about the limitations of their data. As a complementary perspective, data hunches focus on knowledge about data, capture that knowledge from a diverse set of stakeholders, and are embeddable in the analysis process (instead of the data curation pipeline). We argue that this perspective offers a breadth of new opportunities for recording and communicating data hunches in support of richer data analysis.

Through data hunches, we elevate the role that personal knowledge of the data plays in the process of understanding and analyzing it. More precisely, **we define data hunches as an analyst’s knowledge about how and why the data is an imperfect and partial representation of the phenomena of interest.** These hunches can range from the abstract—expressing concern about the validity of the data set—to the concrete—expressing a numerical value that is closer to the phenomenon of interest than the measured data. The scope of a data hunch can be individual data points, a complete dataset, or anything in between. Data hunches emerge when an analyst interacts with the data, triggering reflection about the ways the data is imperfect and partial based upon their inherent knowledge about the data collection process, domain, and more. A data hunch can be based on the missing context necessary to fully comprehend the phenomenon, discrepancies between a mental model and data, opinions on the quality of the data generation process, and so on. Data hunches are knowledge about sources of qualitative uncertainty. During data analysis, data hunches influence an analyst’s interpretation of the data, derived knowledge, and decisions made.

We argue that data hunches are prevalent, but often implicit, in data analysis, and as important as the data itself. The knowledge experts

bring to data analysis is a vital component of data-driven decision-making [38,49]; however, such knowledge is often recorded outside of visual analysis tools. This disconnect then requires mental gymnastics on the part of an analyst to incorporate back into the data analysis process, if it is not overlooked completely. In this paper, we envision that a set of tools that fully utilize what visualization can offer will provide a more visual and intuitive representation of experts’ knowledge in visualizations. Using data and data hunches in tandem supports a richer representation of a phenomenon, leading potentially to improved analysis. By acknowledging and naming data hunches, we aim to elevate the potential for personal knowledge to actively and explicitly contribute to data analysis.

We purposefully scope data hunches for use within collaborative, expert settings for both pragmatic and ethical reasons. Pragmatically, previous work on collaborative visualizations highlights the value of designing tools that support sharing of expert knowledge. In the context of collaborative analysis sessions, Mahyar et al. [44] showed the importance of recording visualizations and note taking in collaborative visual analysis. Similarly, Walny et al. [69] studied the use of data visualizations on whiteboards in corporate offices and found that visualizations as sketches promote team discussions. In another example of data science workflows that utilize computational notebooks, Wang et al. [71] found that data scientists annotated screenshots of visualizations when collaborating as a way to communicate limitations of a tool. Ethically, scoping data hunches to collaborative expert settings reduces potential harm as we expect experts collaborating with peers to recorded well-reasoned and nuanced hunches. However, experts can still be biased, and data hunches may help to reinforce biases even further. We believe that peer-review tools, such as comments or ratings, may be useful to mitigate the risk of reinforcing biases, yet acknowledge that more research on the topic is necessary. We discuss these and other issues further in Section 10.

By identifying data hunches as productive and insightful expert knowledge, we can re-interpret past work on collaborative visualization tools with this framing. For example, a visualization designer can incorporate commenting and discussion features to promote externalization of data hunches [30,47]; apply provenance tracking to record actions they took based on hunches to wrangle the data [15,19,24,25]; use visualization techniques like linked views and visualization states to show a collection of data hunches [35,46,67]. We see a wealth of opportunities for incorporating data hunches into old and new ways of visually analyzing data.

4 TYPES OF DATA HUNCHES

In our development of data hunch classifications, we listed all the data hunches we encountered in our previous design studies, brainstormed other forms of data hunches that may surface in data analysis, and generalized them based on whether they are qualitative or quantitative, the specificity of the hunches, and the forms they take if expressed. We identify three types of data hunches in support of determining suitable methods for reporting and communicating expert knowledge about limitations of data.

Assessment Hunches: Assessment data hunches speak to the trustworthiness or quality of a dataset or individual data items, or simply provide context. Assessment hunches can take different forms, ranging from ratings (thumbs-up/down, numerical scores, etc.), to written comments about data items or datasets.

Structural Hunches: Structural data hunches state that certain data points or relationships should not be included in the dataset (exclusion) or that a data item or relationship is missing (inclusion). In a network dataset, for example, an inclusion data hunch could be used to indicate that an edge is missing. For many data types, inclusion hunches should also be combined with an estimate of a value of the included hunch. For example, when indicating that an item is missing from a dataset, the data hunch could also contain an estimate of the value of the item.

Value Hunches: Value hunches make a statement about how a specific data value should be different from what is recorded in the dataset. Value hunches apply equally to numerical, categorical, and textual/label data. For example, a value hunch for a category could state that an item should be in category A instead of category B.




Fig. 2: A set of abstract techniques to record data hunches. We distinguish three spaces: data, visual, and abstract. Recording in **data space**, such as through (a) form-based manipulation and (b) model-based manipulation, affords a basic technique where analysts can externalize hunches by either manually inputting items, or expressing them through models. Recording in **visual space**, such as through (c) free-form sketching and (d) direct manipulation, affords visual recording of data hunches, ranging from sketching to dragging graphical elements to represent an analyst’s hunch. Recording in **abstract space**, such as through (e) structured elicitation and (f) textual annotations, enables analysts to record hunches that cannot be recorded in the data and visual spaces.

For practical reasons, we found it useful to further distinguish three methods of expressing value hunches, reflecting different levels of “precision” about a value hunch. *Directional Hunches* express that values should be different (higher or lower) from the recorded value without giving a specific value. They are a middle ground between assessment hunches, which make no statement about directionality, and hunches that give estimates for actual values. *Specific Value Hunches* express exactly how values in a dataset should be different. For example, a value data hunch could state that the value encoded by a bar chart should be 20 instead of 12. *Value Range Hunches* acknowledge uncertainty about the value to be specified. Instead of expressing a specific value, they state that a value should be within a certain range. For example, a range hunch could state that the value of a bar should be between 18 and 22 instead of 12. We acknowledge that different subtypes of any of the higher level types of data hunches could be useful in different contexts.

All these types of data hunches can be expressed with different specificity. For example, a value data hunch could apply to a single point (this should be twice as much), to a few data points (all of these items should be twice as high), or to the whole dataset (all data points should be twice as high). These different types of data hunches also establish that hunches may not always be precise, supporting analysts to record some data hunches across varying levels of knowledge. In the next section, we discuss how we design with data hunches in mind. With an emphasis on visual methods, knowledge about sources of qualitative uncertainty can be represented through graphical elements and interpreted context of the data.

5 RECORDING DATA HUNCHES

A key aspect of data hunches is that they are expressed during analysis by a diverse set of stakeholders. Hence, visualizations of the data are the ideal medium to express, record, and consume data hunches. In this section, we explore the set of approaches that can be used for recording data hunches on top of a visualization.

5.1 Data Space

A basic method to record a data hunch is to manipulate the data in data space: before the data has been mapped to a visual element. We consider form-based manipulation and model-based manipulation as the two main methods for recording data hunches in data space. **Form-based manipulation**, shown in Figure 2a, is concerned with inputting a data value or an attribute of the data point through a form, a table, etc., and is suitable for data hunches for specific data items. **Model-based manipulation**, illustrated in Figure 2b, uses a model to bulk-input or edit data values, e.g., through a mathematical function. Note that manipulating the data in data space does not imply that the original data is overwritten.

Previous works have explored ways to express models and values to record knowledge in visualizations. Marasoiu et al. [45], for example, presented an interface that allows users to sketch models, which then generates data points based on the sketch, as a way to facilitate communication between customers and analysts. Romat et al. [60] included data editing in their digital ink externalization system, a functionality requested by participants. Although this functionality was added post facto, it illustrates a preference for editing data directly in systems.

