Quantifying Visualization Vibes:
Measuring Socio-Indexicality at Scale

Amy Rae Fox* 🆔, Michelle Morgenstern* 🆔 Graham M. Jones 🆔, Arvind Satyanarayan 🆔
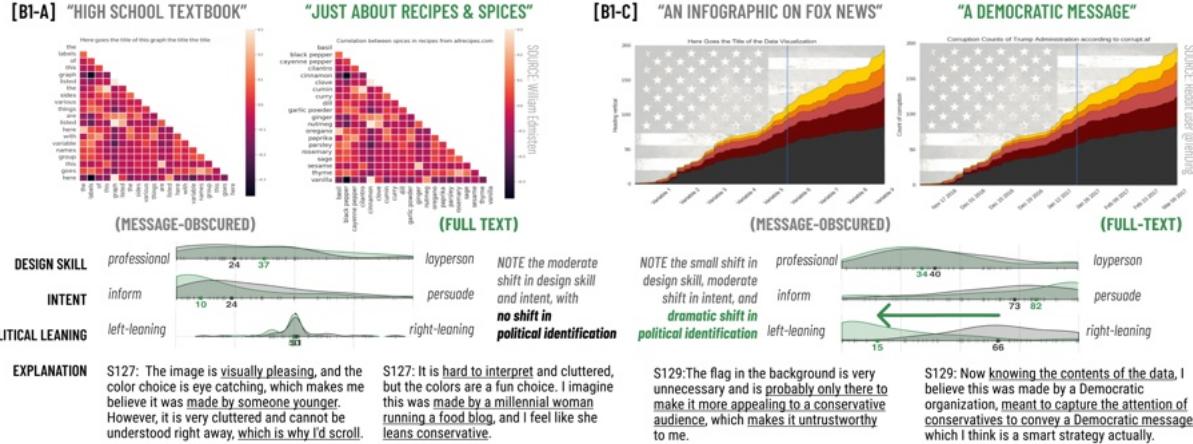
Fig. 1: **Design & Data Context.** Comparing shift in participant to message-obscured and original-text versions of stimuli from the Study 3 social attribution survey illustrates how social inferences arise from the combination of design features and data context.

Abstract—What impressions might readers form with visualizations that go *beyond* the data they encode? In this paper, we build on recent work that demonstrates the *socio-indexical function* of visualization, showing that visualizations communicate more than the data they explicitly encode. Bridging this with prior work examining public discourse about visualizations, we contribute an analytic framework for describing inferences about an artifact's *social provenance*. Via a series of attribution-elicitation surveys, we offer descriptive evidence that these social inferences: (1) can be studied asynchronously, (2) are not unique to a particular sociocultural group or a function of limited data literacy, and (3) may influence assessments of trust. Further, we demonstrate (4) how design features act in concert with the topic and underlying messages of an artifact's data to give rise to such 'beyond-data' readings. We conclude by discussing the design and research implications of inferences about social provenance, and why we believe broadening the scope of research on human factors in visualization to include sociocultural phenomena can yield actionable design recommendations to address urgent challenges in public data communication.

Index Terms—semiotics, socio-indexicality, social provenance, engagement, visualization psychology, public data communication

---◆---

1 INTRODUCTION

Although research has demonstrated that viewers have socially-situated and identity-driven responses to visualizations [26, 58, 64], in offering guidelines for effective design, more attention has been paid to the relatively invariant human factors of graphical perception and graph comprehension [24, 27]. In a recent contribution we argue that visualization not only has a propositional function—conveying insights about depicted data—but *also* a socio-indexical function: conveying impressions of a visualization's social provenance (Morgenstern & Fox et al. [51]). That is, visualizations can generate **socio-indexical inferences: "inferences about a visualization's social provenance based on socioculturally grounded attitudes toward the people, tools, and contexts with which specific visualization features are associated"** [51]. Visualizations convey impressions of the identities

and characteristics of the actors presumed to be involved in its production. These impressions are evoked because the design features of a visualization can point toward—*index*—social categories, contexts, and characteristics [1]. In turn, these socio-indexical inferences can influence how people respond to or engage with a visualization in the first place. The socio-indexical function of visualization is thus critical to understanding how visualizations function as rhetorical objects in the world, and of clear relevance to visualization design.

The socio-indexical function of communication is a concept drawn from linguistic anthropology, which has documented that listeners form impressions about speakers based on the formal features (e.g. accent, lexicon, etc.) of their speech [17, 18, 25]. For example, studies have documented that listeners hearing the same passages spoken with different accents often form different judgements about characteristics of the speakers, such as their intelligence, friendliness, and social status [1, 17, 18, 20, 40, 45]. Importantly, the impressions listeners form of speakers need not be accurate in order to influence behaviour. For example, a classic experiment on language and social cooperation found that compliance with requests for theatre-goers to complete questionnaires differed based on the dialect in which the request was delivered, and the participants' beliefs about the social groups who commonly use it [9]. Such work demonstrates that styles of speech serve as mark-

- * indicates equal contribution. Morgenstern & Jones are affiliated with the MIT Department of Anthropology. Fox & Satyanarayan are affiliated with MIT CSAIL. Email: [mjmorgen, amyfox, gmj, arvindsatya] @mit.edu

[1]This is what distinguishes the function of socio-indexicality (as a semiotic mechanism) from its consequences, such as perceptions of trust.

ers of a speaker's identity, and that listeners make social judgements about speakers' identities in relation to their own, which in turn shapes social interaction. This fundamentally semiotic phenomenon is not limited to natural language, and has been documented in forms of visual communication, including imagery [54], typography [52], and graphic design [53]. Most recently, we offered evidence that data visualizations communicate socio-indexical meaning as well [51]. Specifically, we demonstrated that readers made **social attributions: value-laden inferences about an artifact's social provenance.**

The evidence provided in our initial exploratory work was grounded in ethnographic methods and a purposive sociocultural sample [51]. In this paper, we provide a conceptual replication and extension using survey methods; ultimately demonstrating that inferences about social provenance can be elicited asynchronously. We further integrate our findings with two prior studies examining public discourse [30, 35] to synthesize an analytic framework—a typology—that can be used to construct instruments for eliciting social inferences or conducting content analysis of visualization discourse. We present the results of three attribution-elicitation studies. First we replicate the phenomenon by engaging our prior population (n=78; users of the social media platform Tumblr), followed by a broader sample of US-based English-speakers (n=240; Prolific). After demonstrating that participants in both samples make inferences about social provenance, we test a hypothesis of considerable relevance to visualization design: that inferences about a visualization's makers can intervene in the relationship between its aesthetic-appeal and trustworthiness. Finally, in a third study where participants view message-obscured followed by non-obscured versions of visualizations (n=40; Prolific), we illustrate how readers make social inferences via a combination of a visualization's design features, topic and takeaway messages (see Figure 1). Taken together, we provide converging evidence for the *socio-indexical function* of visualization, demonstrating that visualizations have the capacity to communicate more than the semantic content of the data they encode. Further, we offer quantitative evidence of how such inferences can impact readers' perceptions of trust, and decisions about how a visualization should be engaged. We conclude by describing the implications of readings beyond data for research and design. We discuss how social provenance may be intentionally or unintentionally encoded in a designer's work via the design decisions they make; and call for broadening the scope of human factors research in visualization, from centering readings of data through graphical perception and graph comprehension, to studies of socially-situated visualization behaviour.

2 BACKGROUND AND RELATED WORK

Kennedy & colleagues argue that "the quick extraction of accurate information is only one way of defining what constitutes an effective visualization" [38, pg. 17]. They identify additional criteria, from the extent to which a visualization promotes learning, discourse, or persuasion, to evoking curiosity, surprise or empathy. To develop measures for this broader range of communicative objectives, researchers must extend their attention beyond graphical perception and cognition to a more general class of behaviour described as *engagement*: "the processes of looking, reading, interpreting and thinking that take place when people cast their eyes on data visualisations and try to make sense of them" [38, pg.2]. Here we describe research addressing socially-situated behaviour, including engagement [3, 37, 38, 48], reception [26, 58] framing [7, 41, 42] and most recently, social meaning [51].

2.1 Sociocultural Context & Engagement in Visualization

A growing number of studies are investigating engagement in the context of public data communication, including analysis of online discourse [30, 35], and in-person research with populations difficult to access online [26,38,58]. Analyzing comments made on The Economist magazine's data journalism blog, Hullman & colleagues observed that despite the blog's focus on data, discussions often centered on the broader context of a visualization rather than its content alone [30]. Notably, critiques were frequently directed at the visualization's *creator*, reflecting an engagement with their design decisions and perceived intentions. In similar work, Kauer & colleagues examined discourse

in the visualization subreddit `r/dataisbeautiful` [35]. They developed a framework to categorize different forms of comments, focusing on the scope (e.g. subject) and form (e.g. genre) of public discourse around data. They found that users offered comments of various forms (observations, hypotheses, clarifications, proposals) about various subjects (including the data, topic, and artifact design). In Section 3.5, we extend this framework by including the form *social attribution*, and expanding the range of subjects that can be characterized.

Beyond online discourse, researchers have used interview methods to investigate how social factors shape individual differences in engagement with visualizations. Peck et al. conducted interviews with rural Pennsylvanians, eliciting explanations of how viewers assess the usefulness of different graphs [58]. The most salient theme emerging from their analysis was the extent to which a visualization's usefulness was grounded in assessment of its personal relevance, implicating a complex web of factors informing responses—from a participant's political affiliation and educational background, to beliefs and values grounded in their sociocultural identities. Similarly, He & colleagues explored engagement by interviewing users of public spaces in Canada [26]. They introduced the construct *information receptivity* to describe an individual's differential openness to external representations of information, and interpretations of the information. Importantly, they defined receptivity as a transient state rather than a fixed trait: different stimuli and situations can provoke different reactions from the same individual. We similarly view an individual's response to a graphical artifact as inherently shaped by their sociocultural context. In this work, we zoom in one level of abstraction, aiming to link a participant's reaction to a visualization with their *socio-indexical* response to specific design features. If these features function analogously to indexicality in verbal and visual language, we anticipate they will influence both engagement and interpretive behaviour.

