

Knowing what to look for: A Fact-Evidence Reasoning Framework for Decoding Communicative Visualization

Sahaj Vaidya and Aritra Dasgupta

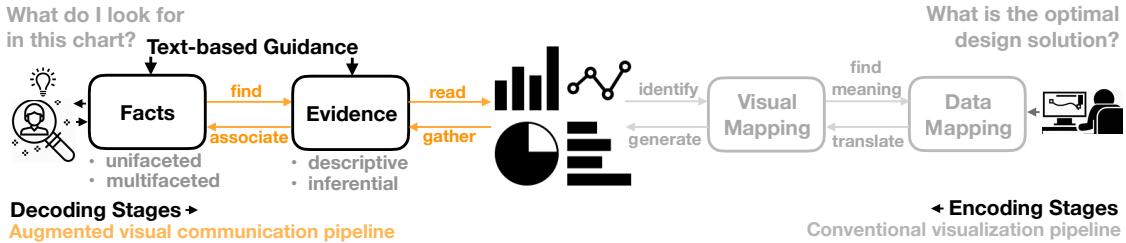


Fig. 1. Our proposed fact-evidence reasoning framework (FaEvR) augments the conventional visualization pipeline by explicitly characterizing the scientific visual communication in terms of decoding facts and associated evidence.

Abstract—Despite the widespread use of communicative charts as a medium for scientific communication, we lack a systematic understanding of how well the charts fulfill the goals of effective visual communication. Existing research mostly focuses on the means, i.e. the encoding principles, and not the end, i.e. the key takeaway of a chart. To address this gap, we start from the first principles and aim to answer the fundamental question: *how can we describe the message of a scientific chart?* We contribute a fact-evidence reasoning framework (FaEvR) by augmenting the conventional visualization pipeline with the stages of gathering and associating evidence for decoding the facts presented in a chart. We apply the resulting classification scheme of fact and evidence on a collection of 500 charts collected from publications in multiple science domains. We demonstrate the practical applications of FaEvR in calibrating task complexity and detecting barriers towards chart interpretability.

Index Terms—Visual communication, scientific communication, graphical reasoning, chart interpretation

1 INTRODUCTION

In his iconic TED talk about fourteen years back [31], Hans Rosling, a Swedish scientist, had opened the door for realizing the hitherto untapped potential of visualization in publicly communicating data-driven facts [33]. Fast-forward to the modern era, when, in the times of a raging global pandemic, charts have been highly consequential for disseminating data-driven facts and evidence. Even beyond these two watershed moments in the history of visual communication, charts have been an integral part of scientific disciplines, via academic publications or presentations. Despite this widespread use and impact of communicative charts in advancing public awareness and scientific discourse, we lack a systematic way of bridging the gap between *what is shown* and *what to look for* in a chart, what we term as a *communication gap*, or the missing link, between the visual representation of data and the mental model of a visualization consumer. This gap is more profound when visualization designers and consumers have different backgrounds or expertise levels (e.g., scientists communicating a message to the public), but also exists when designers and consumers have similar expertise levels (e.g., scientists communicating a message to their peers via publications).

To bridge this gap, we aim to characterize the cognitive effort required for decoding communicative charts using a *fact-evidence* reasoning framework (FaEvR). FaEvR is motivated by the need to characterize the starting point of a visualization decoding pipeline (Figure 1) in terms of the outcomes of scientific data analysis processes. While terms like *messages* or *insights* are used for this purpose, they often have unclear definitions and are loosely connected to the downstream perceptual

tasks and the high-level cognitive processing stages for consuming the information presented in a chart. Our framework addresses this missing link by serving the dual purpose of input and output-oriented reasoning, when decoding and encoding, respectively, from the perspectives of a visualization consumer and a designer. As shown in Figure 1, we augment the conventional, encoding-focused visualization pipeline [13] with the additional input stages of understanding facts and finding evidence, which trigger the process of reading charts and ultimately, lead to the outcome of gathering evidence and associating them back with the facts. The bidirectional arrows in the extended pipeline capture the reasoning processes while decoding and encoding charts. Decoding effort can be calibrated in terms of a visual search process for gathering evidence: a communicative chart could be deemed as most *effective*, when the signals representing relevant evidence are maximized and the noise corresponding to irrelevant evidence can be minimized, a goal that is analogous to optimizing the signal-to-noise ratio in data-driven predictions [35]. When we analyze visualization design using this fact-evidence lens, we can ask questions that can guide us towards calibrating the degree of interpretability of a chart, like: *are the facts represented true or false?*, *are the evidence presented necessary and sufficient for associating them with intended facts?*, and *is the chart expressive* [27] *enough with respect to the presented evidence?*.