Data space is not well suited to communicate the data hunches that have been recorded in the context of visualizations, as e.g., a tabular

representation of data hunches would be detached from the visualization of the data. Instead, designers will have to consider methods to visualize data hunches provided in data space in visual space.

5.2 Visual Space

Recording data hunches in visual space on top of a visualization provides a direct connection between a data hunch and the visualization. Analysts can think about the data hunch in visual terms as they manipulate it, and consider other data points that are visualized while recording their hunch. Another key benefit of recording data hunches in visual space is that, to a large extent, the same encodings can be used for recording and communicating data hunches. We suggest two techniques for recording data hunches in visual space: free-form sketching and direct manipulation.

Free-form sketching refers to adding visual elements directly to a visualization, using approaches such as pen/mouse-based sketching, or adding elements to a visualization using functionality similar to a drawing program (Figure 2c). The technique provides freedom for analysts to express their data hunch in the way they see fit and can record a variety of data hunches, including structural exclusion through crossing out data points, value/directionality (e.g., a value should be higher in reality) through drawing an arrow, categorical value through shading an area in a color, and value range through drawing an area where a value is expected to be. Visual markups have been used for note taking and communicating thought process in visualizations [39, 45, 60], as well as conversation starters [72]. The process of graphical externalization helps with the understanding of and reasoning about visual information [31] and supports reading and reflecting on the visualization [1, 40, 70]. A downside of markup is that it cannot (easily) be converted into structured data, and hence is only connected to the visualization, but not to the underlying data, making reuse of these hunches in other visualizations of the same dataset exceedingly difficult.

Direct manipulation, illustrated in Figure 2d, involves moving, resizing, removing, adding, or otherwise changing parts of the visualization that encode data. While restricting analysts to the marks and channels of the visualization, direct manipulations offer beneficial affordances: dragging a bar element is easier with a mouse than sketching a new bar, for example. Manipulated marks are also straightforward to translate into data space. Previous works have suggested direct manipulation on visual encodings is a viable way to edit data and provide visual demonstrations of thought processes. Baudel [5] presented editing single or groups of data items in a dataset using graphical manipulations in data visualizations. Saket et al. [61] used graphical manipulation, through repositioning, resizing, and recoloring marks in visualizations, to help users express their expected visualization with increments in direct manipulations, and in turn, the system suggests visual transformations. A drawback of direct manipulation is that each possible manipulation has to be designed and implemented for each chart type, in contrast to free-form sketching, which can be implemented once and re-used for all types of hunches and charts.

5.3 Abstract Space

Assessment data hunches can be expressed only through text, comments, or ratings. For example, an analyst might know that a data source is unreliable, but might not have a concrete idea on what the true data should be. To record such a hunch, they want to add comments

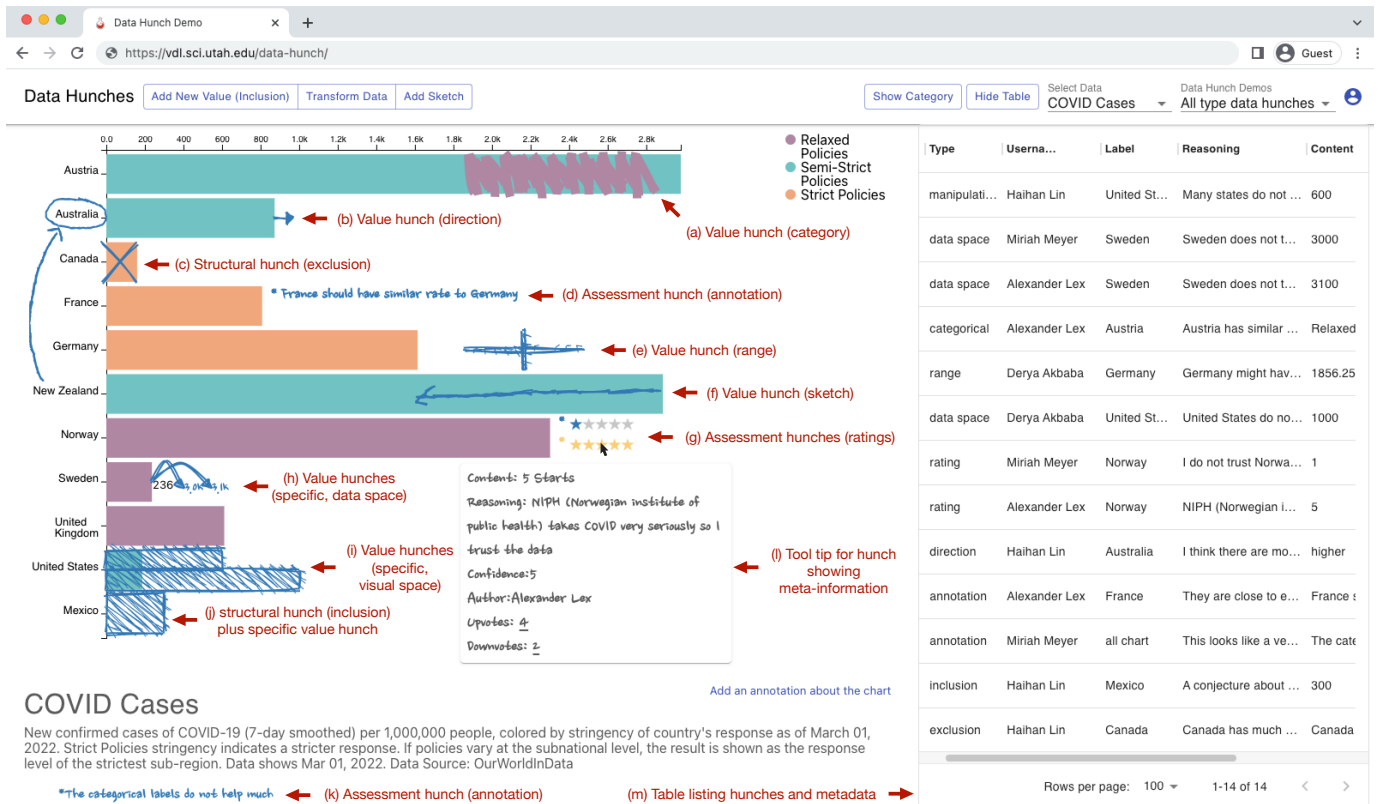


Fig. 3: Our prototype showing a variety of data hunches expressed using tailored methods and designs. Note that red annotations have been added to explain the figure; everything else is as it appears in the prototype. The main chart, showing COVID-19 data and data hunches, is on the left. In the chart, we demonstrate (a) a categorical value hunch that indicates Austria should be colored in purple as it has relaxed policies that are not reflected in the data. (b) A directional value hunch expresses that the data value should be higher, without specifying details. (c) A structural hunch indicates that Canada should not be part of this chart. (d) An assessment hunch is used to comment on the values associated with France. (e) A range value hunch indicates that German values should be higher within a certain range. (f) A value hunch expressed using free-form sketching indicates that values for New Zealand should be lower. (g) Two assessment hunches show diverging perspectives about the quality of Norway’s data. (h) Two value hunches expressed in data space indicate that values in Sweden should be much higher. They are rendered as arrows/labels as they would break the scale. (i) Two specific value hunches show alternative opinions on how high the values for the United States should be. (j) A structural hunch indicates that Mexico should be included in the chart and gives an estimate for a value. (k) An assessment hunch comments on the whole of the dataset. (l) A tool-tip shows information for a selected hunch, including reasoning, confidence, endorsements, and rejections. (m) A (hide-able) table lists all hunches. [View Demo.]

to the data and the data visualization. Such hunches are recorded in “abstract space”, as they do not directly suggest a different structure or value. We identify two methods through which hunches can be recorded in abstract space: **structured elicitation** and **textual annotations** (in addition to e.g., hand-writing using free-form sketching features).

