

Crossing the Chasm: Bridging Visual Augmentations and Designer Intent

Luke S. Snyder, Maureen C. Stone, Jeffrey Heer

Abstract—To direct attention and communicate context, visualization designers often employ augmentations such as annotations, animated transitions, and stylistic layouts. While graphical perception principles can helpfully prescribe effective low-level representations of data, they lack guidance for augmentations that service higher-level communication goals. To bridge this gap, we contribute a design space that frames designer intent as viewer-oriented cognitive behaviors, grounding communicative aims in actionable visualization design techniques, including both visual encodings and augmentations. This design space consists of common augmentation tactics (annotation, animation, stylized encodings, etc.) that implement design strategies (spotlighting, sequencing, association, etc.) to achieve higher-level design goals (observe, interpret, introspect), often simultaneously. We demonstrate the analytic and generative value of our design space with examples across varied designer objectives. We discuss how our contributions help align designer intent with reader takeaways, pave the way for readers to learn more effectively, and enable future (semi-)automated systems to support visualization designers in achieving their communication goals.

Index Terms—Visual augmentations, communicative visualization, visualization design.

I. INTRODUCING THE CHASM

VISUALIZATIONS are used as communicative aids in contexts such as teaching, business, news, and scientific papers. During the design process, designers must consider their audience, what message they want to convey, and how they want to convey it. Designers typically convey this intent by augmenting their visualizations (e.g., annotations, animation, semantic color encoding, etc.), but how can they ensure that their intent was correctly decoded by the reader? Graphical perception literature provides grounding for effective low-level representation of data [1], [2], but we lack similar guidelines to help designers express communicative intent.

We use “communicative” and “designer” intent interchangeably throughout the paper. This can be defined at multiple levels, each increasingly difficult to assess. For instance, the designer’s intent may be rhetorical [3] or affective [4] in nature: to persuade or evoke an emotional response. On the other hand, the goal may be for readers to learn something important [5] or modify their future behavior based on new knowledge. Each of these goals, however, may be secondary to less savory intents such as deceit or profit. We frame

communicative intent as explicit properties of the design: **general cognitive actions that the designer wishes to induce on behalf of the reader**, such as *observing* an important fact about the data or *reflecting* on their own knowledge. Though these behaviors may still be in service of higher-level goals, this more abstract framing has two primary advantages: (1) we are able to more concretely reason about design objectives from a varied set of augmentations used throughout the literature, in contrast to more specific and complex objectives (e.g., learning, emotional, etc.) that are difficult to relate to each other and explicate without discussions or assessments with designers and readers; (2) designers are encouraged to reflect on the viewer’s behavior, using our objectives as actionable guides to help achieve their ultimate intents.

Consider the political issues that US voters associate with the Democratic Party across past decades (Figure 1). After the data has been encoded, a number of decisions were made that affect communicative intent [5], [6]. For instance, color not only helps relate annotated text with the data, it also exploits cultural associations to ground meaning (i.e., visual metaphors [3]): green for climate change, blue for working class (“blue-collar”), and pink for healthcare (breast cancer ribbon). The placement of the annotated text also orders the content left-to-right to convey the temporal shift from working class to other issues more clearly. However, if additional design constraints were present (e.g., more data, heavy overlap of lines, etc.), other sequencing techniques might be used instead, such as animation (e.g., fading to layer in each line), annotated labels (to enumerate content), faceting, a stepper, etc. Additional annotations with historical context could help *explain* voters’ shift, but would alter the communicative intent. Each of these decisions—or lack thereof—make use of **visual augmentations** (defined in section II) to direct readers’ attention and convey context. These objectives and supporting augmentations move beyond standard graphical perception concerns to influence the message that the designer conveys and what readers may learn.