2.2 Framing & Bias: The Role of Text and Data-Topic

A key issue in public data communication is the extent to which visualizations can be designed to minimize bias. Although both critical [2, 38, 63] and semiotic [22, 59, 65] approaches clarify that a representation can never be entirely without bias, it is helpful to understand what features of a visualization trigger well-known cognitive effects like confirmation bias [16], as well as how viewers conceive of bias with respect to data, visualizations, and narrative interpretations. Media and communication scholars have studied the role that titles in particular play in directing readers' attention [6, 44, 66] and framing interpretation [31, 66, 72]. Visualization researchers have found these effects to be true of charts as well [7, 41, 42]. In fact, when the content of a title is slanted or even contradictory to the visualization's depicted data, readers typically recall gist messages aligned with the *title*, rather than the data [41, 42]. Most relevant to the issue of social provenance, Kong & colleagues demonstrate that readers can differentiate between the credibility of a title, a visualization, and its data: the more slanted a title is, the more likely readers are to identify the title as biased, while maintaining that the information is impartial [42]. This echos He et al.'s argument that readers can have different receptions to information (vs) interpretation, where a title acts as an interpretation of the ostensibly impartial encoded data [26]. In our exploratory inquiry we presented participants with a combination of original and message-obscured stimuli to gauge the extent to which inferences about social provenance are derived from a visualization's textual content versus design features [51]. We found that participants could make detailed inferences based on aesthetic features and *also* made inferences about the kind of people who would make charts about particular topics. These indexical associations often unfolded into further inferences about the politics and value alignment between the reader and imagined makers. In the present work, we extend these findings via a repeated-measures study where we first elicit social attributions from a message-obscured image. We then present viewers with the original image and ask which if any of the impressions have changed. In §4.3 we discuss these results and how aesthetics and data-topic interact to give rise to social inferences.

2.3 Social Meaning and Beyond-Data Reading

In linguistic anthropology, *socio-indexical meaning* is understood as "the constellations of inferences that can be drawn about speakers based on how they talk" [5, pg.1]. In a recent study inspired by research in linguistic anthropology on the semiotics of pragmatics [1, 32, 36, 70], we provide evidence for the socio-indexical function of visualization (Morgenstern & Fox et al. [51]). This descriptive work reveals that, like natural language, visualization not only has a referential function, whereby it explicitly encodes propositional meaning (insights about data), but also has a *socio-indexical function* whereby formal features of visualization (i.e. structural and aesthetic design elements, such as chart type, colour, and typography) point to—that is, index—the social contexts, categories, and characteristics with which they are associated. Drawn from C.S. Peirce's triadic distinction between three ways a sign can make meaning [59], indexicality is representation based on real or imagined causal connection or co-occurrence. While visualization design draws heavily on the iconic and symbolic functions of Peircean semiotics, we offer some of the first[2] empirical evidence for the presence and consequences of socio-indexical readings in visualization. Through a series of semi-structured interviews, we document how, in response to open-ended prompts, readers regularly offered comments about a visualization's "vibes"—the social contexts, categories, and characteristics evoked by a visualization's design. From "Microsoft Office vibes" to a "Boomer conservative vibe", readers made social attributions that ranged from concrete identifications of particular actors or types of actors (human and nonhuman) involved in a visualization's production, to elaborate descriptions of their characteristics. We emphasize that both seemingly straightforward *identifications* and clearly subjective *characterizations* can be equally value-laden. Depending on a reader's own social positioning and sociocultural context, simple identifications of age/generation, political party, or even preferred social media platform can imply a host of far-from-neutral qualities. Importantly, we document that part of visualization's socio-indexical function is its impact on behaviour. While our study relied on self-reported impressions, we document how readers described their socio-indexical readings as motivating decisions about how to interact with visualizations, as well as their reception and disposition toward the makers. For example, one participant commented: *"I'd probably just move along [without reading] before it made me angry,"* and another stated, *"I would automatically distrust that person's opinion"* [51].

3 EXPLORING SOCIO-INDEXICALITY AT SCALE

The social attributions reported in our prior work were elicited during semi-structured interviews facilitated by an ethnographer [51]. However, the demand characteristics of interviews are such that participants may be driven to respond in ways they perceive to conform to the interviewer's expectations. Thus, our first aim is to determine whether it is possible to elicit social attributions without a live interviewer. Failure to do so would provide evidence that the prior results may have been compelled, or that the phenomena is so counter to dominant social norms that it cannot be elicited without an explicit permission structure scaffolded by an interviewer. Survey instruments can provide converging evidence for the existence of the phenomenon, insofar as participants respond asynchronously, without the encouragement of an interviewer. Furthermore, respondents have the opportunity to offer indications of confidence, as well as critique or non-response, through free-response text. As a direct-elicitation technique, surveys also have demand characteristics; however, survey respondents experience less social pressure than interviewees (see § 5.1). Moreover, the patterns of variance in numeric responses between participants and across stimuli—especially in concert with free-response explanations—can serve as indications of random or biased responding, or alternatively, of salient features of the phenomenon. In the present studies, we aim to extend the observations of our prior work, determining if inferences about social provenance can be studied asynchronously, and if so, characterize the patterns that emerge from quantifying such inferences. Specifically, we ask:

RQ1 Can social inferences be elicited asynchronously, from a broader demographic sample? If so, what patterns of variance and invariance emerge?

RQ2 Do any types of inference affect the relationship between assessments of a chart's aesthetic *beauty* and *trust*?

RQ3 How does the availability of the main message (e.g. data topic & framing) of a visualization affect social inferences?

To address these questions, we conducted three studies by constructing attribution-elicitation surveys. **Study 1** conceptually replicates [51], eliciting social attributions from users of the social media platform Tumblr using a survey instrument to determine if the prior findings were resultant from the demand characteristics of an interview protocol (n = 78; Tumblr; RQ1, RQ2). **Study 2** extends these findings beyond the original socio-cultural group, using the same survey instrument with a broader demographic sample to determine if the prior findings were idiosyncratic to participants in the sociocultural milieu of Tumblr, a platform known for its high metasemiotic awareness [50] (n = 240; Prolific; RQ1, RQ2). **Study 3** explores if social attributions are affected by a graph's *text* in combination with its design features. We shortened the survey instrument to gather responses to each visualization twice: first, a message-obscured version followed by the original untreated image (n = 40; Prolific; RQ3).

3.1 Study Design

Each study was implemented as a Qualtrics survey and framed to participants as "an exploration of the salience of images on social media." Because our interaction with visualizations is fundamentally contextual, it is necessary to instantiate a specific situational context for engaging with stimuli [21, 29]. We chose the grounding of a social media news feed because social media is a familiar context for encountering visualizations in the wild, and to afford comparison with data from our prior work [51]. We used the framing of visual salience (rather than reading or interpretation) to reinforce the context of social media, and set the expectation that participants would be viewing images and offering impressions, rather than being evaluated on their graph reading skills. In each study, participants viewed a series of (v = 5) visualizations, and answered questions eliciting their impressions of each image.

3.2 Participants

Study 1 leveraged the sampling approach of [51], recruiting via blog post on the social media platform Tumblr, yielding (n = 78) participants (36% Female, 5% Male, 40% Non-binary / third gender, 17% prefer to self-describe, 3% prefer not to say) each compensated via $10.00 gift card. **Study 2** used the same survey with a broader demographic sample, recruiting (n = 240) US-based English-speaking adults via Prolific (54% Female, 42% Male, 3% Non-binary, 1% prefer not to say) each paid $15.00/hour. **Study 3** recruited (n = 40) US-based English-speaking adults via Prolific (50% Female, 47.5% Male, 2.5% Non-binary) paid $15.00/hour.

3.3 Procedure

Fig. 2-A(*left*) illustrates the survey procedure. To establish the situational context, participants selected one of five social media platforms and were instructed to imagine engaging with that platform for the remainder of the survey. They then saw an image of a news feed from the platform, a pattern repeated between every trial to designate the transition to a new stimulus and reinforce the situational context of scrolling through a social media feed. Participants then proceeded to complete a sequence of five trials. In each, the stimulus was anchored to the top of the screen, with survey questions scrolling underneath such that the stimulus was always visible. Following the last trial, participants completed a short demographic questionnaire. In Studies 1 & 2, participants completed five trials (randomly assigned to one stimulus block) and response times ranged from 11 to 227 minutes, $(M = 45, SD = 26)$. In Study 3, participants completed five trials (stimulus Block 1), repeating each trial *twice*: first viewing the message-obscured version of the stimulus, followed by the same set of questions while viewing the original (un-obscured) version. Response times ranged from 13 to 110 minutes, $(M = 41, SD = 20)$.