Structured elicitation (Figure 2e) uses form-based UI elements, ratings, and up/down votes, whereas textual annotations (Figure 2f) are free-formed notes. Previous works have explored the use of rating, structured form, and textual annotations in data visualizations. Quispel & Maes [59] used ratings to investigate preferences of visualization types between groups of people. McCurdy et al. [47] used structured forms to elicit data hunches from domain experts, and Goyal et al. [27] offered more freedom to users by allowing them to use a notepad for free-form notes during their experiment. Structured elicitation is different from form-based manipulation in data space: although both can use forms, structured elicitation is about assessment, whereas form-based manipulation is used to express concrete hunches in data space.

6 GUIDELINES FOR DESIGNING FOR DATA HUNCHES

To demonstrate the feasibility of the techniques we proposed in Section 5 and to explore possibilities of communicating and recording data hunches, we developed a prototype, shown in Figure 3, allowing users to record their data hunches on a simple bar chart. As described in Section 8, we used an iterative design process to develop the prototype. We describe insights gleaned from this process for designing visualizations that support recording and displaying data hunches in the form of design guidelines.

Do not change the original data. Techniques to express and communicate data hunches aim to enable analysts to express their knowledge about the data, but not to alter the dataset. Data hunches and original data are different entities and should be clearly separated, to both avoid confusion about the difference between data and data hunch and to retain the integrity of the original data. Data hunches are also valuable only in the context of the recorded data that they refer to. For example, for a value hunch expressed in Figure 3h, we use a sketchy font and arrows to indicate a different value from the original data, and place the value hunch next to the original value. Furthermore, while transparently expressing data hunches, such as doubts or ideas about what a data point should be, are valuable contributions to the analysis process, editing data can be considered deceptive or even fraudulent.

From a technical aspect, designers should treat recorded data hunches as another dataset completely that only references (elements in) the original dataset. In our prototype, for example, data hunches are recorded as a separate dataset, which is shown in a table next to the visualization (Figure 3m). An unfortunate consequence of separating hunches from data is that off-the-shelf visualization systems and libraries are unlikely to make it easy to render data hunches in addition to the underlying data.

Make data hunches distinct. Consistent with our arguments about not changing the original data, data hunches also need to be clearly distinguishable from the original data within the visualizations. Furthermore, the encoding used for data hunches should not only be distinct from the primary encoding, but also should clearly communicate




Fig. 4: Original, non-sketchy design for data hunches. While the data hunches were distinguishable from the original data by color, the distinction was not immediately obvious and could be confused with additional data values that are part of the original data. Hence, we abandoned this design in favor of sketchy renderings.

that what is shown is not the original data.

In our prototype, we use sketchy rendering [74] for visual elements and handwriting-style fonts to make data hunches discernible from the crisp, clean lines of the original visualizations (see Figure 3). The goal of using sketchiness was to make it obvious, even to first time users, that data hunches are not original data but that instead they represent people’s thoughts and knowledge. Wood et al., for example, speculate that “sketchy [...] visualization has a role to play in constructing visualization narratives where an author’s voice is important” [74]. In our first designs, shown in Figure 4, we attempted to render data hunches using a different color, and experimented with gradients and blur to communicate uncertainty and ranges. However, we abandoned this design as we realized that it could lead to confusion, as the marks could be interpreted as belonging to the original data. To avoid this confusion, we developed a rule we applied throughout our design process that all data hunches should look as if they were hand-drawn or written on top of a visualization, to emphasize the humanness of the hunches.

We do not argue that sketchiness is the only, or even the best choice to communicate data hunches. Other designs, tailored to other visualization techniques and affordances, are conceivable. This reasoning applies to all the implemented features we describe in this section.

Another consideration is to be mindful of how disruptive data hunch encodings could be when placed over the original visualizations. The original visualizations must remain visible to properly interpret data hunches. For example, in an early design, we rendered a bar representing a larger data hunch over the original bar. We abandoned this idea in favor of hatched bars that ensure that the underlying data point remains visible.

But, make data hunches similar. Data hunches should use the same or similar visual encodings as the visualizations of the original data. While this guideline seems like a direct contradiction of the guideline on making data hunches distinct, we believe that data hunches and the original data should be read together without the need for mental conversion. For example, a value data hunch could be expressed as a written numerical value on top of a bar chart. However, we argue that this would make it difficult for an analyst to judge the relative differences between the original value and the data hunch. Comparisons are easier if both are expressed using the same visual channel (size/position in the case of a bar chart). Our prototype uses sketchy bars on top of regular bars for numerical specific value hunches (Figure 3i), and hatched color for categorical hunches (Figure 3a), in both cases using the same encoding channel as the original data. However, some hunches cannot be expressed using the same visual encoding. For example, a range value hunch (providing an estimate that a value should be in a certain range) is not compatible with a bar chart encoding used for the original data. To address this, we use a position encoding on the same scale as the bars, showing the middle and the extend of the range (Figure 3e). We use similar techniques for directional value hunches (Figure 3b) and hunches that do not fit on the scale of a chart (Figure 3h).

To reconcile this guideline with the guideline of making hunches distinct from visualizations of the original data, in our demo we relied

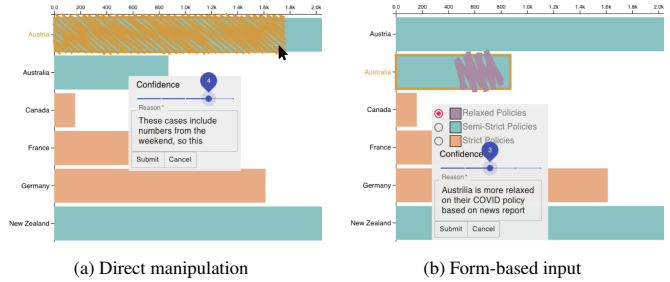


Fig. 5: Recording methods for data hunches, reasoning and confidence. (a) Recording a data hunch using direct manipulation of a bar. Confidence rating and reason are also elicited through a form. (b) Form-based input for a categorical hunch. Selecting an option immediately shows a preview. The input forms are placed close to the selected data point for clear association with the data point.

on using an additional visual channel that was not used in the original visualization: sketchy texture.

Keep data hunches close. As assessment data hunches (comments, ratings, etc.) are not expressed in data space, using the same visual encoding is not feasible. However, designers should ensure that assessment hunches can still be read easily together with the original data. For example, instead of showing assessment hunches in a table, they could be rendered next to the element they are referring to, illustrated for annotations in Figure 3d and for ratings in Figure 3g. If a textual hunch requires more space than is available in the chart, we truncate the comment and reveal the full text in a tooltip. We also considered what to do with assessment hunches referring to the whole dataset and opted for placing a note and an asterisk next to the chart title (Figure 3k); our reasoning is that analysts might read the title and caption as they are attempting to understand the visualization.

Use direct manipulation. Data hunches emerge when analysts explore the data and examine the data visualizations. Hence, the thinking and analysis process happens in visual and data space. While we lay out different approaches for recording data hunches in Section 5, we argue that recording of data hunches should be done as close to the way the data and the data hunch are presented as possible. In our prototype, we trigger recordings by right-clicking on a mark or legend whenever possible and provide methods to record a data hunch through direct manipulation of the data hunch as it will appear once it is recorded (Figure 5a). When recording a hunch in data space, or when recording an assessment hunch, a visualization can provide visual feedback for the analysts (Figure 5b). Also, in our prototype, we place the input forms for data hunches right next to the visual elements in the chart.

Design for data hunches. Not all visualization techniques are equally well suited to visualize data hunches. In our prototype, we have chosen a bar chart with categorical values because bar charts are an important class of visualizations, and because they have affordances that are compatible with data hunches. For example, an analyst could express a value data hunch on top of an individual bar without affecting a neighboring bar. Other visualization techniques, however, do not equally support such similarities. For example resizing a segment in a pie chart, or in a stacked bar chart, requires affecting the other data marks, or overplotting.