A gap remains between design intent and the specific visual augmentations: how might a designer’s communication goals be achieved within a vast space of varied augmentations? To bridge this gap, we contribute a design space that characterizes connections between *how* designers augment their charts and *why* they employ these augmentations. We survey literature in narrative visualization, graphical perception, communication, and learning sciences to distill a design space with three tiers: a common set of *augmentation tactics* (annotation, animation, stylized encodings, etc.) that implement *design strategies* (spotlighting, sequencing, association, etc.) to achieve higher-

This is the author’s version of the article that has been published in IEEE Transactions on Visualization and Computer Graphics. The final version of this record is available at: <https://doi.org/10.1109/TVCG.2026.3671259>.

This research was supported in part by an NSF IIS Award (2402718).

The authors are with the University of Washington, Seattle, WA, USA (e-mail: snyderl@cs.washington.edu; mestone@cs.washington.edu; jheer@cs.washington.edu).

level *design objectives* (observe, interpret, introspect). We find that many tactics and strategies are reused across other tiers, often achieving multiple goals simultaneously. In addition, our design space can be aligned with prior work on learning objectives [5], [6] to support more systematic design and assessment in the future.

Our design space serves as both an analytic tool to evaluate and compare communication aims across different visualizations, as well as a generative aid for exploring new designs with competing augmentation strategies and objectives. We unpack examples in communicative visualization to demonstrate these contributions, noting trade-offs and opportunities.

By grounding communicative intent in visual augmentations, designers gain more actionable techniques—including visual encodings and augmentations—that they can use as communicative “levers” to convey different messages. These “levers” may also encourage reflection about what readers will take away, better aligning designers and readers. We discuss this alignment in greater depth, along with how our contributions can help pave the way for readers to learn more effectively, enable future (semi-)automated systems to support visualization designers in achieving their communication goals, and help scaffold and evaluate visual augmentations.

In summary, we make the following contributions:

- 1) We identify and characterize the gap between visual augmentations and design intent as a disconnect between *how* designers augment their charts and *why*.
- 2) We formally bridge this gap by surveying the visualization literature to distill a design space of *what* augmentations designers use (tactics), *how* they use them (strategies), and *why* (to what aim) they use them (objectives).
- 3) We demonstrate the analytic and generative value of our design space across examples with varied designer intents, identifying both trade-offs and opportunities.

II. CHARACTERIZING THE CHASM: BACKGROUND

To illustrate the gap between visual techniques and designer intent, we first define visual augmentations, using motivating scenarios (Figure 1, Figure 2, Figure 3) to surface potential types of designer objectives and supporting aids.

Defining Visual Augmentation. We define visual augmentation as *any use or modification of existing or supplemental graphical elements designed to communicate information beyond denotative graphical decoding*.

Our definition broadens the scope of augmentations explored by prior work, such as annotations [7] (which constitute a subset of visual aids) and augmentations used to highlight data relationships [8], [9] (which represent one kind of design goal). In addition, we recognize that non-supplemental, standard encoding choices (not only new ones [10]) made during the initial data mapping may also serve communicative goals (e.g., in addition to categorizing data, red and blue colors connote the Republican and Democratic parties in the United States). We fully operationalize our definition in section IV, identifying augmentation tactics, design strategies, and higher-level goals to form a design space.




Fig. 1. Augmented visualization by Burn-Murdoch [11], which illustrates the change in perceptions of Democratic Party values over time (A). (B) The placement of annotated text subtly encourages left-to-right reading; (C) Color hue is used to relate annotated text with line marks and connote subject matter (e.g., blue for blue-collar workers or green for climate); (D) Footnotes provide context for how the data was sourced.

Example: Democratic Party Trends. Perceptions of political party values have shifted over time. This shift within the Democratic Party is demonstrated in Figure 1, with voters’ increasingly associating the party with issues related to marginalized groups, healthcare, and climate regulation. The title (Figure 1(A)) and footnote (Figure 1(D)) serve as **guides** for additional **context** on what the data represents and how it was sourced, aiding **interpretation**. **Annotated text** (Figure 1(B,C)) provides additional **context** for each of the line marks. The placement of these annotations **sequence** the reader left-to-right, prompting them to **observe** the shift in associations from working class to other topical issues. Finally, color hue (Figure 1(B,C)) assists **interpretation** by **relating** annotated text with each line mark and connoting higher-level **associations** (e.g., blue for blue-collar workers, pink for breast cancer awareness, and green for the climate).