For instantiating FaEvR, we collected 500 charts from three science domains, like, energy, climate science, and healthcare. We focus on two main contributions in this paper: i) a theoretical cognitive processing framework organized around the classification of facts and evidence in static, communicative charts, and ii) demonstration of the practical value of the framework in preemptively calibrating task complexity and in analyzing barriers towards effective visual interpretation, by applying the framework on our collected scientific charts.

2 RELATED FRAMEWORKS AND MODELS

In this section, we discuss the frameworks and models related to our fact-evidence based visual communication pipeline.

Critical thinking and reasoning: We know both from real-world expe-

• Sahaj Vaidya is with NJIT. E-mail: ssv47@njit.edu.

• Aritra Dasgupta is with NJIT. E-mail: aritra.dasgupta@njit.edu.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: [xx.xxxx/TVCG.201x.xxxxxx](https://doi.org/10.1109/TVCG.201x.0xxxxxx)

rience and from empirical evidence that critical thinking is an essential component of scientific inquiry and education [7, 22]. Charts representing scientific findings should facilitate such inquiry by engaging visual consumers in deliberative reasoning [26]. In terms of our pipeline, this means reasoning about facts, and assimilation and introspection about associated evidence. In the process of decoding facts and evidence, one can also potentially be engaged in counterfactual thinking [28] and mentally simulate what-if scenarios.

Task taxonomies: Existing visualization task taxonomies and classification schemes [4, 10, 34] mostly address interactive, exploratory data analysis scenarios. Here, the target is often unknown and the goal is to traverse through the *why*, *what*, and *how* [10] pipeline of data and visual mappings for deriving hypotheses and findings. On the other hand, the starting point for decoding scientific charts are the findings, which we characterize as facts and associated evidences. Our framework, therefore, naturally links to the existing task taxonomies by providing a way for visualization consumers, often unfamiliar with the data and domain, to translate their intent to visualization tasks and ultimately decode meaning from the data. The fact-evidence characterization serves as an abstraction layer [36], providing incentives to users to perform the necessary perceptual tasks for gathering evidence and interpreting facts. Analogous to the knowledge precepts proposed by Amar and Stasko for data analysis scenario [5], we provide an abstraction bridge for visual communication. It can be used for a more meaningful analysis of the effects of visualization design on decoding communicative charts, going beyond the exclusive focus on encoding principles [18].

Bridging Perception and Cognition: Visualization can be conceptualized as a communication channel between the data space and the mental space of the intended audience [17]. The encoding side of visualization, based on principles of graphical perception, has been well studied. Starting from the seminal work of Bertin [9], defining the building blocks of data visualization, and the work of Cleveland and McGill [14], providing principles grounded in psychophysics to make informed decisions in visualization design, researchers have developed a wealth of knowledge on how to effectively depict data [8, 11, 23, 25, 37] from the perspective of human perception. Recently, studies and frameworks have also been proposed for analyzing graphical interpretation and cognition [20, 39], connecting visualization interpretation with graphical inference. FaEvR complements these empirical approaches by providing a systematization of the scientific visual communication outcomes and processes. Our work is similar in spirit to the human cognition framework proposed by Patterson et al. [30], however, in comparison, FaEvR provides a more accessible and operational framework to visualization non-experts like domain scientists and general information consumers alike.