---

[2]Schofield et.al. [67] discuss how the semiotic index might challenge the sign-world dichotomy implied by the visualization pipeline [12], though our work specifically addresses the *socio*-indexical function [51].
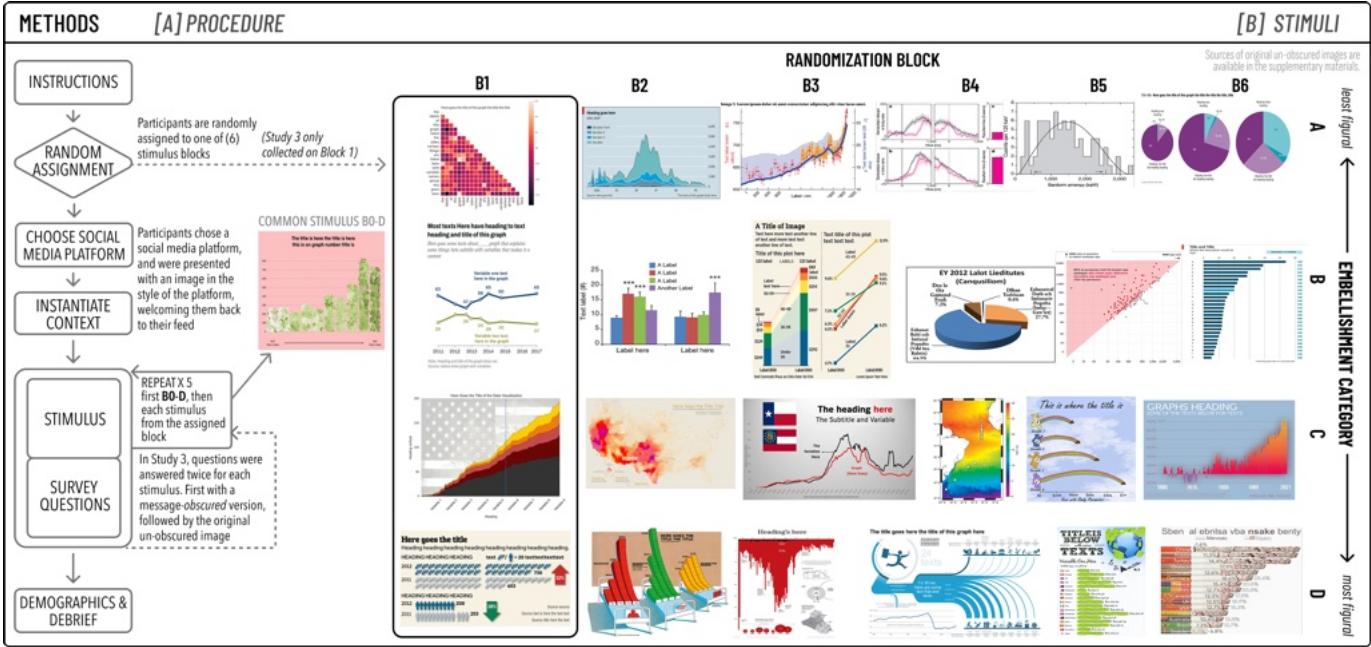
Fig. 2: **Survey Procedure & Stimuli** Part A—*left* illustrates the structure of a participant's survey session. Part B—*right* depicts the full set of stimuli.

3.4 Materials: Visualization Stimuli

To explore the extent to which participants are able to make social attributions *without* the context of data topic, we began by extending the corpus of stimuli used in [51]. We first selected images that elicited particularly detailed indexical readings in the interviews, and then extended the corpus with visualizations from the MassVis dataset [8], optimizing for a variety of chart types and publishing sources. We developed *message-obscured* versions of each image where all text was replaced with placeholders, mimicking typeface and spatial layout. To engage with prior work on aesthetics [3, 4], we organized the images into four groups, positioned along a continuum of increasing aesthetic 'embellishment', from *most abstract* (group A: containing primarily marks as required by the chart type) to *most figural* (category D: containing iconic imagery such as pictograms from which the data topic might reasonably be inferred). We selected one image that elicited consistently detailed indexical responses in [51] and designated that as a common stimulus to be seen by each participant (B0-D, see Figure 2). We organized the remaining stimuli into (6) randomization blocks containing (4) images, one from each embellishment group. In assigning stimuli to blocks we sought to maximize variance in chart type, publication source, and our own subjective assessment of visual style, such that no participant would see four stimuli that were similar along any one of those dimensions. Figure 2-B*(right)* depicts the complete set of stimuli. Each participant in Studies 1 & 2 was randomly assigned to one of the (6) blocks, and thus responded to questions about (5) images: the common stimulus (B0-D) and one image from each embellishment group in randomized order. In Study 3 we sought to determine how the context of data-topic affects social inferences, and elected to begin our exploration of this research question with data collection limited to one stimulus block. We chose Block 1 because three of the four stimuli include elements we believed may cause social inferences to change once the data topic was revealed. Thus each participant in Study 3 was assigned to Block 1, and answered questions about each image twice: first the *message-obscured* version, followed by the original (un-obscured) visualization.

3.5 Materials: Analytic Framework & Survey Questions

To aid in the development of survey questions, we developed an *analytic framework* that can be used to describe the semantic content of discourse about social provenance. Kauer et al.'s contribution analyzing discourse on the `r/ dataisbeautiful` subreddit served as the inspiration for this endeavour, and we began by evaluating the applicability

of their framework, which identified the SCOPE and TYPE of a comment made about a visualization; where SCOPE characterizes the *subject* of a comment, and TYPE characterizes its *form* [35]. We chose to extend the SCOPE portion of this taxonomy to include the social ACTORS identified in our prior work [51]. We reviewed the social attributions expressed in those interviews, and descended a level of detail to identify as many PROPERTIES of named ACTORS as possible. For example, interviewees described the Maker's gender, age, occupation, interests, and political values, among others. We construe these as PROPERTIES of the MAKER. We added three additional ACTORS: (1) DATA, (2) the visualization (ARTIFACT), and (3) AUDIENCE; to better reflect the

ACTOR the **subject**/ topic of the statement	PROPERTIES aspect(s) of the **ACTOR** being described	EXAMPLES (from survey free responses)
MAKER (ambiguous) *responsible for artifact (role in making not specified)*	identification (specific maker), age, gender, trustworthiness, political affiliation, occupation, hobbies, interests, intentions, ... etc	*It has to be a dinosaur kind of company like IBM [B2-B] Some useless bureaucratic office designing propaganda for the government [B1-D]*
MAKER (principal) *responsible for having the artifact made*		*This has Turning Point USA written all over it. (B1-C) Buzzfeed or some less serious news source (B6-D)*
MAKER (author) *responsible for creating the artifact*		*[...] a boomer who created it. It is most likely a graphic designer with some sort of understanding of data analysis. (B1-D)*
MAKER (animator) *responsible for distributing and/or sharing the artifact*		*These types of pictures are usually shared by political, confrontational people. (B0-D)*
MODE *venues, sources, channels where an artifact is distributed*	identification (e.g. specific name), medium (e.g. online, print), circulation, popularity, platform, ... etc.	*a magazine that you could find at a news stand. (B0-D) I imagine this is more of a highschool textbook kind of chart . [B2-D]*
TOOL *instrument(s) used in the artifact's production*	identification (specific tool), purpose, complexity, features, accessibility	*It's not like any R or Python output I've seen. [...] I'm thinking Canva or something similar that non-designers use to pretend to be a designer. (B0-D)*
AUDIENCE *individuals who receive and engage with the artifact*	identification (specific audience), occupation, data literacy, level of education, motivation, political affiliations, ... etc.	*The picture looks childish and given the little bears and bright colors makes me think a parent did this for their child then shared it with the class. (B5-C)*
DATA *entities explicitly encoded in the artifact*	identification (specific dataset), topic, source, structure (e.g. timeseries, geospatial), ...etc.	*Looks like a cell phone coverage map (B2-C) large business's sales team - this feels like a graph revolving around sales (B2-A)*
ARTIFACT *the material (physical or digital) substance of the artifact itself*	identification (i.e. the specific artifact), chart type, marks, spatial layout, complexity, typography background color...etc.	*This looks very monotonous again. There's no creativity, and it looks like it was made strictly for information. [B6-A]*

Fig. 3: **Analytic Framework for Describing Social Provenance.** Inferences about the social provenance of a visualization are composed statements about the PROPERTIES of ACTORS, often expressed as chains of indexical associations, such as ACTOR(PROPERTY) → ACTOR(PROPERTY)

scope of inferences that might be made about a visualization's social origins and purpose. The resulting typology is illustrated in Figure 3 [3] . To construct our survey instrument, we developed questions centered on the PROPERTIES we thought were most likely to yield actionable insights for visualization design; especially the viewer's beliefs about the identity, values, and competencies of the MAKER. The result was a combination of: (1) multiple choice questions (eliciting identifications), (2) semantic-differential scales [28] (eliciting characterizations), and (3) free-response text (eliciting explanations of the previous responses). Together these constituted the 'long-form' survey used in Studies 1 & 2. We developed a 'short-form' version for Study 3 to accommodate the repeated-measures design without increasing participant fatigue. We removed questions about TOOLS, Behaviours, and a subset of questions about the MAKER and ARTIFACT that were strongly correlated (e.g. ARTIFACT(LIKE) was almost precisely correlated with ARTI-FACT(BEAUTY)). In the interest of space we illustrate the short-form survey questions in Figure 5, while the long-form is available in the supplemental materials.

3.6 Data Analysis

Consistent with the descriptive nature of this work, we conducted an exploratory data analysis of quantitative survey measures, aiming to describe patterns of systematic variance and invariance across participants and stimuli. All statistical analyses were conducted in R [61], and quantitative measures were standardized via z-score prior to modelling, though distributions are visualized in this paper using the original response scales for ease of reading. We also performed a thematic analysis of free response data, updating the analytic framework (Fig. 3) with additional ACTOR(PROPERTIES) that were present in respondents' explanations, and extracting representative quotations illustrating salient patterns in the quantitative results. *Survey data, stimuli, model specifications and statistical tests are available alongside a reproducible analysis notebook in the supplemental materials.*

4 RESULTS

In this section, we report the results of our attribution-elicitation studies organized thematically by research question. We begin by describing how the content of free-response text (§ 4.1.1), structure of variance in participant confidence scores (§ 4.1.2), and pattern of response to semantic-differential questions (§ 4.1.3) demonstrate that social attributions can be elicited asynchronously. We describe an exploratory factor analysis indicating a latent structure underlying the elicited social attributions (§ 4.1.4), and statistical models of the relationship between a viewer's assessment of a chart's beauty and the maker's trustworthiness, demonstrating how inferences of the maker's skill in data analysis, intent, and alignment with the viewer's own values moderate this relationship (§ 4.2). We conclude by describing which impressions of a visualization shift when formed with (vs) without the context of a graph's textual content (§ 4.3). *Quotations of free-response text are referenced by the survey ID and stimulus (e.g. S101:B1-C).*

4.1 RQ1:Social Inferences Can be Elicited Asynchronously

4.1.1 Free-Response Explanations

If participants were generally unable to make attributions, or could only do so at the prompting of an interviewer, we would expect to see a high degree of short or nonsense text in the free-response data. In reviewing the explanations gathered in Studies 1 & 2, we found respondents stated specific social attributions, similar in form to those reported in our interview study [51], as well as indications when they were not able to make attributions (i.e. answer the semantic differential questions with certainty.) For example, in response to the stimulus B1-B (a dual-time series with simple colour palette), participant S323 indicated 95% confidence in their identification of the MAKER(TYPE) as a news outlet, explaining, *"This looks like a newspaper like The New York Times."* However, there were also respondents who indicated reluctance to make attributions for certain ACTOR(PROPERTIES). Participant S328

indicated high confidence in their identification of the same stimulus as produced by a news outlet, but indicated 0% confidence in their identification of MAKER(GENDER), writing *"I have no inclination one way or the other about the gender of the writer because that has nothing to do with [...] the source of the image."* This increased our confidence that, while the free-response questions followed the semantic-differential scales and thus primed respondents' attention toward social provenance, participants could nonetheless indicate (via the free-response) if they found the questions about social provenance difficult to answer. Across both studies, free-response explanations frequently described inferences about MAKER(TYPE) (e.g. news outlet, political party, company, etc.) based on a combination of chart type and "editorial style" [3] of the aesthetics. In contrast, participants gave fewer explanations of MAKER(GENDER), unless there was a specific feature (such as the pink background of stimulus B0-D, or "aggressive colors" of B1-C—(see Fig. 5) that strongly index a particular gender stereotype. This is consistent with the claim that an individual may not have socio-indexical associations between design features and every ACTOR(PROPERTY), and further, that stimuli differ in their degree of "indexical-salience" between participants [51]. In Figure 3—*rightmost column* we illustrate the variety of social inferences explained by our participants in the free-response text, using the structure of our analytic framework for social provenance.