Fig. 2. Augmented visualization by Tufte [12] that shows weekly cholera deaths during the 1854 London cholera outbreak. (A) Casket symbol that represents 50 deaths; (B) Pictures of Queen Victoria and John Snow; (C) Annotated text that contextualizes when the handle from the Broad street pump was removed and by whom; (D) Annotated connector that identifies the when the pump was removed; (E) Aggregation to display weekly deaths.

Example: Cholera Outbreak. London suffered a cholera epidemic in 1854. Physician John Snow surmised that the source was contaminated water, removing the Broad Street pump to test his hypothesis. Figure 2 visualizes Snow’s actions within the context of the epidemic, with several visual augmentations employed to support communication goals. The visualization makes use of **textual** and **connector** annotations (Figure 2(C, D)) that provide **context** (Figure 2(C)) as well as **spotlight** (Figure 2(D)) when the Broad Street pump was removed.

Rather than use standard




Fig. 3. Augmented visualization by Victor [15], demonstrating interactive and explorable explanations of a car simulation algorithm. (A) Linked color highlighting to indicate mouse hovering, which sets the algorithm parameter (degrees of turn) in views (D) and (E); (B) Linked color highlighting to indicate mouse clicking, which sets a specific parameter value in views (D) and (E) and focuses the region of interest (C); (D) Slider to explore different parameter values.

bar marks, the designer encoded the number of deaths with symbolic caskets (Figure 2(A)) and includes pictures of historical figures (Figure 2(B)). These symbolic representations may help the reader **associate** the meaning of the data with known referents, which can aid **interpretation** of the subject matter. It is worth noting that the decision to include these associations was strongly criticized by Tufte [12]. Indeed, this decision may have little effect on the reader's ability to accurately decode the information presented in the chart, a topic of some debate [13], [14]. However, with respect to higher-level communication objectives, symbols may more directly convey a visualization's thematic content by engaging a reader's prior associations.

Finally, the chart **aggregates** data by week (Figure 2(E)) to help the reader **observe** the reduction in deaths after the handle was removed. This tactic must be used with care: depending on the context, aggregation achieves rhetorical aims or it may mislead the reader by hiding data that conflicts with a chosen narrative [3], [12]. Here, the number of *daily* deaths were already decreasing before the handle was removed.

Example: Exploring Algorithms. Explorable Explanations surface underlying models in the form of interactive content for readers to investigate, test, and refine both the assumptions that the model and themselves make [16]. One example of this type of content is demonstrated in Figure 3, which communicates different layers of abstraction in a car simulation algorithm [15]. As in previous examples, repeated use of linked colors help **relate** different content and functionality, aiding **interpretation**. Red coloring (Figure 3(A,C,E)) is used consistently throughout the article to help the reader **observe** specific interactive abilities (e.g., mouse hovering in Figure 3(A)), as well as **spotlight** selected content in **faceted** views (Figure 3(C,E)). Underlined blue text (Figure 3(B)) represents snapshots of different model states, which the reader can click to **spotlight** corresponding content (Figure 3(C)).

Finally, an interactive slider (Figure 3(D)) enables the reader to explore and test the algorithm across faceted views. Critically, this functionality encourages the reader to **introspect** by **self-reflecting** and **repairing** their knowledge. Specifically, the reader engages in a cycle in which they make tacit **predictions** about the algorithm, use the interactive slider to test these predictions, and revise their assumptions if the **simulated**

content differs from the expected outcome.

Summary of Examples. These examples illustrate varied design tactics (encodings, annotations, interactive simulations), strategies (spotlight, contextualize, relate), and objectives (observe, interpret) to scaffold *how* and *why* visual augmentations are employed. Crucially, a single tactic can serve different strategies and objectives, sometimes simultaneously. For example, as in Figure 1(A), annotations may **sequence** data to guide readers in **observing** important data insights or phenomena, as well as provide **context** to connect and explain these insights, aiding **interpretation**. Consistent use of colors can both **relate** and semantically **associate** content for **interpretability** (Figure 1(B,C)).