Another consideration is the complexity of a chart: the more complex, and the more visual channels are used, the more difficult it will likely be to find a suitable design for data hunches. Similarly, visualizations that give overviews of large amounts of data, like cluster heatmaps, will require different approaches for data hunches, as the manipulation of individual data items is less relevant and the visualizing of the hunches more challenging.

However, even when using a visualization technique that is well suited for data hunches, like bar charts, line charts, or scatterplots (see Figure 2), there are suggestions we recommend that a visualization designer keep in mind to better support data hunches. For example, our original design used vertical bars (see Figure 4). However, we quickly found that vertical bars are problematic for rendering longer comments (assessment hunches) next to the bar due to the text orientation, so we

switched to horizontal bars. Likewise, our original design had a chart title and a subtitle at the top, which made it difficult to find a suitable place for comments on the whole chart. Hence, we moved the title below the chart, and reserved space below the subtitle for comments. Finally, we found it useful to leave white-space from the beginning, so that data hunches can be easily expressed and rendered. For example, our prototype has large margins to the right of the bar chart, so that larger value hunches and comments can be effectively rendered. The designer can also use binning and grouping to organize several data hunches and associate them with a data point, a technique that Badam et al. [3] adopted in FacetNotes. While these specific examples may not directly translate to other visualization techniques, the larger lesson of thinking about position, layout, and space for data hunches as a designer creates a visualization, holds.

Design for collaboration. Data hunches are predominantly a medium to communicate knowledge about data to others, and hence, data hunches are inherently collaborative. Data analysis activities are also commonly collaborative efforts in the first place. We argue that data hunches should be designed with collaboration in mind.

Our prototype acknowledges the importance of collaboration by allowing multiple people to log in and review data hunches. However, we also speculate that enabling multiple collaborators to record only data hunches might be insufficient, as collaborators might also want to endorse, reject, or comment on others’ data hunches. To address this aspect of collaborating, we introduce features to endorse or reject a particular data hunch, using a thumbs-up or thumbs-down metaphor, illustrated in Figure 6. Another more expressive method is to add capability to comment on data hunches. This way, a team member can express their sentiment about a data hunch without having to re-specify a data hunch they endorse.

Elicit context and accountability. *What* a data hunch says about the data is different from *why* a person has the hunch. The context of a data hunch is as critical for its interpretation as the context of the data. Similarly, the reasoning and identity of the data hunch author can effect how a data hunch is perceived and trusted.

We propose that along with recording data hunches, context is important for establishing trustworthiness. This may include reasoning about the data hunch or the identity of person making the hunch. As designers work with stakeholders to determine how data hunches are recorded, they should also explore how important contextual information can be recorded and shared. In our prototype, we require that analysts provide reasoning for and express their confidence in a data hunch when recording it (Figure 5a). Additionally, the identity of the data hunch author is recorded as an attribute of the data hunch. We then visualize these attributes in a tool-tip (Figure 3), although other more salient approaches are conceivable. For example, it might be worth exploring the use of opacity to encode the confidence of a data hunch. These attributes not only enrich the recording of data hunches, but also allow for features such as filtering and sorting.

We acknowledge a tension between revealing identities, to ensure accountability and leverage networks of trust, and the desire to be anonymous to record inconvenient opinions or facts. In prior work, for example, we found that experts in an organization were unwilling to record hunches under their name due to tensions in the organization. In addition to logged-in recording of data hunches, our prototype also provides a “log-in as guest” option to record a hunch without revealing one’s identity. We further discuss this issue in Section 10.

7 PROTOTYPE IMPLEMENTATION

Our web-based prototype (available at <https://vdl.sci.utah.edu/data-hunch/>) implements visualization methods for all types of hunches we describe (see Section 3). For specific value hunches, we also provide dedicated methods for larger numbers of data hunches. The prototype also implements all appropriate methods for recording data hunches (see Section 5) for each type. A specific hunch, for example, can be recorded in data space through form-based manipulation, through model-based manipulation, as well as in visual space through direct manipulation and free-form sketching.




Fig. 6: Collaboration and scalability. Collaboration features, such as endorsing, rejecting, and commenting on data hunches, help in reducing the need for logging similar data hunches in the first place. Different visual encodings account for scalability in value data hunches. If more than three data hunches are logged for a single bar, we replace the sketchy bars with sketchy ticks.

A concern for designing hunches was scalability: how we can show many different data hunches reliably on top of a visualization. We address scalability in two ways: first, we implement dedicated encodings for hunches in case too many hunches are recorded. Figure 6 shows specific value hunches on top of three bars. We use sketchy bars for one or two hunches on top of a bar, but switch to sketchy ticks for more hunches, as shown on the top. However, switching from bars to ticks violates the guideline for making hunches appear similar to the original marks and hence is a trade-off we have to make. Second, we support collaborative features. If a data hunch is already present, another person sharing the hunch could endorse, reject, or comment on data hunches instead of logging a new data hunch, thereby reducing the number of hunches that need to be visualized simultaneously.

We use three example datasets: COVID-19 case counts in selected countries, greenhouse gas emissions across the food supply chain (both downloaded from OurWorldInData), and the size of research areas at the School of Computing at the University of Utah. Data is loaded from CSV files, where data hunches are stored in a Firebase database. Data hunches can be freely added by guests or after signing-in via Google. A tabular visualization (Figure 3) gives an overview of all data hunches and their meta-data, although all relevant information about data hunches is also available through the visualization interface.

Our prototype is open source; the code is available at github.com/visdesignlab/data-hunches-package. We use React and D3 for rendering the UI and the visualizations, and the RoughJS library (roughjs.com), which is based on the techniques for sketchy rendering developed by Wood et al. [74], to render visual elements for data hunches in sketchy patterns.

8 METHODOLOGY

Our methodology for theorizing about data hunches and developing a framework for recording and communicating data hunches was based on reflective practices [21, 62]. We began by reflecting on our experiences working with a variety of domain experts who have rich knowledge about their data, knowledge that was not captured in their datasets. Through group discussions about our experiences, we recognized the missing formalization of personal knowledge and its impact in data analysis. We began mapping out the scope of data hunches, the relationship between data hunches and existing visualization concepts, and how hunches have been reported in the existing literature. This process included a literature search into data feminism, critical data studies, and uncertainty, as well as searching works on design studies and reviewing any reported sources of qualitative uncertainty in previous design studies.

After investigating the landscape of data hunches, we iteratively developed our understanding of data hunches. The iterations critically reflected illustrative examples from our prior experiences, and design spaces proposed for interactive visualization interfaces, uncertainty visualization, and collaborative sense-making. We additionally received feedback from our research lab and colleagues and made adjustments accordingly. We used the design space to re-imagine visualization systems presented in several design studies [30, 43, 47].

Initially, we considered our framework for recording data hunches as a medium for collecting input and knowledge about a dataset from a general audience. Our reasoning was that the crowd might have insights about datasets based on their own experiences, such as collective and local knowledge about a COVID-19 dataset. We came to appreciate that