3 THE FACT-EVIDENCE REASONING FRAMEWORK (FAEvR)

FaEvR was developed with the intuition that the primary goal of scientific charts, found in publications, presentations, and abundantly in news stories in the COVID-19 era, is to *inform* people about facts. These facts are results of scientific data analysis processes, which naturally generate both facts and associated evidence. We adopt the following definition of fact and evidence: the Cambridge dictionary [1] defines fact as “*something that is known to have happened or to exist, especially something for which proof exists, or about which there is information.*” Evidence [1] is defined as “*anything that helps to prove that something is or is not true.*” Our framework aims to formalize these concepts in the context of visual communication. When presented with a chart, **text-based guidance**, in the form of captions, axis labels, annotations, and chart titles, help users get an idea of the intent of the designer, i.e., the intended fact. Facts need to be backed up by evidence and users need to **find** evidence by **reading** a chart, leading to them performing a set of visualization tasks, followed by **gathering** and **association** of the evidence with the facts. This process of gathering evidence and associating them back with the facts is a hybrid between the classical exploratory data analysis and confirmatory data analysis processes [2]. Here, the complexity of the reasoning tasks is dependent on the nature of facts and evidence. To instantiate FaEvR, we selected three scientific domains where visualization is an integral part of the

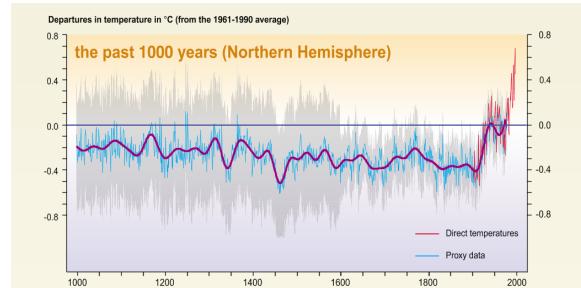


Fig. 2. **Descriptive evidence.** Simulations of the Earth’s temperature variations and comparing the results to the measured temperatures [38]. This is an example of descriptive visual evidence, where, With minimal cognitive effort, a chart consumer can compare the two measures and associate the evidence with the fact about global warming.

scientific communication process, especially to a non-expert audience, like the public and policy-makers. For collecting charts, we consulted with senior domain scientists in each of these domains, like climate science, healthcare, and energy, who are experts working in national labs. They pointed us to the main publications in their respective domains which are used for knowledge dissemination purposes. From these publications, we extracted the charts, along with their captions, and ended with a sample of 500 charts. Next, we performed thematic and qualitative analysis [32] for organizing the facts and evidence based on our classification scheme, as described below.

3.1 Metadata-based Classification of Facts

We use the chart metadata for classifying the facts presented in communicative charts. A dimension contains discrete values such as year, geographical locations which are used to categorize, segment, or group data items. A measure represents a numerically quantifiable piece of information. Temperature, Sales, Profit, Retention Rate, are all examples of specific measures. They represent observations about the data or the calculated values like budget invested for renewable sources of energy, average cost, profit revenue, GDP growth of a country per capita.

We categorize the measures according to their number and their scales by extracting this information from a chart. In the case of charts where there are more than one measure, we further drill down according to the similarity of their scales. By scale here we mean the range of values which suggests the relative size or the extent of a quantitative attribute. In summary, we classify charts based on whether they represent a single measure or multiple measures, and also if, for multiple measures, they are on the same or different scales. Both of these are important factors for triggering the downstream tasks, especially when combining multiple measures for drawing an inference.

Multiple dimensions usually represent different facets of a data item. We adopt the concept of faceted search [24], traditionally used for describing interactive user interfaces, for a high-level classification of facts, expressed in static charts, based on dimensions and measures. We consider **unifaceted facts** as those which depict one dimension corresponding to given measures in a given visualization. **Multifaceted facts** depict more than one dimension or more than one measure with different scales in a visualization. For example, the line chart in Figure 3a, which depicts the rise in temperature involves one dimension (the data from simulations) over a continuous scale of time. Hence we classify this fact as a **unifaceted** one. For the stacked bar chart in Figure 3b, multiple measures and their changes are shown across the time dimension, percentage of yield measure, a continuous measure, and the range of yield change, which is discretized into two directions of change, and hence can be treated as another dimension. Since associating multiple dimensions is needed here, we classify this fact as a **multifaceted** one. Faceted search for evidence gathering is a useful concept for linking the presented facts with the difficulty level of the reasoning and retrieval tasks, as we will see in the following sections.