During the design process, authors must consider (1) their communicative objectives, (2) what encoding and augmentation techniques are available to them, (3) which techniques are appropriate for their objectives, and (4) trade-offs and interactions between different techniques. Currently, the literature provides guidance on points (1) and (2), but stops short of relating them to achieve points (3) and (4).

III. EXAMINING THE CHASM: RELATED WORK

We review literature on graphical perception, visualization taxonomies, communicative visualization, and learning sciences to note current gaps in scaffolding communicative aids and providing designers with actionable design guidance.

A. Graphical Perception

While the creative and effective use of visualization to leverage human cognition can be traced back as early as William Playfair's *The Commercial and Political Atlas* [17], Bertin's *Semiology of Graphics* [1] arguably contributes the first formalism for visual encodings and the appropriate mappings from different types of data (e.g., categorical vs. quantitative). These principles have been empirically studied [18], [19] and mapped to effectiveness rankings for automatic systems [20], smart defaults in visualization systems [2], [21], [22], and general criteria for appropriate chart design (e.g., Mackinlay's expressiveness and effectiveness principles [2]). Recent work has shown that perceptual principles which govern our visual system (e.g., Gestalt principles [18]) can even affect higher-level comprehension, as readers focus on certain data relationships (e.g., trends) which affects their insights and interpretation of a chart's intended message [23], [24]. In part, this work empirically validates early efforts of psychologists to understand how humans comprehend graphics, as well as why various chart types and encodings are better suited for certain tasks than others (e.g., using bar chart to compare overall trends vs. a line chart to comprehend changes in slopes) [18].

Graphical perception principles helpfully prescribe effective low-level representations of data, such that readers can (on average) accurately decode visualized data. However, visualizations within a communicative context may be viewed more appropriately as a learning problem [25]. As such, while perceptual principles might help readers *quickly* and *accurately* comprehend a chart (applying a cognitive-efficiency perspective), they fail to guide designers in reasoning about

and achieving higher-level communication and learning goals. Within the context of visual augmentations, graphical perception work may guide *how* the augmentation should be implemented (e.g., use color to encode categorical data or animation to convey change [26]), but not *why* in the context of communication (e.g., use red / blue color to culturally associate data with Republican / Democratic parties, or an interactive stepper that sequences the order of content to improve the reader’s interpretation of the narrative).

B. Visualization Tasks & Evaluation

Most visualization taxonomies consider the role of the data analyst, providing useful guidance for appropriate types of charts [9] or interactions [27] based on that context. Low-level task taxonomies operationalize graphical perception principles as different analytic tasks that the user might perform (e.g., compare distributions, retrieve / identify values, etc.) and the appropriate representations for these tasks (e.g., scatter plot for detecting outliers) [9]. Brehmer and Munzner’s typology of abstract visualization tasks [27] mirror our efforts in the analytic space, bridging the *why*, *what*, and *how* for analysis tasks. However, this typology relegates all communicative objectives to a single abstract category: *present*. As Adar et al. [5] note, although this decision can reflect the constant cycle between analysis and communication [28], [29], it fails to delineate nuanced interactions among varied communication aims. Naturally, this design process involves a different, possibly more complex, set of questions beyond the scope of analysis – namely, *what* to communicate, *how* to communicate it, *why* it should be communicated and *to whom* it will be communicated (i.e., the audience). By bridging design intent with augmentation tactics (i.e., how designers may realize their intents), our design space provides an initial foundation to help answer these questions, including reasoning about tradeoffs and interactions between visual augmentations which can be formally assessed in future studies.

Commonly, the effectiveness of a visualization is evaluated with a suite of cognitive-efficiency metrics (e.g., accuracy, speed, size) that focus on data insights [23], [24]. However, in isolation these metrics are often not the most appropriate for communicative visualization. Rather, the primary concern is whether the reader *actually* learned what the author intended [5], as opposed to how quickly or accurately they could decode insights, some of which might be tangential to the original intent. Although the use of visual augmentations outlined in our design space might require additional effort and time to decode (e.g., annotations, explorable explanations [16], etc.), we expect that what the reader *learned* will be closer to the designer’s intent. We leave this evaluation as future work.