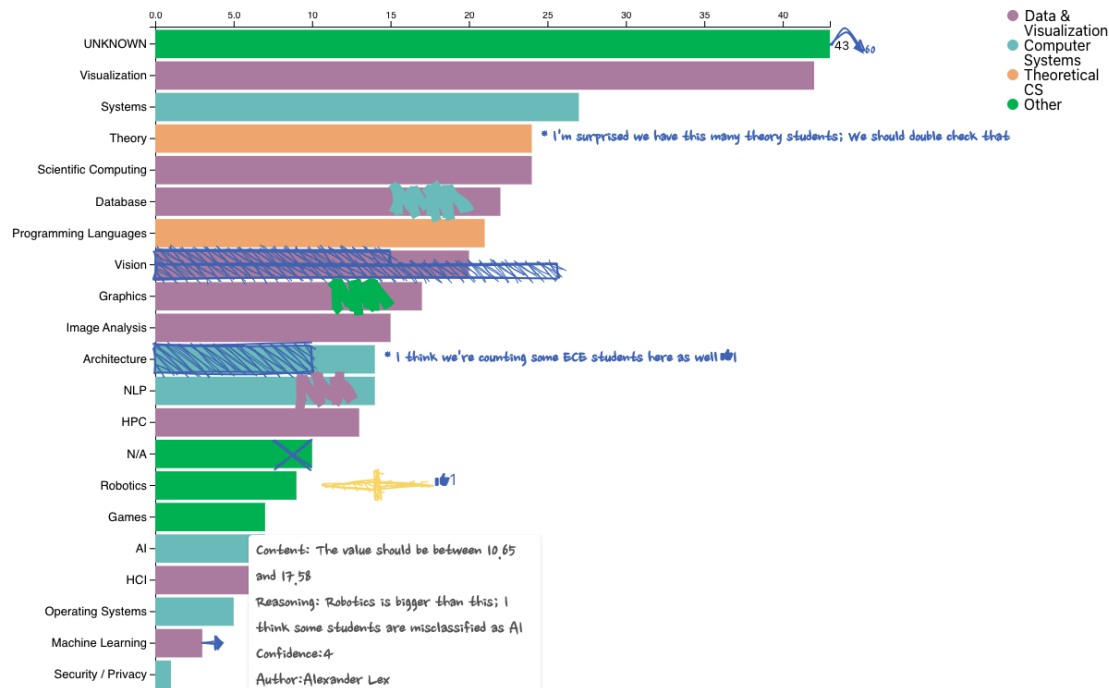


Fig. 7: A usage scenario showing the data hunches of two faculty members about the size of research areas in a CS department. [View Demo]

a key challenge of data hunches, however, is that they could be used to explain away inconvenient data points, or that they could exacerbate the problem of confirmation bias [20, 42]. We thus made the decision to argue for scoping data hunches to collaborative settings with groups of experts who are supported by networks of trust [55]. We discuss the potential benefits and harms that could be associated when data hunches are implemented for general audience systems in Section 10.

Finally, we used an iterative design and development process for our prototype system, with the intent of embodying our ideas and evaluating their feasibility. Over a period of six months, we sketched alternative design ideas and implemented promising solutions, which we then evaluated within our team of four authors and discussed within our research lab. We designed and implemented many variations that we subsequently abandoned for various reasons (see the [supplementary material](#) for examples). We describe the lessons learned from the entire process that stretched over more than a year as guidelines in Section 6.

9 USAGE SCENARIO

Here we illustrate how data hunches could be used in a scenario with real data about the size of research areas in the University of Utah’s School of Computing, shown in Figure 7. To begin, a faculty member, also a co-author of this paper, first noticed that while the largest bar was for *UNKNOWN* students, there was another bar labeled *N/A*. He recorded a hunch that *N/A* should be removed (a structural hunch) and that *UNKNOWN* should be bigger (a value hunch).

Upon reviewing the classification of research areas into larger fields, he expressed concern that *Databases* is classified as *Data and Visualization*, given his knowledge about the type of research conducted by the database groups at the University of Utah, and that it should rather be in the *Computer Systems* category — he recorded his hunch about the different classification of the bar (a value hunch). He expressed a similar data hunch for *Graphics*, which should be in the *Other* bin.

When reviewing the number of *Architecture* researchers, he was surprised by the seemingly high number. He provided a value hunch that he considered closer to reality, and added a comment speculating that some Electrical and Computer Engineering students advised by Computer Science faculty are included in this count. Finally, he realized that *Robotics* is likely shown smaller than it is, probably because some students are incorrectly classified as *AI* instead. He left a value range hunch also noting his reasoning. This faculty member then passed on the visualization to a second faculty member, another co-author of this paper, who reviewed his hunches, upvoted several, and left some additional hunches about her own thoughts on where the data did not

reflect the make-up of the department.

10 DISCUSSION AND FUTURE WORK

Inspired by the breadth of opportunities that data hunches open up, in this section we present a series of discussion threads and future work possibilities. Some of these threads reflect on ethical considerations for ensuring data hunches are used in productive and positive ways, and others consider a number of technical and design challenges for consideration. We end with a brief statement about the opportunity that data hunches point to for considering new perspectives on knowledge.

Challenges for Designing with Data Hunches. As a proof of concept, we used a bar chart to demonstrate how data hunches can be implemented in data visualizations. As we developed the prototype, we realized how quickly the additional layers of data hunches can complicate the visualizations. It was challenging to add data hunches in the chart while keeping the original visualization legible. For data hunches to become established, we need to develop designs for a wide range of visualization types. We believe that our guidelines apply broadly, but the specific design decisions in our prototype might not easily translate to certain types of space-filling visualizations, such as pie charts, treemaps, or icicle plots. Also, while our framework allows for structural hunches through inclusion and exclusion, tabular data is a simple case compared to network data, which has much more complex structural (topological) information. Hence, we believe that a good amount of design work remains to be done to make data hunches work well with such datasets.

Potential for Harm. In our advocacy for data hunches, we focus only on the use cases for analysts with rich knowledge about their field. Narrowing the target audience ensures that the data hunches are based on analysts’ knowledge and experience of the field, and can provide a richer view of the phenomenon that the data represents. The limitation can also avoid misuse and misinformation in data hunches. In an ideal world, our definition of data hunches implicitly assumes that all users are positively contributing to the visualizations when expressing their opinions and knowledge of the data. However, even with good intentions, there is a risk that data hunches could be used to explain away inconvenient data points, or to reinforce an analyst’s preconceived ideas. Because of such risk, it is critical that data hunches are treated fundamentally different from data, even if they are recorded in the same space. A design for data hunches should go to great lengths to avoid any confusion between the data and the data hunch. Such risk also makes it important that data hunches come with explanations

and justifications. Using these techniques, analysts can evaluate a data hunch holistically and judge whether it is reasonable.

Much remains to be explored when considering how to support the public with the ability of recording data hunches on visualizations. Data hunches can encourage open conversations about data and visualizations, similar to current online Q&A forums. Previous research concluded that identity-based trust, social feedback, and exposure all have positive effects on knowledge contribution [28]. Visualizations supporting data hunches are similar to online forums, where users can communicate their knowledge or opinions about the data through a set of techniques and receive feedback (such as upvotes, downvotes, and comments) from other users in the community, and the association with identity, feedback, and exposure can positively promote conversation and knowledge sharing about the data. However, what if a hunch is wrong? Or worse, what if a hunch is maliciously intended as misinformation? Previous research has discovered that data can be used in contradicting ways, depending on how people understand the phenomenon [42]. It is possible that a system that supports recording and visualization of data hunches can be used with malicious intent to fulfill a personal agenda. Moderators, allowing voting, and providing reports mechanism can potentially help with the issue, but it remains an open question that requires further investigations.

Trust and Privacy While we suggest that identities behind data hunches can play a big role in building trust for data hunches, privacy can become an issue and concern for analysts. In some settings, the politics of an organization or field could cause vulnerable people to remain silent about contradictory hunches, depriving others of important perspectives [48]. Anonymity, however, can be equally caustic by invoking negative behavior toward others with opposite views [4]. How to ensure the value, credibility, trustworthiness, and transparency of data hunches is an important, yet open, question.

Another interesting, open question is what happens to someone's trust in a visualization and the underlying data when data hunches are communicated in a tool. Previous work [41] has reported that *social information* can affect a user's trust and memorability about the data visualization. We anticipate similar effects with the inclusion of data hunches. We argue in this paper that data is an imperfect representation of reality, and making that imperfection visible is one goal of our work. However, if data hunches make people less trusting, will designers avoid including them, as they sometimes do with uncertainty [32]?

A reader may trust the visualization more when data hunches are provided by experts, or when data hunches are highly rated. On the other hand, if too many data hunches disagree with the original data, the reader may trust the source of the visualization less. In the end, the goal of conceptualizing data hunches and proposing a design space for them is to formally recognize the role of personal knowledge in understanding data and empower users to express their views. Designers should fully consider the possible impacts of data hunches before committing to including or excluding them. The work we present in this paper is only the first step in exploring a rich space about how, why, and when to include personal knowledge about data in visualizations.