3.2 Reasoning-based Classification of Evidence

We take inspiration from reasoning frameworks [15], models of graph comprehension [12] and concepts of statistical data analysis [2] for

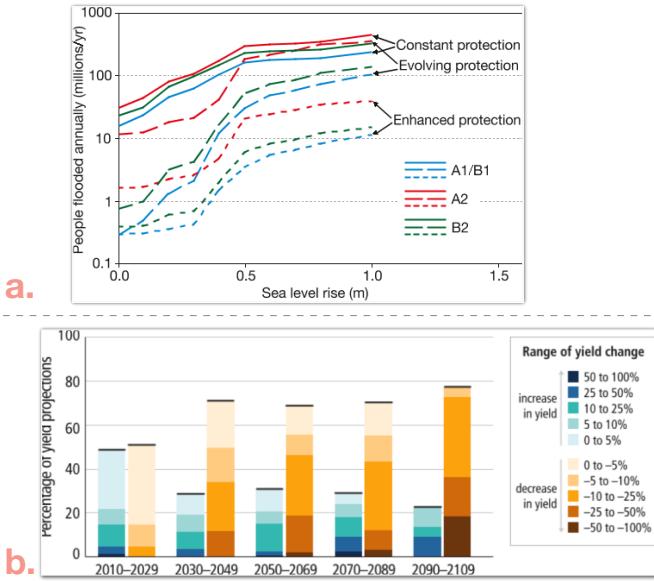


Fig. 3. **Inferential evidence.** Multiple estimates of people flooded in coastal areas due to sea level rise (a) [29] and comparison of projected changes in crop yields across multiple time periods due to climate change (b) [3]. These are examples of inferential visual evidence, as decoding these charts requires mental computation and deductive reasoning for gathering the evidence and associating them with the presented facts.

classifying the different forms of evidence into two broad classes: descriptive evidence, where minimal cognitive effort needs to be spent for decoding a chart and caters to the system 1 type of quick thinking [26] and inferential evidence, where one needs to engage in the process of deliberative reasoning and it caters to the system 2 type of slow thinking [26] for piecing together the evidence. It should be noted that these are not mutually exclusive categories of evidence types as in many cases, both types of evidence can be present in a chart. Descriptive evidence can be considered as the minimal amount of evidence that substantiates a fact, whereas, inferential evidence necessitates an added degree of reasoning for processing the information. To simplify our classification, we avoid creating a hybrid class of evidence and wherever inferential evidence is present, we classify it as such, with the assumption that additional descriptive evidence maybe also present for substantiating the presented fact.

Descriptive Evidence: The concept of descriptive evidence is analogous to that of descriptive statistics for summarizing properties of data. Similarly, when an evidence is shown directly without an additional level of inference required, we classify it as descriptive evidence. An analogy can also be drawn here with WYSIWYG editors where one directly observes what one writes. In the case of charts, the visual evidence directly describes the facts and does not demand a thoughtful cognitive processing on the user’s end to establish the fact. Deriving descriptive evidence is related to recognizing characteristics and relationships which can be identified quickly either through visual means or using text-based guidance. As shown in Figure 2, the chart aims to convey the surge in temperature over Northern Hemisphere for the past 1000 years. Chart consumers can readily notice the upward trend and can associate this with the fact that global temperature has been escalating for decades.

Inferential Evidence: The concept of inferential evidence is similar to that of inferential statistics where one must draw inference based on the metrics (e.g. the degree of uncertainty) presented about the data. Similarly, in a chart, inferential evidence requires the chart consumer to piece together multiple aspects of the information, drawing from the visual encoding of dimensions and measures and using that information to update their mental model about the presented facts [12]. Inferential evidence demands introspective and critical thinking as the user tries to understand: how to derive the evidence, how to link them together, and how they can be associated with the fact that the chart intends to communicate. This requires continuously alternating between the

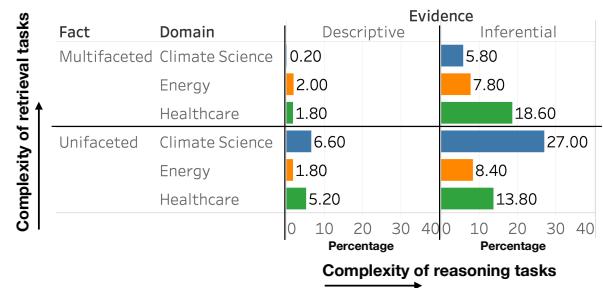


Fig. 4. **Applying the classification of facts and evidence.** A large proportion of the scientific charts we collected exhibited inferential evidence, for which the complexity of reasoning tasks is high.