4.1.2 Variance in Confidence

While participants can directly signal reluctance to answer questions to an interviewer, survey respondents cannot. In all surveys, participants were asked to indicate their confidence after answering the MAKER(TYPE), MAKER(AGE) AND MAKER(GENDER) multiple choice questions. If participants were generally unable to make attributions, or could only do so at the prompting of an interviewer, we would expect to see low variance in confidence scores between stimuli (with means clustered around 0 or 50%)—indicating survey takers had consistently low confidence—*or* uniform distributions for each stimulus—indicating a random response bias [77]. Figure 4 illustrates the distribution of each confidence question (for Block 1 stimuli) in Studies 1 & 2. Note that although the distribution of confidence ratings aggregated over all stimuli are similar, when broken out by stimulus, we see interesting patterns of variance. These patterns indicate that confidence in each kind of identification *varied in response to different stimuli*—consistent with the claim that social attributions are made in response to the features of a particular visualization.
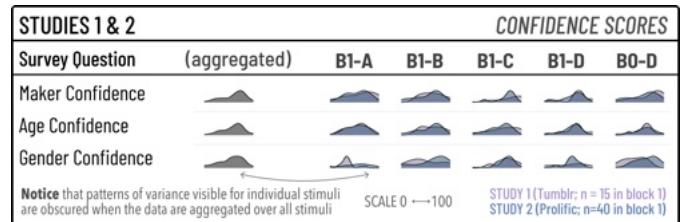
Fig. 4: **Variance in Confidence Scores.** Responses to Studies 1& 2 confidence questions illustrate how confidence varies across the stimuli.

4.1.3 Variance in Identifications & Characterizations

To explore how the features of a visualization's design index social attributions, we must zoom in from our aggregated data to view responses to individual stimuli. Figure 5 illustrates the distribution of responses for the (n = 95) participants assigned to Stimulus Block 1 in all three studies. **These stimuli demonstrate a number of patterns salient in the other stimulus blocks, described in Figure 5-(annotations 1:8).** *Descriptive statistics and visualizations for each stimulus block are available in the supplemental materials.* If participants were generally unable to make attributions we would expect to see either: (1) a pattern of *consistent response* in the semantic differential questions for each stimulus (e.g. uniform distributions for each stimulus, indicating random-answering), (2) stimulus-level clusters at 50% (neutral response bias), or (3) stimulus-level bi-modal distributions (extreme response bias) [77]. Inspection of the stimulus-level

---

[3] See the supplemental materials for a detailed translation between our structure and those offered by Kauer & colleagues [35] and Hullman & colleagues [30].

The questions below appeared on the short form version of our survey. At *right* are the stimuli from Block 1. Descriptive statistics of the semantic differential questions are visualized as density ridge plots **(with medians)**, and multiple choice questions as bar charts.

Are the person(s) most likely responsible for creating this image:

Study 3 responses to **message-obscured** versions of stimuli are in black

Study 3 responses to **untreated** versions of stimuli are in green

B1-A "Spicy Heatmap"
adapted from: William Edelstein

B1-B "Newsy Lines"
adapted from: Pew Research Center

B1-C "Aggressive Flag"
adapted from: SOURCE: Reddit user @TehDing

B1-D "Tiny Icon Guns"
adapted from: The Wall Street Journal

B0-D" Pasted Plants"
adapted from: Mona Chalabi

Maker Design Skill
Are they likely _or_ in graphic design?
professional ⟷ *layperson*

professional

① Images of **complex chart types**, and scaffolding such as **subtitles**, and illustrations were often attributed to makers more **professional** in design

Images that reflected **tool defaults**, or very '**bespoke' elements** more often attributed to **laypersons** in design ①
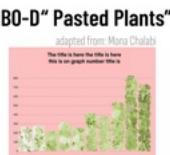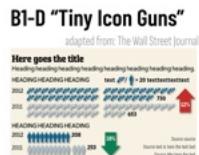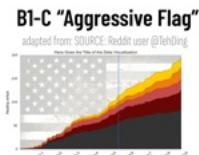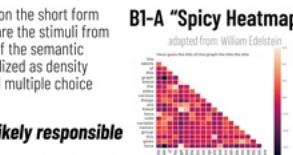
layperson

Data Skill
Are they likely _or_ in data analysis?
professional ⟷ *layperson*

professional

② Attributions and explanations about **data-skill** were often associated with the complexity of the chart type, or complexity of the underlying data or statistical analyses

Simple chart types (especially bar charts) were often associated with makers have *less skill* in data analysis ②

layperson

Politics
Are their politics more likely _or_?
left leaning ⟷ *right leaning*

left-leaning

③ NOTE a distinct difference in the pattern of **politics** ratings between more abstract images (categories A, B) and more embellished images (C,D), where the reader can make reasonable inferences about the topic

Ⓐ Note the dramatic **left-ward shift** in inference about the **maker's political leaning**, when the original text of the image was revealed

right-leaning
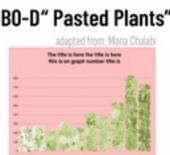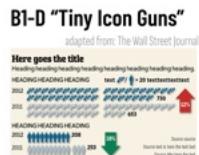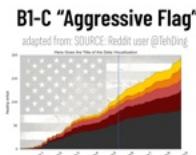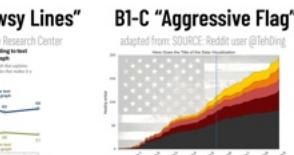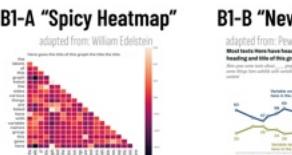
Alignment
Do they likely share your values?
does NOT share ⟷ *DOES share*

does NOT share

④ ratings of **alignment** tend to reflect correspondence with ratings of the maker's political values, in the context of the viewer's own political beliefs

DOES share
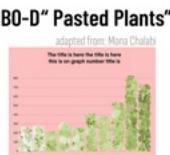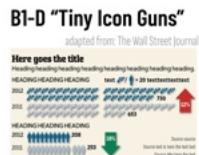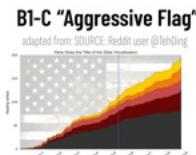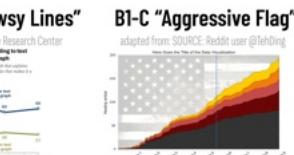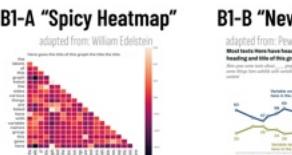
Intention
Is their intent more likely to _or_ ?
inform ⟷ *persuade*

to inform

⑤ ratings of chart **intent** largely corresponded with (categorical) Maker Identity; where political and business makers created images to persuade, where educational and news images were more likely intended to inform

to persuade

Trust
How trustworthy are they?
untrustworthy ⟷ *trustworthy*

untrustworthy

trustworthy

Beauty
How aesthetically pleasing is the image?
not at all ⟷ *very much*

not at all

very pleasing

⑥ inferences about the type or **TYPE/IDENTITY of the MAKER** varied widely between stimuli, largely with explanations that mentioned a combination of the chart type and aesthetics

Maker Identity
Who do you think is most likely responsible for having this image created?

individual
organization
education
business
news
political

Ⓑ **68%** of S3 participants changed Maker ID (many to 'an individual')

Ⓒ **53%** changed Maker ID (many to 'news')

Ⓓ **48%** changed Maker ID (many to news)

Ⓔ **55%** changed Maker ID
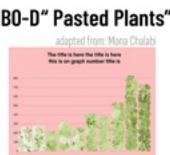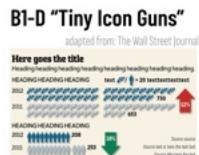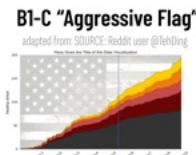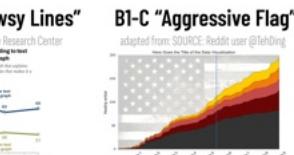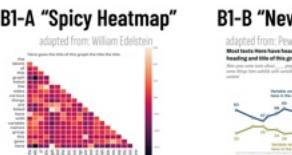
Maker Age
What generation are they most likely from?

boomer
gen-x
millenial
gen-z

⑦ Observe here a stark difference in inference about **gender** of the stimulus B1-C (largely male; described as 'aggressive') and B0-D (largely female; described as 'feminine colors')
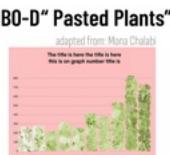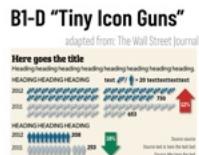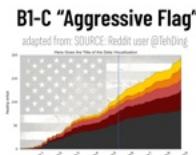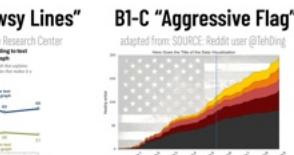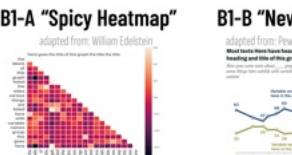
Maker Gender
What gender do they most likely identify with?