C. Communicative Visualization & Design

A large body of work helps scaffold the design dimensions of communicative visualization across multiple contexts, including data storytelling and narration [30], business dashboards [31], interactive articles [32], [33], [34], and research papers [35]. Although these design spaces address *how* designers might author a visualization, they remain disconnected

from *why* (to what aim). As a result, designers lack guidance when reasoning from high-level design goals, including how their final design might be improved to better align with reader takeaways. For instance, when are *explorable explanations* appropriate (which have been used incorrectly, much to Bret Victor’s chagrin [16]), and what are alternative techniques depending on the viewing context?

We take direct inspiration from this line of work, asking, “*How do varied techniques relate to designers’ communicative goals?*”, and use literature in this area to scaffold our design space, noting similarities in how divergent tactics can be used to accomplish similar goals, often simultaneously. A central benefit to bridging visual augmentations with design intents is the ability to identify potential mismatches with what readers actually learn. For example, under what contexts (e.g., background knowledge, domain, learning goals etc.) do narrative systems (e.g., Idyll [32], explorable explanations [16], playable guessing games [36]) provide learning benefits to readers? To answer this question empirically, we must first unpack the intended uses and interactions behind such augmentations, including *when* and *why* an augmentation outperforms others.

There exists a long-standing debate as to under which circumstances visual augmentations are beneficial to readers and to what the extent they are mere “chart-junk” [14]. Tufte’s well-known prescription to maximize the data-to-ink ratio of a chart [37] (i.e., representing the most amount of (necessary) data with the smallest amount of “ink” in a chart) has been shown to not always equate with maximal comprehension or learning [13], a central aim of communicative graphics. Hullman et al. [13] characterize these differences as visual difficulties, such as charts with lower data-to-ink ratios that are more likely to engage the reader. As an alternative to cognitive-efficiency evaluations, these benefits may arise (in certain contexts) by inducing active processing and self-explanation to better manipulate internal representations and expectations of data [38]. Visual augmentations may also provide novelty, aesthetics, personalization, challenge, and reflection, promoting engagement [13]. By supporting learning outcomes at the expense of cognitive efficiency, we expect future assessments will help formalize the opportunity cost of visual augmentations, including diminishing returns (e.g., too many augmentations, additional augmentations needed for more complex types of content or domains, etc.).

D. Visualization Objectives & Extant Design Spaces

Prior work classifies visualization objectives at multiple levels: perceptual, descriptive, and cognitive. Perceptual goals are often insight-driven [39], [40], [41], [42], ranging from low-level tasks such as summarization, comparison, and identification [7], [9], [27], [43] to high-level comprehension (e.g., reader takeaways) [18], [19], [44], [45] or aesthetics such as legibility [37], [45]. Most design spaces, however, often only relate these objectives to a single *tactic* in our design space, such as annotations [7], text [43], [46], or encodings [2], [47].

Descriptive objectives delineate *what* a visualization or augmentation is intended to accomplish, which roughly correspond to *strategies* in our design space. For example, prior

work in narrative visualization [30], [48], [49] relates specific implementation features (e.g., zooming, progress bar) to structural devices affecting narrative flow (e.g., highlighting, transitions, sequential ordering). We organize implementation features into tactics to more clearly delineate important differences in how divergent tactics are used (e.g., zooming, animated transitions, and annotations can each be used to highlight or sequence content, though annotations may be used in other ways, such as to provide external context).