perceptual properties of the visualization and the underlying semantics [6] of the evidence. Deriving inferential evidence to support the fact requires a greater amount of cognitive effort on the user’s end, because of two reasons. First, the nature of the dimensions could be such that there is an inherent hierarchy, implying that part-to-whole relationships need to be derived either through visual cues or through mental operations. As shown in Figure 3a, the fact represented is about the correlation between people flooded annually and sea-level rise, across different model types (A1, B1, A2, B2). The model types are hierarchical, based on the different protection types. In the line chart, while this is a unifaceted chart with apparently less complexity, one has to spend some time and effort to understand how much variance there is across the protection sub-categories and how that variance can be associated back with the fact. Second, one might need to associate multiple measures at once, which have different scales and semantics, for drawing conclusions. As shown in Figure 3b, the fact is about projected changes in crop yield over given years. Here, the two directions of change need to be associated with the magnitude percentage of yield projections by estimating the relative heights of the different stacks. Combining these two pieces of information, by connecting the end-points of the bar charts and by observing the diverging colors, is needed to verify the trajectory of the temporal changes across multiple years and measures. An increasing number of facets can lead to greater task complexity while deriving inferential evidence, while the presence of prior knowledge about the fact can mitigate some of those costs associated with the decoding process.

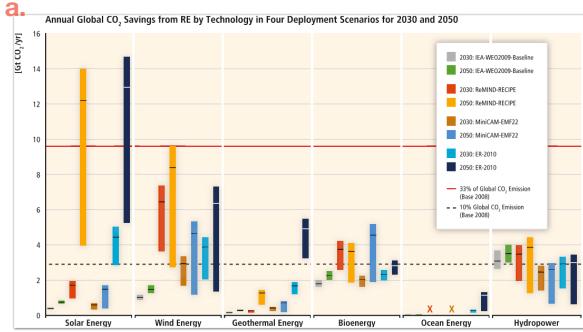
4 APPLICATIONS OF FAEVR

We developed FaEvR with the vision that the classification of scientific findings in terms of facts and evidence will help both chart consumers and chart producers, like domain scientists. The latter can leverage this framework for better design decisions by staying close to their mental model and not spend their time and effort applying encoding-based design principles. This addresses two recurring issues. First, the principles of visual encoding might not fully capture the goals and nuances of what communicative charts need to optimize for. Second, scientists themselves might be skeptical of adopting new techniques or principles owing to the familiarity barrier [16]. FaEvR addresses these by creating an abstraction, using which domain scientists will be able to better anticipate the consequences of their design choices by preemptive task-level assessment of the decoding effort and an evaluation of interpretability barriers in the design outcomes. In this section, we demonstrate how FaEvR can be operationalized by applying the classification scheme on our collection of 500 charts (faevr.njitvis.com) across energy, climate, and healthcare domains.

4.1 Task-level assessment of decoding effort

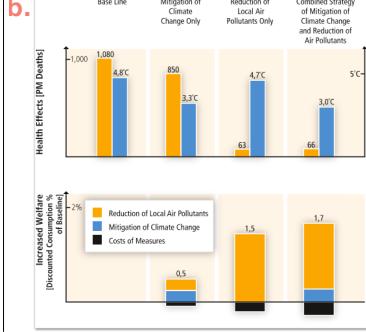
FaEvR is an extension of the conventional data visualization pipeline (Figure 1) and it naturally connects the cognitive reasoning tasks with the lower-level perceptual tasks. Here, we demonstrate that FaEvR helps calibrate the complexity of cognitive and perceptual tasks by using the fact and evidence classifications. To facilitate this analysis, we group the charts based on four quadrants, derived from the fact and evidence types (Figure 4). The figure summarizes two trends we extracted by applying the framework on our collection: i) moving

Presence of visual cues

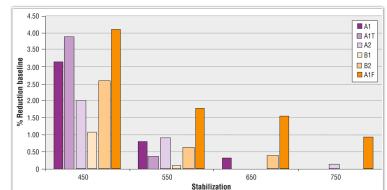


Fact: Multifaceted Evidence: Inferential

Absence of visual cues



C.