other
female
male

Ⓕ **45%** of participants changed Maker Gender, largely from MALE to FEMALE

⑧ Thematic analysis of explanations for whether to **stop and read** (vs) **scroll past** an image indicates individual differences in how different participants share or engage with content *on their social media platform of choice, rather than features of the visualizations.*
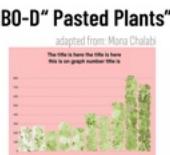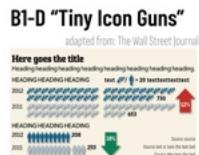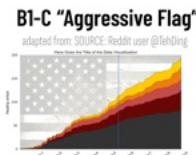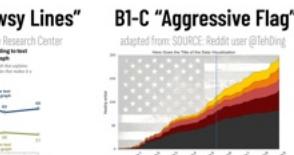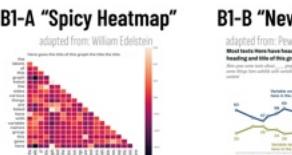
Encounter
As you're scrolling through your feed, you see this image. What would you do?

scroll
read
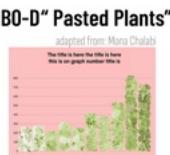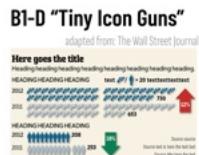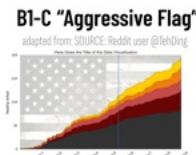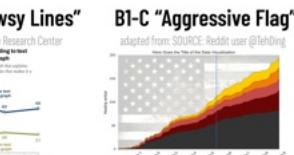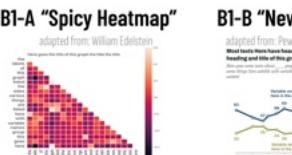
Fig. 5: **Survey Questions Response Distributions for Block 1 Stimuli.** At *left* find the short-form survey questions, and at *right*, visualizations of the distribution of responses for each study (vertically stacked: S1, S2, S3) and each stimulus in Block 1 (as columns).

distributions for each question (Figure 5) suggests this is not the case. To quantitatively verify that the patterns of variance observed do not result from these forms of response bias, we fit a linear mixed effects model on the semantic differential questions, predicting response VALUE by a linear combination of STUDY and interaction between QUESTION and STIMULUS. A model ANOVA indicates both significant main effects, and interaction term between STIMULUS and QUESTION ($F(240) = 15.5, p < 0.001$)—indicating that participants' responses were influenced by both the Stimulus and Question being posed, which we would *not* expect to find if participants responded following a pattern of neutral or extreme response bias [77]. Further, the main effect of STUDY was *not significant* ($F(1) = 2.1, p = 0.15$)—indicating that responses did not systematically vary as a function of the sample population (Study 1: Tumblr, Study 2: Prolific). Taken together, the patterns of variance in confidence and semantic differential questions, along with our inspection of free-response text across both studies, suggest that: (1) participants were in fact offering social attributions rather than nonsense responses, (2) that such attributions are not unique to users of Tumblr, and (3) the patterns of responding illustrate contours of the underlying phenomenon: inferences about the social provenance of a visualization. In the sections that follow, we combine responses from Studies 1 & 2 for exploratory analysis.

4.1.4 A Latent Structure of Social Attributions

To explore the structure of the survey data and identify potential latent constructs underlying the attributions measured via the semantic differential questions, we conducted an exploratory factor analysis[4]. We first evaluated the suitability of the data by calculating the Kaiser-Meyer-Olkin measure ($KMO = 0.83$) (indicating a high level of sampling adequacy), noting Bartlett's test of sphericity was statistically significant ($\chi^2(21) = 2671, p < 0.001$) (supporting the factorability of the correlation matrix). We used the maximum likelihood extraction method to determine the number of factors to retain (verified by a scree plot indicating a clear decrease in power after the third factor), and applied an oblimin rotation to enhance factor interpretability, appropriate for the high correlation between measures. The resulting factor loadings matrix is described in Figure 6. The analysis identified three latent factors accounting for a cumulative 55% of variance in the data:

1. **A Trust/Alignment factor explains (23%) of variance:** with attributions of the makers' political learning, trustworthiness, alignment with the viewer's values, and beauty of the image.

2. **A Design / Beauty factor explains (18%) of variance:** with the maker's skill in graphic design, and beauty of the image.

3. **A Data-Skill/Intent/Trust factor explains (14%) of variance:** with ratings of maker trust, markers' intent (inform ↔ persuade) and skill in data analysis (professional ↔ layperson).

The high proportion of variance explained by the first Trust/Alignment factor indicates that these especially value-laden characterizations of the maker (including their political leaning and extent to which the maker shares the reader's values) are important in relation to trust. We address this relationship further via statistical modelling in (§ 4.2).

4.2 RQ2: Predicting Trust: Inferences About Makers Matter

Recent work in social psychology has extended a well-known effect from research on human faces—more attractive faces are deemed more trustworthy—to the case of visualization, offering crowd-sourced evidence that more beautiful visualizations are rated as more trustworthy [46]. However, participants in our interview study described some graphs as being so well designed they were read as advertisements and thus untrustworthy, and alternatively, relatively 'ugly' graphics as created by scientists and thus more trustworthy [51]. Following this evidence, we hypothesize that the impressions a reader forms

---

[4]Like PCA, EFA is a dimensionality reduction technique; though EFA assumes variance in the observed data comes from unobserved *latent constructs* (unobservable states/traits), and is used in psychometrics to determine how items on a measurement instrument are related to the psychological processes they measure; especially in the case of instruments with highly correlated questions.

Fig. 6: **Exploratory Factor Analysis.** Factor loadings for Studies 1 & 2; observations of (q = 7) semantic differential questions (n = 318)

about the skills, values, and intentions of a visualization's maker likely moderate any relationship between trust and beauty. To explore this hypothesis, we compared a series of linear mixed effects models predicting TRUST by BEAUTY, as well as ALIGNMENT (does *NOT* ↔ DOES share my values), INTENT (inform ↔persuade), and DATA-SKILL (professional ↔layperson), with PARTICIPANT as random intercept. We built models in a step-wise fashion deciding to keep successive predictors via χ^2 and likelihood ratio tests. An initial model predicting TRUST by BEAUTY alone explained ($R^2conditional = 22\%$) of variance in TRUST, with ($R^2marginal = 12\%$) explained by the main effect of BEAUTY. By comparison, the best fitting model (shown in Figure 7) explains ($R^2conditional = 59\%$) of variance in TRUST, with ($R^2marginal = 51\%$) explained by the following fixed effects:

- A main effect of BEAUTY such that more *attractive* graphs are rated as *more* trustworthy.

- A main effect of ALIGNMENT such that graphs from makers who *share the viewer's values* are rated as *more* trustworthy.

- A main effect of INTENT such that graphs intended to *inform* are rated as *more* trustworthy than those intended to *persuade*.

- A main effect of DATA-SKILL such that graphs from makers more *professional* in data analysis are rated as *more* trustworthy.

- An interaction of BEAUTY and INTENT, such that the effect of BEAUTY on TRUST is *minimized* when the maker's intent is to *inform* rather than persuade (*i.e. A graph intended to inform does not need to be as attractive as a graph intended to persuade in order to be assessed as trustworthy*).

- An interaction of ALIGNMENT and INTENT such that the effect of INTENT on TRUST is *minimized* when the maker's values *are aligned* with those of the viewer (*i.e. If I think the maker shares my values, their intent has less effect on my assessment of trust. If they do not share my values, then the trustworthiness of a graph intending to persuade is much lower than that of a graph intending to inform*).

Taken together, these results demonstrate that in addition to a visualization's aesthetic appeal, a viewer's inferences about a maker's skills, values and intentions strongly influence trust. In our best fitting model the variable INTENT contributed to explaining ($IR^2 = 38\%$) of variance in trust, while BEAUTY contributed only ($IR^2 = 13\%$). To better understand how readers draw conclusions about a visualization's trustworthiness, we believe it is necessary to explore the communicative expectations of particular sociocultural groups, exploring differences in what design features index a visualization's communicative intent, and what makes an image 'beautiful' in the eyes of a particular beholder.

4.3 RQ3: Design Features & Data Context

How do the combination of design features with topic & takeaway messages of a visualization affect social inferences? In Studies 1 & 2, participants viewed message-obscured versions of the stimuli in order to explore to what extent viewers make inferences about social
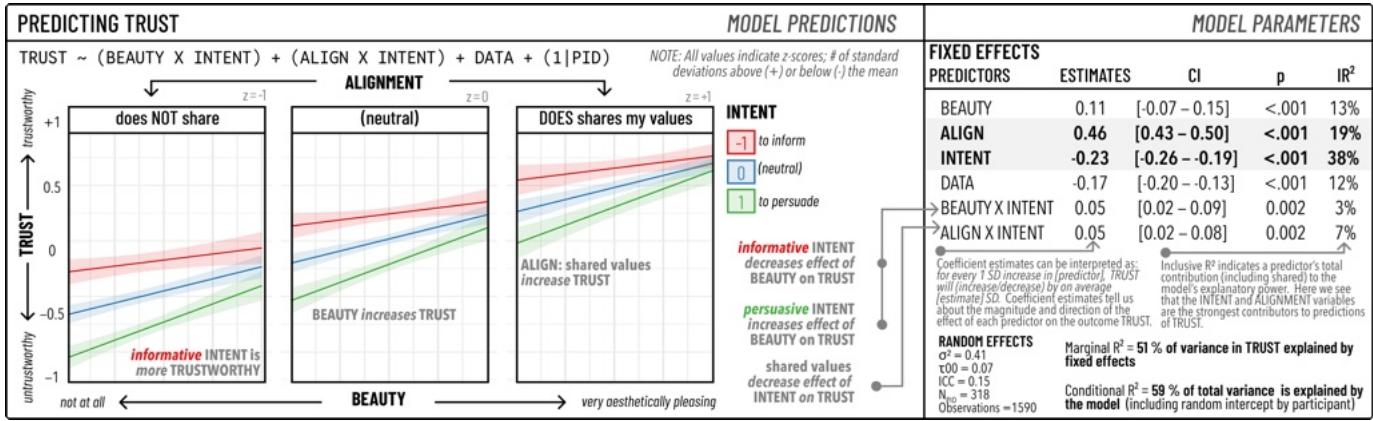
Fig. 7: **Social Inferences Complicate the Relationship Between Beauty & Trust.** This linear mixed effects model explains $R^2(cond) = 51\%$ of variance in TRUST by main effects of BEAUTY, ALIGNMENT, INTENT and DATA-SKILL, & interactions (INTENT and BEAUTY) and (INTENT and ALIGNMENT).