Application Scenarios Another important consideration is what types of visualization systems and scenarios are most appropriate for implementing mechanisms that support data hunches. We believe that most visual data analysis involves hunches, but designing and developing tools that support externalization and communication of multiple data hunches is likely to require significant effort. We see two possible application scenarios that we think justify this effort. On the one hand, data hunches could be integrated in widely used off-the-shelf visualization libraries. For example, adding the capability to visualize and record data hunches to a library such as Altair [66], an interactive visualization library for Python and Jupyter notebooks, could make data hunches accessible to a wide range of audiences. A second possible application is bespoke systems that support recording and communicating data hunches to be implemented primarily in scenarios where the topic of the data is of shared interest among larger communities of experts. For example, a recent project elicited feedback from the scientific community on an animation of the SARS-CoV-2 protein structure [36]. Unlike visualization tools designed for an individual research

lab, such applications target a wider audience with shared interests, where visualizing data hunches can lead to deeper impact compared to casual visualizations. Finally, as designs for data hunches become more common and libraries to add data hunches to visualization become available, it might be feasible to integrate recording and visualization of data hunches into a wider set of visualization tools.

Data Hunches as Structured Data Depending on the type of hunch and the method for recording it, data hunches can also become structured data. Value hunches recorded via data space or direct manipulation, for example, are either directly available in the same space as the original data, or can be easily translated back into data space. They are, hence, different from e.g., annotations provided on top of a figure, as they can be reused whenever a dataset is reused. For example, if a data hunch is recorded for a visualization of a dataset in a Jupyter notebook, the data hunch could be shown not only in that one visualization, but could also be propagated to all subsequent and prior visualizations that are based on that dataset, thereby surfacing the hunch at all stages of the analysis process. Data hunches could also be preserved as datasets are updated, as long as the structure is preserved. For example, if a dataset is refined over time, possibly because of discrepancies recorded as data hunches, a new version of a dataset could be overlaid with data hunches recorded for the old version of the dataset, to see whether the hunches expressed still apply. Overlaying data hunches could be combined with an explicit comparison of datasets [24]. However, such a workflow would incur additional visual complexity and hence would require dedicated methods to manage that complexity.

Epistemology The visualization community has characterized the imperfect and partiality of data as uncertainty; however, this is not the only available classification. Within the fields of critical data studies and digital humanities, data is an artifact of decisions and situated contexts that reflect one captured slice of reality: "*Data are capta*, taken not given, constructed as an interpretation of the phenomenal world, not inherent in it" [18]. Data, then, as an object of decision-making practices, is one representation of many possibilities. This recognition is important because it highlights the non-neutrality of data [11, 16, 17]. These perspectives stem from different epistemological roots. Feminist perspectives on data are based on the theory of *situated knowledges* [29]. This theory posits that knowledge cannot be obtained from a single source, but rather is best derived through a collection and collaboration across partial and overlapping perspectives. Data, similarly, cannot fully represent the natural world. From this theoretical grounding, data and data hunches would capture different perspectives of reality but contribute to a richer more complete picture. This is only a musing, however, leading us to bigger questions such as: In what ways could other epistemologies be productive in visualization research? Could we better characterize and describe the entangled relationship between data, visualizations, and people?

11 CONCLUSION

In this work, we framed the personal knowledge about how representative data is, defining such knowledge as data hunches, and analyzed the implication of supporting data hunches in data analysis. We proposed techniques for recording and communicating data hunches in data visualizations, listed design guidelines, and implemented a prototype to demonstrate data hunches in action. The ultimate goal of this work is to formalize and recognize the significant role that personal knowledge has in understanding data, which many works overlook, and elevate this personal knowledge into another form of information that can be explicitly recorded and utilized. Through this work, we seek to question the notion of data being the gold standard of representing phenomena in the world, and open up the potential to grow visualization research beyond constrained notions of data.

ACKNOWLEDGMENTS

We wish to thank Anders Ynnerman, Ben Shneiderman, Ryan Metcalf, and the Visualization Design Lab for fruitful discussions and feedback. This work was supported by the National Science Foundation (OAC 1835904, IIS 1751238), ARUP Laboratories, and by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