Fact: Unifaceted Evidence: Descriptive

Fig. 5. **Illustrating barriers to chart interpretability in simple bar charts.** (a) Variation in expected range of annual global CO₂ savings from renewable energy for multiple scenarios and years [21]. The presence of visual cues, like indicators of mean values or missing values help in reducing the cognitive load for inferential reasoning. In (b), though there are no obvious encoding problems, it requires gathering of evidence across multiple facets, namely the welfare scenarios and the strategies [21], and across measures like temperature rise, health effects, and increased welfare. The absence of appropriate visual cues increases cognitive load for the multiway comparison needed for decoding the chart. In (c), showing global average GDP reduction for alternative stabilization targets and multiple reference scenarios [29], a chart consumer has to mentally organize the presence or absence of data points and the absence of cues can make it difficult to interpret the fact correctly.

from *unifaceted* to *multifaceted* along the Y-axis of fact types, the **complexity of information retrieval tasks** increases and ii) moving from along the X-axis of evidence types from *descriptive* to *inferential*, the **complexity of reasoning tasks** increases. The greater complexity of retrieval tasks is correlated with the larger number of dimensions and measures encoded in multifaceted visualization, implying greater the need for greater decoding time and effort. The greater complexity of reasoning tasks is correlated with the need to spend more time on drawing inferences for extracting semantics from the encodings and make mental calculations and reasoning to assimilate the information. Figure 4 also shows that most charts across the three domains belonged to the inferential evidence category, implying a high degree of average task complexity for the charts we collected. This categorization also helps us to distinguish among the **sequence of perceptual tasks** triggered by the types of facts and evidence. For the 37.97% charts in the *multifaceted - inferential* quadrant (exemplified by Figure 3b), user will generally perform the following sequence of tasks. They will start with the organization of the data values by *comparing* the different dimensions with each of the given measures and *browse* through the values for making necessary derivations. Here, the time involved to understand data distribution is correlated with the number of dimensions. A high amount of cognitive effort is spent in **finding inferential evidence**, by *comparing* the relationships among multiple facets and then **reading** the chart by using the inferences from the previous step. This helps in generating an *explanation* about the semantics of these deductions and ultimately, **gather evidence by associating** the patterns found in the derived *evidence* with the presented *fact*, thus completing the loop. In contrast, for the 11.83% charts in the left most *unifaceted - descriptive* quadrant (exemplified by Figure 3a), since only one dimension is involved, the user mainly aims to *identify* relevant dimensions and patterns related to measures to **find** the *descriptive evidence*. This requires less cognitive effort, as compared to the *multifaceted - inferential* quadrant, and the effort is spent in mainly *summarizing* and *retrieving* the facets to **read** the chart because mapping is to be done only within the one dimension.

4.2 Understanding barriers to chart interpretability

The quadrant-based classification of charts based on FaEvR helps discover several chart interpretability barriers, going beyond the conventional criteria of effectiveness and expressiveness [27] and the usual decoding suspects, like clutter, over-plotting, etc. The challenges of task complexity, encountered during the *mental computation that one needs to perform for inferring the evidence and associating the evidence with the fact*, can be mitigated by use of **visual cues**. Mental computations here involve performing aggregation operations and performing multi-way comparisons [19] across different dimensions and measures. Figure 5a illustrates an example from our collection where visual cues