provenance on the basis of the graph's design features alone, without reference to the messages of the underlying data. While complete obfuscation of data-topic is not possible when embellishments include iconic signs, we noted in the free-response data that even when viewers made reference to the presumed topic of the graph, the iconicity of the figural elements alone did not specify the perspective or 'takeaway message'. For example, while the guns in B1-D and hospital beds in B2-D indexed the topics of violence and medical care respectively, they did not indicate if they were communicating 'pro-gun' or 'universal healthcare' messages. This provided the motivation for Study 3, where we used a repeated-measures design to ask participants the same set of questions about each image (in stimulus Block 1) twice: first, when viewing the message-obscured image, followed by the image with original text. On the second viewing participants were also asked to explain how their answers did or did not change.

To address our third research question we began by determining which attributions changed when viewed with the artifact's full text. To examine how *characterizations* (semantic differentials) changed, we evaluated a linear mixed effects model predicting continuous outcome SHIFT (change in answer between text obscured/unobscured) by a linear interaction between QUESTION and STIMULUS, with participant as a random intercept. As expected, an ANOVA indicated significant main effects of QUESTION and STIMULUS and significant interaction. Pairwise comparisons indicate that answers shifted the most for stimuli B1-A (a colorful heatmap), and B1-C (stacked area chart atop a US flag), and across all stimuli, the most labile questions were those about the maker's POLITICS, TRUST, and to what extent the maker's values ALIGN with the viewer's. We also evaluated a mixed effects logistic regression model predicting (binomial) SHIFT in answer to the *identification* type questions (MAKER TYPE, AGE, GENDER) and similarly found significant main effects of both predictors, with post-hoc tests indicating participants most frequently changed their identification of the MAKER TYPE, and that of the stimuli in block 1, participants most frequently changed their identifications for stimulus B1-A. *Significant effects are indicated in Figure 5—(annotations A:F).*

Inspecting the free response data helps us understand how data context influences social inferences. For many readers the initial identification of B1-A (a colourful heatmap) was an educational or scientific maker. However, when the text revealed the *topic* of the graph was Spices, many changed the maker TYPE to *an individual* and the GENDER to *female*, while the INTENT shifted toward *inform* and the graph became more *trustworthy*, despite the individual maker shifting to more *amateur* in DATA analysis. Respondent S47 at first explains, *"I think its definitely intended to inform as there is no wording or descriptions to persuade. I think a business probably created this. Can't tell if man or woman or left vs right without more information."* But upon viewing the full-text image, they write, *"My answers changed when i saw the chart had to do with spices and is from allrecipes. This informs me an individual created [it]. I would guess most likely a woman as they cook more than males and probably search allrecipes more also."*

We also observed a dramatic shift in inferences about the maker of stimulus B1-C (stacked area chart atop a US flag). In all studies the obscured version of this graph was overwhelmingly attributed to a right-leaning maker with persuasive intent, described as: *"intentionally confrontational"*(S359), *"antagonistic"*(S1479), *"pushing some narrative"*(S304). However the un-obscured text reveals the data concern types of corruption in the first administration of US-President Donald Trump (see Fig. 1). With this text, readers' characterizations of the makers' POLITICS shifted dramatically to the *left* with smaller but also significant shifts in alignment of values, but little change across the other attributions. As S85 explains, *"I was definitely wrong about the political leaning [...] but I feel like most of my answers remained the same, based on its aesthetics."* The striking aesthetic decisions made by B1-C's maker and consequently strong (and consistent) inferences made by viewers indicate an effect with substantial implications for visualization design. It is possible, through aesthetic design decisions to *imply* that the identity and values of the maker *align* with those of the target reader, even when that is not the case. If it is true that humans prefer to engage with information that reinforces their existing world views, then representing conflicting data via aesthetics that lead the reader to infer the artifact was made *by* and *for* someone *like them*, there is a pathway to engaging adversarial audiences. This effect is not unique to a particular political stance, as evidenced by this evocative response to an obscured stimulus B1-B: *"That's a New York Times graph if I've ever seen one. It gave me a fight or flight response" (S324).*

5 DISCUSSION

In this paper we contribute a conceptual replication and extension of our recent work arguing that visualizations carry more than the semantic meaning of the data they encode; they also convey socio-indexical meaning (Morgenstern & Fox et al. [51]). We describe three attribution-elicitation studies demonstrating that: (1) social inferences can be elicited asynchronously via surveys; (2) such inferences are not unique to Tumblr [51], and can be elicited from a broader population sample; (3) social inferences affect viewers' assessments of trust; and (4) social inferences are drawn from a combination of a visualization's design features, data-topic, and salient messages, in the context of the viewer's sociocultural identities. Taken together, we provide converging evidence for the *socio-indexical function* of visualization. When encountering visualizations, viewers not only read to extract insights about depicted data, but also make rich and nuanced inferences about an artifact's social origins and purpose in the world. We call for broadening our perspective of how visualizations work—beyond the iconic and symbolic semiotic functions, to include *indexicality*. We argue this phenomenon opens a door to understanding situated behaviour with visualizations, such as how and why an individual might choose to engage or not engage with an artifact, and why artifacts might provoke adversarial readings. To catalyze future research we contribute an analytic framework useful for eliciting or analyzing discourse about social provenance, or as a reference for designers considering what social meaning their design choices might convey.

5.1 Limitations

Although surveys are less subject to the demand characteristics imposed by live social interaction, they are nonetheless subject to social-desirability and other forms of response bias. This can limit a participant's willingness to express ideas that might be perceived as stereotypes. This presents a challenge to any explicit measure of socio-indexical phenomena. Developing associations between formal/design features of communication and the subjective qualities and characteristics of those producing them involves social categorization, and invoking learned stereotypes associated with members of that social group. Acknowledging such stereotypes can be socially fraught. In order to mitigate such effects, we recommend that future work explore *implicit* measures of evaluative responses, drawing on a similar line of research in sociolinguistic cognition [14] in order to determine if social attributions are made spontaneously, or only when prompted. Beyond response bias, the present studies are also limited in their ability to reveal systematic variance in social attributions that may exist *between* sociocultural groups. By using a broad demographic sample our results likely reflect patterns in social attributions that are relatively consistent across populations. However, this sampling approach can mask indexical associations of crucial importance to developing targeted communication strategies for particular audiences. It may be beneficial for basic research to continue with broader samples to explore mechanisms or identify shared visualization ideologies. However, we recommend that work seeking to develop design interventions take a more targeted approach. That is, mapping the indexical fields for the target genre of communication and intended audience first through ethnographic or contextual inquiry. Most importantly, readers should be aware that the specific results of these studies are grounded in the sociocultural identities of participants sampled, the corpus of stimuli, and the situational context of engagement (i.e. a social media feed). If social inferences function in a way that is substantively similar to socio-indexical inferences in (natural) language, then a wealth of further work is required to describe: (1) the nature of visual language varieties, and (2) range of evaluative responses we may hold toward them.

5.2 Extending the Framework for Social Provenance

The analytic framework we contribute is structured to aid thinking about INFERENCES composed of ATTRIBUTIONS about PROPERTIES of ACTORS. It reflects the structure of social provenance of visualizations in which we participate as researchers. However, we note that its *content* (named Actors and Properties) is grounded in the design decisions of our studies, as well as the prior work which it extends [30, 35, 51]. Research employing different populations, stimuli, and situational contexts may yield discourse revealing: (1) additional ACTORS, and (2) emphasizing inferences about different ACTOR(PROPERTIES). For example, as we previously foregrounded interaction on social media [51], those interviews surfaced the MAKER—ANIMATOR as a salient ACTOR. We would not necessarily expect this to be as salient in contexts where artifacts are not regularly shared by ACTORS uninvolved in their creation. Similarly, we expect populations with specialized expertise in visualization design (including technical communication) to yield more PROPERTIES of TOOLS with second-order inferences about the MAKER's competencies based on their tool use (such as programming languages & libraries). We invite scholars studying engagement—especially in applied settings—to actively contribute to extending the framework, giving the community a more powerful tool for thinking about social provenance.

5.3 Implications for Design: Encoding Social Provenance

Intended or not, we argue that every decision a designer makes, from chart type, to colour, framing text and modes of distribution, serve both to communicate desired insights about data, as well as impressions of the artifact's social history. Of course designers are not ignorant to the idea of understanding one's audience; this is central to design education and designerly ways of knowing [15, 56, 57]. Nonetheless it seems that the degree of tuning (especially of aesthetics) to the expectation of one's audience that is prevalent in domains like graphic design, sits in tension with certain modern sensibilities prioritizing minimalism in visualization design [13, 43, 68, 73–75]. We believe this is a tension best resolved by the design community, but can be productively informed by adding a new dimension to empirically derived guidelines for visualization design that actively accounts for the role of socio-indexical inferences in how an artifact is received. An urgent example of the need for this collaboration comes from our exploration of trust. Our quantitative data indicate that social inferences can disrupt an otherwise straightforward hypothesis about beauty and trust: that the more *aesthetically-appealing* a graphic is, the more trustworthy it will be. Our evidence demonstrates this relationship is more complicated. Our data shows that assessments of beauty are moderated by a reader's inferences about both the maker's intent and presumed alignment with their own values. Our survey respondents described particular images as *so* 'well-designed' they believed them to be advertisements, which in turn were not trustworthy, and not interacted with (e.g. "trying to sell me something → keep scrolling"). Alternatively, the most bland graphics using the default settings of basic word-processing programs were described by some as "earnest" and "more objective". These data suggest that designers and science communicators alike have a thin needle to thread with respect to communicating competency and professionalism, without being "so slick" as to trigger attributions of persuasive intent. This challenge is further complicated by the fact that what qualifies as informative versus persuasive differs by individual, as does an individual's openness to data and its interpretation [26].