REFERENCES

- [1] G. Aisch, A. Cox, and K. Quealy. You Draw It: How Family Income Predicts Children's College Chances. *The New York Times*, May 2015.
- [2] M. Arnold, R. K. E. Bellamy, M. Hind, S. Houde, S. Mehta, A. Mojsilović, R. Nair, K. N. Ramamurthy, A. Olteanu, D. Piorkowski, D. Reimer, J. Richards, J. Tsay, and K. R. Varshney. FactSheets: Increasing trust in AI services through supplier's declarations of conformity. *IBM Journal of Research and Development*, 63(4/5):6:1–6:13, July 2019. doi: 10.1147/JRD.2019.2942288
- [3] S. K. Badam, S. Chandrasegaran, and N. Elmquist. Integrating annotations into multidimensional visual dashboards. *Information Visualization*, 21(3):270–284, July 2022. doi: 10.1177/14738716221079591
- [4] C. P. Barlett. Anonymously hurting others online: The effect of anonymity on cyberbullying frequency. *Psychology of Popular Media Culture*, 4(2):70–79, Apr. 2015. doi: 10.1037/a0034335
- [5] T. Baudel. From information visualization to direct manipulation: Extending a generic visualization framework for the interactive editing of large datasets. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology*, UIST '06, pp. 67–76. Association for Computing Machinery, 2006. doi: 10.1145/1166253.1166265
- [6] G.-P. Bonneau, H.-C. Hege, C. R. Johnson, M. M. Oliveira, K. Potter, P. Rheingans, and T. Schultz. Overview and State-of-the-Art of Uncertainty Visualization. In C. D. Hansen, M. Chen, C. R. Johnson, A. E. Kaufman, and H. Hagen, eds., *Scientific Visualization*, pp. 3–27. Springer, London, 2014. doi: 10.1007/978-1-4471-6497-5_1
- [7] N. Boukhelifa, A. Bezerianos, T. Isenberg, and J.-D. Fekete. Evaluating Sketchiness as a Visual Variable for the Depiction of Qualitative Uncertainty. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2769–2778, Dec. 2012. doi: 10.1109/TVCG.2012.220
- [8] N. Boukhelifa, M.-E. Perrin, S. Huron, and J. Eagan. How Data Workers Cope with Uncertainty: A Task Characterisation Study. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pp. 3645–3656. Association for Computing Machinery, 2017. doi: 10.1145/3025453.3025738
- [9] K. Brodlić, R. Allendes Osorio, and A. Lopes. A Review of Uncertainty in Data Visualization. In J. Dill, R. Earnshaw, D. Kasik, J. Vince, and P. C. Wong, eds., *Expanding the Frontiers of Visual Analytics and Visualization*, pp. 81–109. Springer London, London, 2012. doi: 10.1007/978-1-4471-2804-5_6
- [10] E. Brynjolfsson and K. McElheran. The Rapid Adoption of Data-Driven Decision-Making. *American Economic Review*, 106(5):133–139, May 2016. doi: 10.1257/aer.p20161016
- [11] M. Correll. Ethical Dimensions of Visualization Research. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, pp. 188:1–188:13. Association for Computing Machinery, 2019. doi: 10.1145/3290605.3300418
- [12] M. Correll and M. Gleicher. Error Bars Considered Harmful: Exploring Alternate Encodings for Mean and Error. *IEEE Trans. Visual. Comput. Graphics*, 20(12):2142–2151, Dec. 2014. doi: 10.1109/TVCG.2014.2346298
- [13] M. Correll, D. Moritz, and J. Heer. Value-Suppressing Uncertainty Palettes. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pp. 642:1–642:11. Association for Computing Machinery, 2018. doi: 10.1145/3173574.3174216
- [14] B. A. Covitt and C. W. Anderson. Untangling Trustworthiness and Uncertainty in Science. *Sci & Educ*, Feb. 2022. doi: 10.1007/s11191-022-00322-6
- [15] Z. T. Cutler, K. Gadhav, and A. Lex. Ttrack: A Library for Provenance Tracking in Web-Based Visualizations. In *IEEE Visualization Conference (VIS)*, pp. 116–120. IEEE, 2020. doi: 10.1109/VIS47514.2020.00030
- [16] C. D'Ignazio and L. F. Klein. Feminist Data Visualization. In *Workshop on Visualization for the Digital Humanities (VIS4DH)*, 2016.
- [17] M. Dörk, P. Feng, C. Collins, and S. Carpendale. Critical InfoVis: Exploring the politics of visualization. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '13, pp. 2189–2198. Association for Computing Machinery, 2013. doi: 10.1145/2468356.2468739
- [18] J. Drucker. Humanities Approaches to Graphical Display. *Digital Humanities Quarterly*, 5(1), 2011.
- [19] M. Feng, C. Deng, E. M. Peck, and L. Harrison. HindSight: Encouraging Exploration through Direct Encoding of Personal Interaction History. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):351–360, Jan. 2017. doi: 10.1109/TVCG.2016.2599058
- [20] F. Fischer. Knowledge politics and post-truth in climate denial: On the social construction of alternative facts. *Critical Policy Studies*, 13(2):133–152, Apr. 2019. doi: 10.1080/19460171.2019.1602067
- [21] R. Fleck and G. Fitzpatrick. Reflecting on reflection: Framing a design landscape. In *Proceedings of the 22nd Conference of the Computer-Human Interaction Special Interest Group of Australia on Computer-Human Interaction*, OZCHI '10, pp. 216–223. Association for Computing Machinery, 2010. doi: 10.1145/1952222.1952269
- [22] N. S. Foundation. NSF's 10 Big Ideas - Special Report. https://www.nsf.gov/news/special_reports/big_ideas/harnessing.jsp.
- [23] M. Franke, R. Barczok, S. Koch, and D. Weltecke. Confidence as First-class Attribute in Digital Humanities Data. In *Proceedings of the 4th VIS4DH Workshop*, 2019.
- [24] K. Gadhav, Z. T. Cutler, and A. Lex. Reusing Interactive Analysis Workflows. *EuroVis (to appear)*, 2022. doi: 10.31219/osf.io/udqjr
- [25] K. Gadhav, J. Görtler, Z. Cutler, C. Nobre, O. Deussen, M. Meyer, J. M. Phillips, and A. Lex. Predicting intent behind selections in scatterplot visualizations. *Information Visualization*, 20(4):207–228, Oct. 2021. doi: 10.1177/14738716211038604
- [26] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. D. Iii, and K. Crawford. Datasheets for Datasets. *Communications of the ACM*, 64(12):86–92, Dec. 2021. doi: 10.1145/3458723
- [27] N. Goyal, G. Leshed, and S. R. Fussell. Effects of Visualization and Note-taking on Sensemaking and Analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pp. 2721–2724. Association for Computing Machinery, 2013. doi: 10.1145/2470654.2481376
- [28] T. Guan, L. Wang, J. Jin, and X. Song. Knowledge contribution behavior in online Q&A communities: An empirical investigation. *Computers in Human Behavior*, 81:137–147, 2018. doi: 10.1016/j.chb.2017.12.023
- [29] D. Haraway. Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies*, 14(3):575, 1988. doi: 10.2307/3178066
- [30] J. Heer, F. B. Viégas, and M. Wattenberg. Voyagers and voyeurs: Supporting asynchronous collaborative visualization. *Commun. ACM*, 52(1):87–97, Jan. 2009. doi: 10.1145/1435417.1435439
- [31] M. Hegarty and K. Steinhoff. Individual differences in use of diagrams as external memory in mechanical reasoning. *Learning and Individual Differences*, 9(1):19–42, Jan. 1997. doi: 10.1016/S1041-6080(97)90018-2
- [32] J. Hullman. Why Authors Don't Visualize Uncertainty. *IEEE Trans. Visual. Comput. Graphics*, 26(1):130–139, Jan. 2020. doi: 10.1109/TVCG.2019.2934287
- [33] J. Hullman, X. Qiao, M. Correll, A. Kale, and M. Kay. In Pursuit of Error: A Survey of Uncertainty Visualization Evaluation. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):903–913, Jan. 2019. doi: 10.1109/TVCG.2018.2864889
- [34] J. Hullman, P. Resnick, and E. Adar. Hypothetical Outcome Plots Outperform Error Bars and Violin Plots for Inferences about Reliability of Variable Ordering. *PLOS ONE*, 10(11):1–25, Nov. 2015. doi: 10.1371/journal.pone.0142444
- [35] P. Isenberg, N. Elmquist, J. Scholtz, D. Cernea, K.-L. Ma, and H. Hagen. Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization*, 10(4):310–326, Oct. 2011. doi: 10.1177/1473871611412817
- [36] J. Iwasa, M. Meyer, A. Lex, J. Rogers, A. H. Liu, and M. Riggi. SARS-CoV-2 Visualization and Annotation Project. <https://animationlab.utah.edu/cova>, 2020.
- [37] C. R. Johnson and A. R. Sanderson. A Next Step: Visualizing Errors and Uncertainty. *IEEE Computer Graphics and Applications*, 23(5):6–10, Sept. 2003. doi: 10.1109/MCG.2003.1231171
- [38] B. Karer, H. Hagen, and D. J. Lehmann. Insight Beyond Numbers: The Impact of Qualitative Factors on Visual Data Analysis. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1011–1021, Feb. 2021. doi: 10/gghxc23
- [39] Y.-S. Kim, N. Henry Riche, B. Lee, M. Brehmer, M. Pahud, K. Hinckley, and J. Hullman. Inking Your Insights: Investigating Digital Externalization Behaviors During Data Analysis. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces*, ISS '19, pp. 255–267. Association for Computing Machinery, 2019. doi: 10.1145/3343055.3359714