can provide a reference for facilitating efficient visual comparisons across multiple facets and associated evidence. However in Figure 5b, there is no visual cue available to assist users with finding the associated fact and a chart consumer has to make more inefficient comparisons across the facets of welfare scenarios and strategies and across multiple measures with different scales, for deducing the information. From our chart collection, it was found that around 3% charts in *unifaceted - inferential* and 13% charts in *multifaceted - inferential* guided user's attention towards the relevant details using visual cues. The second factor influencing the interpretation of visualizations is **contextual explanation**. When people try to comprehend a visualization, they are encountered with the task of forming reasonable inferences from the visuals. In these cases, the presence of textual information draws user's attention towards what needs to be seen and they don't have to put extra efforts to understand the graphic as shown in Figure 5a. From the charts we analyzed, there were around 9% *Multifaceted - Descriptive*, 8% *Multifaceted - Inferential* 7% *Unifaceted - Descriptive* and 2% *Unifaceted - Inferential* where having an explanatory text reduces user's efforts to decode the given visualization by providing them with additional context to grasp the information shown. Both the visual features and the cognitive efforts impact the user's ability to get a sense of the message being conveyed. The third barrier to chart interpretability we identified is **missing information** potentially leading to incorrect or incomplete derivation of fact as shown in Figure 5c. When the visualization does not contain information about all the dimensions and measures represented, it obstructs the user from forming a complete understanding of the fact. From the charts we examined, 3% in *Multifaceted - Descriptive* and 5% from *Multifaceted - Inferential* were found to contain insufficient information. This impedes the user from forming a complete sense of the message being conveyed and establishing a link between the fact and its corresponding evidence.

5 CONCLUSION

Fact and evidence-based understanding of information are imperative in a world that is getting inundated with misinformation every day, especially via social media. Visual communication of data-driven facts, demonstrated by the widespread use of charts for creating public awareness about COVID-19, will become an increasingly important tool for ensuring high-quality information dissemination and consumption. By grounding visual communication in a fact-evidence framework, we have laid the foundation for a deeper understanding of how communicative visualization can be optimized for inference-based reasoning.

6 ACKNOWLEDGEMENT

This work was partially funded by the NSF grant 1928627. We would like to thank Gabriel Aquende and Tanvirul Islam for their effort and perseverance with the chart extraction process.

REFERENCES

[1] *Cambridge Dictionary*. Cambridge University Press, 5th ed.

[2] K. Abt. Descriptive data analysis: a concept between confirmatory and exploratory data analysis. *Methods of information in medicine*, 26(02):77–88, 1987.

[3] I. Adopted. Climate change 2014 synthesis report. 2014.

[4] R. Amar, J. Eagan, and J. Stasko. Low-level components of analytic activity in information visualization. In *Proceedings of the IEEE Symposium on Information Visualization*, pp. 111–117. IEEE Computer Society, 2005.

[5] R. A. Amar and J. T. Stasko. Knowledge precepts for design and evaluation of information visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 11(4):432–442, 2005.

[6] K. Andy. Data visualization. *A Handbook for Data Driven Design*. LA London: SAGE, 2016.

[7] S. Bailin. Critical thinking and science education. *Science & Education*, 11(4):361–375, 2002.

[8] R. Beecham, J. Dykes, W. Meulemans, A. Slingsby, C. Turkay, and J. Wood. Map lineups: effects of spatial structure on graphical inference. *IEEE transactions on visualization and computer graphics*, 23(1):391–400, 2016.

[9] J. Bertin. Semiology of graphics; diagrams networks maps. Technical report, 1983.

[10] M. Brehmer and T. Munzner. A multi-level typology of abstract visualization tasks. *IEEE transactions on visualization and computer graphics*, 19(12):2376–2385, 2013.

[11] S. K. Card and J. Mackinlay. The structure of the information visualization design space. In *In Proceedings of Information Visualization and Parallel Rendering Symposium*, pp. 92–99. IEEE, 1997.

[12] P. A. Carpenter and P. Shah. A model of the perceptual and conceptual processes in graph comprehension. *Journal of Experimental Psychology: Applied*, 4(2):75, 1998.

[13] E. H.-h. Chi. A taxonomy of visualization techniques using the data state reference model. In *In Proceedings of IEEE Symposium on Information Visualization*, pp. 69–75. IEEE, 2000.

[14] W. S. Cleveland and R. McGill. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American statistical association*, 79(387):531–554, 1984.

[15] G. Cohen. Inferential reasoning in old age. *Cognition*, 9(1):59–72, 1981.

[16] A. Dasgupta, S. Burrows, K. Han, and P. J. Rasch. Empirical analysis of the subjective impressions and objective measures of domain scientists' visual analytic judgments. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 1193–1204. ACM, 2017.