5.4 Implications for Research: From Graphical Perception & Cognition to Socially-Situated Behaviour

A tremendous volume of research has addressed the psychology of visualization, including: early-stage graphical perception [23, 62, 71], higher-order graph comprehension [21, 60, 69], and subsequent judgement, reasoning and decision-making with extracted information [33, 39, 55]. Empirical research has yielded design guidelines for affording fast and accurate insights *about data* [24, 27, 76]. What existing models of cognition struggle to account for, however, is how viewers' social and situational contexts shape interaction with visualizations [21, 22, 64]. While a growing body of work explores individual differences & sociocultural influences on *reading data* from graphs [3, 34, 39, 49, 78] we join colleagues studying *engagement* more broadly [26, 37, 38, 48, 58] to document distinct forms of behaviour. **Social inferences are *socio-indexical*, rather than *semantico-referential*, graph "readings".** We argue that to understand the broader context of engagement with visualizations outside the laboratory, we must document these socially-situated *beyond-data* behaviours and determine their relationship to graphical perception and cognition. Although the descriptive evidence offered in our prior interviews [51] as well as the present surveys provide converging evidence for the *existence* of socio-indexical inferences as a class of behaviour arising during interaction with visualizations, both methods involve direct elicitation, and are limited in their ability to illuminate the *mechanisms* of this behaviour. What kind of cognitive activities lead to social inferences, and how is this activity related to graphical perception, comprehension, and decision-making? From the present work, we cannot describe the temporal dynamics of social inferences, nor predict if they occur spontaneously, or only upon elicitation. We suggest a logical next step is to join researchers at the intersection of sociolinguistics and cognitive science studying the mechanisms of language attitudes [14]. This would involve adapting implicit measures of social attribution (such as sociolinguistic guise [19, 47] and implicit association tests [10, 11]), and other experimental paradigms. Such measures could reveal: the timescale of social inferences, what stages of cognitive information processing they affect, and the consequences of inferences not explicitly expressed. If social inferences are made automatically, or in the timescale of object-recognition and attention allocation, this has profound implications for what we know about graphical perception and comprehension. If social inferences are more volitional or occur on a longer timescale, they are nonetheless critical to understanding how and why people choose to interact with visualizations in particular ways, and how inferences about makers may bias our decisions and dispositions toward their encoded data.

SUPPLEMENTAL MATERIALS

Supplemental materials are available on OSF at https://doi.org/10.17605/OSF.IO/23HYX.

REFERENCES

[1] A. Agha. Stereotypes and registers of honorific language. *Language in Society*, 27(2):151–193, Apr. 1998. doi: 10.1017/S0047404500019849 1, 3

[2] G. Aiello. Inventorizing, Situating, Transforming: Social Semiotics and Data Visualization. In M. Engebretsen and H. Kennedy, eds., *Data Visualization in Society*, pp. 49–62. Amsterdam University Press, 2020. doi: 10.2307/j.ctvzgb8c7.9 2

[3] M. Alebri, E. Costanza, G. Panagiotidou, and D. P. Brumby. Embellishments Revisited: Perceptions of Embellished Visualisations Through the Viewer's Lens. *IEEE Transactions on Visualization and Computer Graphics*, 30(1), 2024. doi: 10.1109/TVCG.2023.3326914 2, 4, 5, 9

[4] S. Bateman, R. L. Mandryk, C. Gutwin, A. Genest, D. McDine, and C. Brooks. Useful Junk? The Effects of Visual Embellishment on comprehension and Memorability of Charts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 2573–2582. ACM, New York, NY, USA, 2010. doi: 10.1145/1753326.1753716 4

[5] A. Beltrama and F. Schwarz. (im)precise personae: The effect of socio-indexical information on semantic interpretation. *Language in Society*, pp. 1–28, 2024. doi: 10.1017/S0047404524000320 3

[6] J. N. Blom and K. R. Hansen. Click bait: Forward-reference as lure in online news headlines. *Journal of Pragmatics*, 76:87–100, Jan. 2015. doi: 10.1016/j.pragma.2014.11.010 2

[7] M. A. Borkin, Z. Bylinskii, N. W. Kim, C. M. Bainbridge, C. S. Yeh, D. Borkin, H. Pfister, and A. Oliva. Beyond Memorability: Visualization Recognition and Recall. *IEEE Transactions on Visualization and Computer Graphics*, 22(1), 2016. doi: 10.1109/TVCG.2015.2467732 2

[8] M. A. Borkin, A. A. Vo, Z. Bylinskii, P. Isola, S. Sunkavalli, A. Oliva, and H. Pfister. What Makes a Visualization Memorable? *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2306–2315, Dec. 2013. doi: 10.1109/TVCG.2013.234 4

[9] R. Y. Bourhis and H. Giles. The Language of Coperation in Wales: A Field Study. *Language sciences*, 42(13-16), 1976. 1

[10] K. Campbell-Kibler. Sociolinguistics and Perception. *Language and Linguistics Compass*, 4(6):377–389, 2010. doi: 10.1111/j.1749-818X.2010.00201.x 9

[11] K. Campbell-Kibler. The Implicit Association Test and sociolinguistic meaning. *Lingua*, 122(7):753–763, May 2012. doi: 10.1016/j.lingua.2012.01.002 9

[12] S. K. Card, J. D. Mackinlay, and B. Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Academic Press, San Diego, CA, US, 1999. 3

[13] A. Cesal. What Are Data Visualization Style Guidelines? https://medium.com/nightingale/style-guidelines-92ebe166addc. 2021. 9

[14] J.-P. Chevrot, K. Drager, and P. Foulkes. Editors' Introduction and Review: Sociolinguistic Variation and Cognitive Science. *Topics in Cognitive Science*, 10(4):679–695. doi: 10.1111/tops.12384 9

[15] N. Cross. *Designerly Ways of Knowing*. Springer-Verlag. doi: 10.1007/1-84628-301-9 9

[16] 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*, 2018. doi: 10.1109/TVCG.2018.2872577 2

[17] M. Dragojevic. *Language Attitudes*. Oxford University Press. doi: 10.1093/acrefore/9780190228613.013.437 1

[18] M. Dragojevic, F. Fasoli, J. Cramer, and T. Rakić. Toward a Century of Language Attitudes Research: Looking Back and Moving Forward. *Journal of Language and Social Psychology*, 40(1):60–79. doi: 10.1177/0261927X20966714 1

[19] M. Dragojevic and S. Goatley-Soan. The Verbal-Guise Technique. In L. Zipp and R. Kircher, eds., *Research Methods in Language Attitudes*, pp. 203–218. Cambridge University Press, Cambridge, 2022. doi: 10.1017/9781108867788.017 9

[20] P. Eckert. *Language Variation as Social Practice: The Linguistic Construction of Identity in Belten High*. Blackwell, 1999. 1

[21] A. Fox. Theories and Models of Graph Comprehension. In M. Chen, D. A. Szafir, R. Borgo, L. M. K. Padilla, D. J. Edwards, and B. Fisher, eds., *Visualization Psychology*. Springer, 2023. doi: 10.1007/978-3-031-34738-2_2 3, 9

[22] A. Fox and J. D. Hollan. Visualization Psychology: Foundations for an Interdisciplinary Research Programme. In M. Chen, Szafir, Danielle Albers, R. Borgo, L. M. K. Padilla, D. J. Edwards, and B. Fisher, eds., *Visualization Psychology*. Springer, 2023. doi: 10.1007/978-3-031-34738-2_9 2, 9

[23] S. L. Franconeri. Three Perceptual Tools for Seeing and Understanding Visualized Data. *Current Directions in Psychological Science*, 30(5):367–375, Oct. 2021. doi: 10.1177/09637214211009512 9

[24] S. L. Franconeri, L. M. Padilla, P. Shah, J. M. Zacks, and J. Hullman. The science of visual data communication: What works. *Psychological Science in the Public Interest*, 22(3):110–161, 2021. doi: 10.1109/TVCG.2023.3326917 1, 9

[25] L. Hall-Lew, E. Moore, and R. J. Podesva, eds. *Social meaning and linguistic variation: theorizing the third wave*. Cambridge University Press, Cambridge, United Kingdome, 2021. doi: 10.1017/9781108578684 1

[26] H. A. He, J. Walny, S. Thoma, S. Carpendale, and W. Willett. Enthusiastic and Grounded, Avoidant and Cautious: Understanding Public Receptivity to Data and Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 30(1), 2024. doi: 10.1109/TVCG.2023.3326917 1, 2, 9

[27] M. Hegarty. The Cognitive Science of Visual-Spatial Displays: Implications for Design. *Topics in Cognitive Science*, 3(3):446–474, 2011. doi: 10.1111/j.1756-8765.2011.01150.x 1, 9

[28] D. R. Heise. The Semantic Differential and Attitude Research. In *Attitude Measurement*, pp. 235–253. Rand McNally, 1970. 5

[29] J. Hullman. In the real world people have goals and beliefs. In a controlled experiment, you have to endow them, 2023. doi: 10.1109/VIS47514.2020.00049 3

[30] J. Hullman, N. Diakopoulos, E. Momeni, and E. Adar. Content, Context, and Critique: Commenting on a Data Visualization Blog. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '15, pp. 1170–1175. ACM, New York, NY, USA, Feb. 2015. doi: 10.1145/2675133.2675207 2, 5, 9

[31] E. Ifantidou. Newspaper headlines, relevance and emotive effects. *Journal of Pragmatics*, 218:17–30, Dec. 2023. doi: 10.1016/j.pragma.2023.09.013 2

[32] J. Irvine. 'style' as distinctiveness: The culture and ideology of linguistic differentiation. In *Style and Sociolinguistic Variation*. Cambridge University Press, 2004. 3