- [40] Y.-S. Kim, K. Reinecke, and J. Hullman. 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, CHI '17, pp. 1375–1386. Association for Computing Machinery, 2017. doi: 10.1145/3025453.3025592
- [41] Y.-S. Kim, K. Reinecke, and J. Hullman. Data Through Others' Eyes: The Impact of Visualizing Others' Expectations on Visualization Interpretation. *IEEE Trans. Visual. Comput. Graphics*, 24(1):760–769, Jan. 2018. doi: 10.1109/TVCG.2017.2745240
- [42] C. Lee, T. Yang, G. D. Inchoco, G. M. Jones, and A. Satyanarayan. Viral Visualizations: How Coronavirus Skeptics Use Orthodox Data Practices to Promote Unorthodox Science Online. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, pp. 607:1–607:18. Association for Computing Machinery, 2021.
- [43] H. Lin, R. A. Metcalf, J. Wilburn, and A. Lex. Sanguine: Visual analysis for patient blood management. *Information Visualization*, 20(2-3):123–137, 2021. doi: 10.1177/14738716211028565
- [44] N. Mahyar, A. Sarvghad, and M. Tory. Note-taking in co-located collaborative visual analytics: Analysis of an observational study. *Information Visualization*, 11(3):190–204, July 2012. doi: 10.1177/1473871611433713
- [45] M. Marasoiu, A. F. Blackwell, A. Sarkar, and M. Spott. Clarifying Hypotheses by Sketching Data. In *Proceedings of the Eurographics / IEEE VGTC Conference on Visualization: Short Papers*, EuroVis '16, pp. 125–129. Eurographics Association, 2016. doi: 10.5555/3058878.3058905
- [46] A. Mathisen, T. Horak, C. N. Klokmose, K. Grønbaek, and N. Elmqvist. InsideInsights: Integrating Data-Driven Reporting in Collaborative Visual Analytics. *Computer Graphics Forum*, 38(3):649–661, 2019. doi: 10.1111/cgf.13717
- [47] N. Mccurdy, J. Gerdes, and M. Meyer. A Framework for Externalizing Implicit Error Using Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):925–935, Jan. 2019. doi: 10.1109/TVCG.2018.2864913
- [48] S. Mohammed. Toward an Understanding of Cognitive Consensus in a Group Decision-Making Context. *The Journal of Applied Behavioral Science*, 37(4):408–425, Dec. 2001. doi: 10.1177/0021886301374002
- [49] M. G. Morgan. Use (and abuse) of expert elicitation in support of decision making for public policy. *Proceedings of the National Academy of Sciences*, 111(20):7176–7184, May 2014. doi: 10.1073/pnas.1319946111
- [50] M. Muller, I. Lange, D. Wang, D. Piorkowski, J. Tsay, Q. V. Liao, C. Dugan, and T. Erickson. How Data Science Workers Work with Data: Discovery, Capture, Curation, Design, Creation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, pp. 126:1–126:15. Association for Computing Machinery, 2019. doi: 10.1145/3290605.3300356
- [51] S. Nowak, L. Bartram, and P. Haegeli. Designing for Ambiguity: Visual Analytics in Avalanche Forecasting. In *2020 IEEE Visualization Conference (VIS)*, pp. 81–85. IEEE, 2020. doi: 10.1109/VIS47514.2020.00023
- [52] L. M. K. Padilla, S. H. Creem-Regehr, and W. Thompson. The powerful influence of marks: Visual and knowledge-driven processing in hurricane track displays. *Journal of Experimental Psychology: Applied*, 26(1):1–15, 2020. doi: 10.1037/xap0000245
- [53] L. M. K. Padilla, M. Powell, M. Kay, and J. Hullman. Uncertain About Uncertainty: How Qualitative Expressions of Forecaster Confidence Impact Decision-Making With Uncertainty Visualizations. *Front. Psychol.*, 11:579267, Jan. 2021. doi: 10.3389/fpsyg.2020.579267
- [54] G. Panagiotidou, R. Vandam, J. Poblome, and A. Vande Moere. Implicit Error, Uncertainty and Confidence in Visualization: An Archaeological Case Study. *IEEE Transactions on Visualization and Computer Graphics*, (1):1–1, June 2021. doi: 10.1109/TVCG.2021.3088339
- [55] S. Passi and S. J. Jackson. Trust in Data Science: Collaboration, Translation, and Accountability in Corporate Data Science Projects. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW):136:1–136:28, Nov. 2018. doi: 10.1145/3274405
- [56] K. Potter, R. M. Kirby, D. Xiu, and C. R. Johnson. INTERACTIVE VISUALIZATION OF PROBABILITY AND CUMULATIVE DENSITY FUNCTIONS. *Int J Uncertain Quantif*, 2(4):397–412, 2012. doi: 10.1615/Int.J.UncertaintyQuantification.2012004074
- [57] K. Potter, J. Kniss, R. Riesenfeld, and C. Johnson. Visualizing Summary Statistics and Uncertainty. *Computer Graphics Forum*, 29(3):823–832, Aug. 2010. doi: 10.1111/j.1467-8659.2009.01677.x
- [58] K. Potter, P. Rosen, and C. R. Johnson. From Quantification to Visualization: A Taxonomy of Uncertainty Visualization Approaches. In *Uncertainty Quantification in Scientific Computing*, pp. 226–249. Springer, London, 2012. doi: 10.1007/978-3-642-32677-6.15
- [59] A. Quispel and A. Maes. Would you prefer pie or cupcakes? Preferences for data visualization designs of professionals and laypeople in graphic design. *Journal of Visual Languages & Computing*, 25(2):107–116, Apr. 2014. doi: 10.1016/j.jvlc.2013.11.007
- [60] H. Romat, N. Henry Riche, K. Hinckley, B. Lee, C. Appert, E. Pietriga, and C. Collins. ActiveInk: (Th)Inking with Data. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, pp. 42:1–42:13. Association for Computing Machinery, 2019. doi: 10.1145/3290605.3300272
- [61] B. Saket, H. Kim, E. T. Brown, and A. Endert. Visualization by Demonstration: An Interaction Paradigm for Visual Data Exploration. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):331–340, Jan. 2017. doi: 10.1109/TVCG.2016.2598839
- [62] D. A. Schön. *The Reflective Practitioner: How Professionals Think in Action*. Routledge, Oxfordshire, England, Mar. 2017.
- [63] J. Thomson, E. Hetzler, A. MacEachren, M. Gahegan, and M. Pavel. A typology for visualizing uncertainty. In *Visualization and Data Analysis 2005*, vol. 5669, pp. 146–157. SPIE, 2005. doi: 10.1117/12.587254
- [64] M. Troilo, A. Bouchet, T. L. Urban, and W. A. Sutton. Perception, reality, and the adoption of business analytics: Evidence from North American professional sport organizations. *Omega*, 59:72–83, Mar. 2016. doi: 10.1016/j.omega.2015.05.011
- [65] A. M. van der Bles, S. van der Linden, A. L. J. Freeman, J. Mitchell, A. B. Galvao, L. Zaval, and D. J. Spiegelhalter. Communicating uncertainty about facts, numbers and science. *R. Soc. open sci.*, 6(5):181870, May 2019. doi: 10.1098/rsos.181870
- [66] J. VanderPlas, B. E. Granger, J. Heer, D. Moritz, K. Wongsuphasawat, A. Satyanarayan, E. Lees, I. Timofeev, B. Welsh, and S. Sievert. Altair: Interactive statistical visualizations for Python. *Journal of open source software*, 3(32):1057, 2018.
- [67] F. B. Viegas, M. Wattenberg, F. van Ham, J. Kriss, and M. McKeon. ManyEyes: A Site for Visualization at Internet Scale. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1121–1128, 2007. doi: 10.1109/TVCG.2007.70577
- [68] W. E. Walker, P. Harremoës, J. Rotmans, J. P. van der Sluijs, M. B. A. van Asselt, P. Janssen, and M. P. K. von Krauss. Defining Uncertainty: A Conceptual Basis for Uncertainty Management in Model-Based Decision Support. *Integrated Assessment*, 4(1):5–17, Mar. 2003. doi: 10.1076/iaij.4.1.5.16466
- [69] J. Walny, S. Carpendale, N. H. Riche, G. Venolia, and P. Fawcett. Visual Thinking In Action: Visualizations As Used On Whiteboards. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2508–2517, Dec. 2011. doi: 10.1109/TVCG.2011.251
- [70] J. Walny, S. Huron, C. Perin, T. Wun, R. Pusch, and S. Carpendale. Active Reading of Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):770–780, Jan. 2018. doi: 10.1109/TVCG.2017.2745958
- [71] A. Y. Wang, A. Mittal, C. Brooks, and S. Oney. How Data Scientists Use Computational Notebooks for Real-Time Collaboration. *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW):39:1–39:30, Nov. 2019. doi: 10.1145/3359141
- [72] W. Willett, P. Goffin, and P. Isenberg. Understanding Digital Note-Taking Practice for Visualization. *IEEE Computer Graphics and Applications*, 35(4):38–51, July 2015. doi: 10.1109/MCG.2015.52
- [73] C. M. Wittenbrink, A. T. Pang, and S. K. Lodha. Glyphs for visualizing uncertainty in vector fields. *IEEE Transactions on Visualization and Computer Graphics*, 2(3):266–279, 1996. doi: 10.1109/2945.537309
- [74] J. Wood, P. Isenberg, T. Isenberg, J. Dykes, N. Boukhelifa, and A. Slingsby. Sketchy Rendering for Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2749–2758, Dec. 2012. doi: 10.1109/TVCG.2012.262