[17] A. Dasgupta, M. Chen, and R. Kosara. Conceptualizing visual uncertainty in parallel coordinates. In *Computer Graphics Forum*, vol. 31, pp. 1015–1024. Wiley Online Library, 2012.

[18] A. Dasgupta, J. Poco, Y. Wei, R. Cook, E. Bertini, and C. T. Silva. Bridging theory with practice: An exploratory study of visualization use and design for climate model comparison. *IEEE Transactions on Visualization & Computer Graphics*, 21(9):996–1014, 2015.

[19] A. Dasgupta, H. Wang, N. O'Brien, and S. Burrows. Separating the wheat from the chaff: Comparative visual cues for transparent diagnostics of competing models. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1043–1053, 2020.

[20] E. Dimara, S. Franconeri, C. Plaisant, A. Bezerianos, and P. Dragicevic. A task-based taxonomy of cognitive biases for information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 26(2):1413–1432, 2020.

[21] O. Edenhofer, R. Pichs-Madruga, Y. Sokona, K. Seyboth, S. Kadner, T. Zwickel, P. Eickemeier, G. Hansen, S. Schlömer, C. von Stechow, et al. *Renewable energy sources and climate change mitigation: Special report of the intergovernmental panel on climate change*. Cambridge University Press, 2011.

[22] R. Ennis. Critical thinking. 2011.

[23] L. Harrison, F. Yang, S. Franconeri, and R. Chang. Ranking visualizations of correlation using weber's law. *IEEE transactions on visualization and computer graphics*, 20(12):1943–1952, 2014.

[24] M. A. Hearst. Clustering versus faceted categories for information exploration. *Communications of the ACM*, 49(4):59–61, 2006.

[25] J. Heer and M. Bostock. Crowdsourcing graphical perception: using mechanical turk to assess visualization design. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 203–212, 2010.

[26] D. Kahneman. *Thinking, fast and slow*. Macmillan, 2011.

[27] J. Mackinlay. Automating the design of graphical presentations of relational information. *ACM Transactions On Graphics (TOG)*, 5(2):110–141, 1986.

[28] D. T. Miller, W. Turnbull, and C. McFarland. Counterfactual thinking and social perception: Thinking about what might have been. In *Advances in experimental social psychology*, vol. 23, pp. 305–331. Elsevier, 1990.

[29] M. Parry, M. L. Parry, O. Canziani, J. Palutikof, P. Van der Linden, C. Hanson, et al. *Climate change 2007-impacts, adaptation and vulnerability: Working group II contribution to the fourth assessment report of the IPCC*, vol. 4. Cambridge University Press, 2007.

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

[31] H. Rosling. Ted talk: Hans rosling shows the best stats you've ever seen, 2006.

[32] J. Saldaña. *The coding manual for qualitative researchers*. Sage, 2015.

[33] C. O. Schell, M. Reilly, H. Rosling, S. Peterson, and A. Mia Ekström. Socioeconomic determinants of infant mortality: a worldwide study of 152 low-, middle-, and high-income countries. *Scandinavian journal of public health*, 35(3):288–297, 2007.

[34] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *In Proceedings of the IEEE symposium on visual languages*, pp. 336–343. IEEE, 1996.

[35] N. Silver. *The signal and the noise: why so many predictions fail—but some don't*. Penguin, 2012.

[36] B. Victor. Up and down the ladder of abstraction: A systematic approach to interactive visualization. *Kill Math*, pp. 204–228, 2011.

[37] C. Ware. *Information visualization: perception for design*. Morgan Kaufmann, 2019.

[38] R. Watson, D. Albritton, T. Barker, I. Bashmakov, O. Canziani, R. Christ, U. Cubasch, O. Davidson, H. Gitay, D. Griggs, J. Houghton, J. House, Z. Kundzewicz, M. Lal, N. Leary, C. Magadza, J. McCarthy, J. Mitchell, J. R. Moreira, and D. Zhou. *Climate Change 2001: Synthesis Report*. Cambridge University Press Cambridge, UK, 01 2001.

[39] C. Xiong, J. Shapiro, J. Hullman, and S. Franconeri. Illusion of causality in visualized data. *IEEE transactions on visualization and computer graphics*, 26(1):853–862, 2019.