[33] A. Kale, M. Kay, and J. Hullman. Visual Reasoning Strategies for Effect Size Judgments and Decisions. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):272–282, Feb. 2021. doi: 10.1109/TVCG.2020.3030335 9

[34] A. Karduni, D. Markant, R. Wesslen, and W. Dou. A Bayesian cognition approach for belief updating of correlation judgement through uncertainty visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):978–988, Feb. 2021. doi: 10.1109/TVCG.2020.3029412 9

[35] T. Kauer, M. Dörk, A. L. Ridley, and B. Bach. The Public Life of Data: InvestigatingReactions to Visualizations on Reddit. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, pp. 1–12. ACM, New York, USA, May 2021. doi: 10.1145/3411764.3445720 2, 4, 5, 9

[36] W. Keane. On Semiotic Ideology. *Signs and Society*, 6(1):64–87, Jan. 2018. doi: 10.1086/695387 3

[37] H. Kennedy and R. L. Hill. The Feeling of Numbers: Emotions in Everyday Engagements with Data and Their Visualisation. *Sociology*, 52(4):830–848, Aug. 2018. doi: 10.1177/0038038516674675 2, 9

[38] H. Kennedy, R. L. Hill, W. Allen, and A. Kirk. Engaging with (big) data visualizations: Factors that affect engagement and resulting new definitions of effectiveness. *First Monday*, Nov. 2016. doi: 10.5210/fm.v21i11.6389 2, 9

[39] Y.-S. Kim, P. Kayongo, M. Grunde-McLaughlin, and J. Hullman. Bayesian-Assisted Inference from Visualized Data. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):989–999, Feb. 2021. doi: 10.1109/TVCG.2020.3028984 9

[40] R. Kircher and L. Zipp. An introduction to language attitudes research. In R. Kircher and L. Zipp, eds., *Research Methods in Language Attitudes*. Cambridge University Press, 2022. doi: 10.1017/9781108867788.002 1

[41] H.-K. Kong, Z. Liu, and K. Karahalios. Frames and Slants in Titles of

Visualizations on Controversial Topics. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pp. 1–12. ACM, New York, NY, USA, Apr. 2018. doi: 10.1145/3173574.3174012 2

[42] H.-K. Kong, Z. Liu, and K. Karahalios. Trust and Recall of Information across Varying Degrees of Title-Visualization Misalignment. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, pp. 1–13. ACM, New York, NY, USA, May 2019. doi: 10.1145/3290605.3300576 2

[43] C. Kostelnick. The Visual Rhetoric of Data Displays: The Conundrum of Clarity. *IEEE Transactions on Professional Communication*, 50(4):280–294, Dec. 2007. doi: 10.1109/TPC.2007.908725 9

[44] J. Kuiken, A. Schuth, M. Spitters, and M. Marx. Effective Headlines of Newspaper Articles in a Digital Environment. *Digital Journalism*, 5(10):1300–1314, Nov. 2017. doi: 10.1080/21670811.2017.1279978 2

[45] H. J. Ladegaard. Language attitudes and sociolinguistic behaviour: Exploring attitude-behaviour relations in language. *Journal of sociolinguistics*, 4(2):214–233, 2000. 1

[46] C. Lin and M. A. Thornton. Visualization aesthetics bias trust in science, news, and social media, Dec. 2021. doi: 10.31234/osf.io/dnr9s 7

[47] V. Loureiro-Rodríguez and E. F. Acar. The Matched-Guise Technique. In L. Zipp and R. Kircher, eds., *Research Methods in Language Attitudes*, pp. 185–202. Cambridge University Press, Cambridge, 2022. doi: 10.1017/9781108867788.016 9

[48] N. Mahyar, S.-H. Kim, and B. C. Kwon. Towards a taxonomy for evaluating user engagement in information visualization. In *Workshop on Personal Visualization: Exploring Everyday Life*, vol. 3, p. 4, 2015. 2, 9

[49] D. B. Markant, M. Rogha, A. Karduni, R. Wesslen, and W. Dou. Can Data Visualizations Change Minds? Identifying Mechanisms of Elaborative Thinking and Persuasion. In *2022 IEEE Workshop on Visualization for Social Good (VIS4Good)*, 2022. doi: 10.1109/VIS4Good57762.2022.00005 9

[50] M. Morgenstern. *Play, Purity, and the Poetics of Moral Counterpublic Discourse on Tumblr.com*. PhD thesis, University of Virginia, 2022. 3

[51] M. Morgenstern, A. Fox, G. Jones, and A. Satyanarayan. Visualization Vibes: The Socio-Indexical Function of Visualization Design. *IEEE Transactions on Visualization and Computer Graphics*, 32(1), 2026. 1, 2, 3, 4, 5, 7, 8, 9

[52] K. Murphy. Fontroversy! or, how to care about the shape of language. In *Language and Materiality: Ethnographic and Theoretical Explorations*, pp. 63–86. Cambridge University Press, Jan. 2017. doi: 10.1017/9781316848418.004 2

[53] K. M. Murphy. Fake News and the Web of Plausibility. *Social Media + Society*, 9(2), Apr. 2023. doi: 10.1177/20563051231170606 2

[54] C. V. Nakassis. A Linguistic Anthropology of Images. *Annual Review of Anthropology*, 52(1):73–91, Oct. 2023. doi: 10.1146/annurev-anthro-052721-092147 2

[55] L. M. Padilla, S. H. Creem-Regehr, M. Hegarty, and J. K. Stefanucci. Decision making with visualizations: A cognitive framework across disciplines. *Cognitive Research: Principles and Implications*, 3(1):29, July 2018. doi: 10.1186/s41235-018-0120-9 9

[56] P. Parsons and P. Shukla. Data Visualization Practitioners' Perspectives on Chartjunk. In *2020 IEEE Visualization Conference (VIS)*, pp. 211–215. IEEE Computer Society, Oct. 2020. doi: 10.1109/VIS47514.2020.00049 9

[57] P. C. Parsons. Design Cognition in Data Visualization. In *Visualization Psychology*, pp. 197–215. Springer. doi: 10.1007/978-3-031-34738-2_8 9

[58] E. M. Peck, S. E. Ayuso, and O. El-Etr. Data is Personal: Attitudes and Perceptions of Data Visualization in Rural Pennsylvania. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, NY, USA, 2019. doi: 10.1145/3290605.3300474 1, 2, 9

[59] C. S. Peirce. *Peirce on Signs: Writings on Semiotic by Charles Sanders Peirce*. James Hoopes. University of North Carolina Press, 1991. 2, 3

[60] S. Pinker. Theory of Graph Comprehension. In R. Freedle, ed., *Artificial Intelligence and the Future of Testing*. Erlbaum, Hillsdale, NJ, 1990. 9

[61] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2021. 5

[62] R. A. Rensink. Visualization as a stimulus domain for vision science. *Journal of Vision*, 21(8):3, Aug. 2021. doi: 10.1167/jov.21.8.3 9

[63] J. W. Rettberg. Ways of knowing with data visualizations. In M. Engebretsen and H. Kennedy, eds., *Data Visualization in Society*, pp. 35–48. Amsterdam University Press, 2020. doi: 10.2307/j.ctvzgb8c7.8 2

[64] W.-M. Roth. *Toward an Anthropology of Graphing*. Springer Netherlands, Dordrecht, 2003. doi: 10.1007/978-94-010-0223-3 1, 9

[65] W.-M. Roth. What is the meaning of "meaning"? A case study from graphing. *The Journal of Mathematical Behavior*, 23(1):75–92, Jan. 2004. doi: 10.1016/j.jmathb.2003.12.005 2

[66] J. M. Scacco and A. Muddiman. The curiosity effect: Information seeking in the contemporary news environment. *New Media & Society*, 22(3):429–448, Mar. 2020. doi: 10.1177/1461444819863408 2

[67] T. Schofield, M. Dörk, and M. Dade-Robertson. Indexicality and visualization: Exploring analogies with art, cinema and photography. In *Proceedings of the 9th ACM Conference on Creativity & Cognition*, C&C '13, pp. 175–184. ACM, New York, NY, USA, June 2013. doi: 10.1145/2466627.2466641 3

[68] J. Schwabish. Why Your Organization Needs a Data Visualization Style Guide. https://policyviz.com/2021/03/16/why-your-organization-needs-a-data-visualization-\style-guide/. 2021. 9

[69] P. Shah, E. G. Freedman, and I. Vekiri. The Comprehension of Quantitative Information in Graphical Displays. In P. Shah, ed., *The Cambridge Handbook of Visuospatial Thinking*. Cambridge University Press, New York, NY, 2005. doi: 0.1017/CBO9780511610448.012 9

[70] M. Silverstein. Shifters, linguistic categories, and cultural description. In *Meaning in Anthropology*. UNM Press, 1976. 3

[71] D. A. Szafir, S. Haroz, M. Gleicher, and S. Franconeri. Four types of ensemble coding in data visualizations. *Journal of Vision*, 16(5):11–11, Mar. 2016. doi: 10.1167/16.5.11 9

[72] P. H. Tannenbaum. The Effect of Headlines on the Interpretation of News Stories. *Journalism Quarterly*, 30(2):189–197, Mar. 1953. doi: 10.1177/107769905303000206 2

[73] E. Tufte. *Envisioning Information*. Graphics Paper Press LLC. 9

[74] E. Tufte. *Visual Explanations: Images and Quantities, Evidence and Narrative*. Graphics Paper Press LLC. 9

[75] E. Tufte. *Visual Display of Quantitative Information*. Graphics Paper Press LLC, Cheshire, CT, 1983. 9

[76] B. Tversky. Cognitive Principles of Graphic Displays. 1997. 9

[77] E. Wetzel, J. R. Böhnke, and A. Brown. Response biases. In *The ITC International Handbook of Testing and Assessment*, pp. 349–363. Oxford University Press. doi: 10.1093/med:psych/9780199356942.003.0024 5, 7

[78] C. Xiong, L. Van Weelden, and S. Franconeri. The Curse of Knowledge in Visual Data Communication. *IEEE Transactions on Visualization and Computer Graphics*, 26(10), Oct. 2020. doi: 10.1109/TVCG.2019.2917689 9