Cognitive goals are the highest level and often the most difficult to assess. These objectives may span the following contexts: learning [5], [6], [50], [51], affective or emotional [4], rhetorical [3], engagement or reflection [34], [52], [53], efficiency [13], and reasoning or decision-making [54]. However, these objectives may comprise sub-goals themselves. In the context of learning, “recall” / “remember” may be viewed as the first level (though this remains controversial [55]) before progressing to higher levels such as “understanding”, “application”, and “analysis”. Though “recall” may be used as a proxy for comprehension [13], [14], [56], this begs the question: *why* does the designer want the viewer to remember certain information? Cognitive goals (affective, rhetorical, learning) may, for example, be in service of specific *outcomes*, such as modifying the viewer’s behavior, improving trust, or profit.

A central issue is that prior work often discusses perceptual and cognitive goals interchangeably at different levels of abstraction that are difficult to define and relate. For instance, graphical perception work discusses “interpretation” in the context of visual encodings and semantics [18], [19] (i.e., which marks are better at denoting different types of data and relationships?), while “interpretation” as a learning goal [6] is often higher-level and depends on both the encodings and visual augmentations (e.g., using annotations to help readers’ interpret a correlation or summarize key temporal stages [5]). Certain learning verbs (e.g., summarize, compare) also overlap with analytic visualization tasks in the context of interaction [27] or chart encodings [9], which obfuscates what the viewer does **cognitively** versus what the viewer *can* do **analytically** (e.g., via interaction).

Conceptually, our visualization objectives (observe, interpret, and introspect) abstract descriptive goals into explicit properties of the chart that reflect what the viewer will *do* as a result of a visual augmentation. These objectives are partitioned in a general manner to help relate perceptual and cognitive aims through our objective hierarchy: observation augmentations serve primarily as attention aids (e.g., to identify, compare, etc.), while interpretation and introspection aids modify what the viewer interprets from this attention and how they may assess their interpretation. Designers can then use these objectives as guides to reflect on the viewer’s behavior and reason about how to achieve higher-level cognitive goals (e.g., learning, rhetorical, etc.).

Our design space reflects general-to-specific structures between low-level verbs and high-level goals from prior taxonomies [27], [30], [49]. However, several gaps remain in extant taxonomies:

G1 Communication goals aimed at viewers are often relegated to a single category [27] or not abstract enough

to relate varied augmentations [30], [49].

G2 Communication goals are described without connection to the full range of tactics [5], [13].

G3 Augmentations are only discussed for a single tactic [7], [57] or context [30], [49].

G4 Roles and relationships among augmentations at different levels are not delineated, including (1) similar uses of divergent tactics, (2) different uses of the same tactic, (3) or using a single tactic to achieve multiple goals simultaneously.

These gaps pose two primary risks: (1) the visualization community cannot adequately analyze and evaluate diverse communication visualizations across a variety of presentation contexts due to a lack of shared goals and language for describing visual augmentations; (2) without a sufficiently nuanced vocabulary, designers lack tools that both encourage reflection about what the viewer does and guide them with different strategies and tactics for achieving these goals, resulting in potentially sub-optimal designs that limit reader comprehension and learning.

Our design space defines augmentations at multiple levels of specificity (tactics, strategies, objectives), enabling us to capture nuanced roles, relationships, and dependencies (G2-4). To provide scaffolding, we abstract perceptual and cognitive goals into more general categories that describe viewers’ cognitive behaviors (G1). Since the designer cannot “force” their ultimate goal (e.g., that the viewer will have learned exactly what they intended), our objectives can also be interpreted as “guides” to help designers achieve such higher-level intents.

E. Communication & Learning Sciences

Our interest in the relationship between the *use* and *intent* of visual augmentations was inspired heavily by Adar and Lee-Robbins [5]. They frame communicative visualization as a learning problem, building off *process*- [58] and *meaning* [1]-focused communication theories and extending Bloom’s taxonomy of educational learning objectives [50]. This framing provides important scaffolding for designers to both describe what they want readers to actually learn and use assessment to highlight gaps in the design. We view our design space as a next step in this direction, seeking to provide more actionable design techniques to control different learning “levers”. While we reason about how our designer objectives, strategies, and tactics (section IV) might relate to Adar’s learning taxonomy [5] (section VI), we leave a more exact mapping to future work given the need for further empirical grounding. One benefit of keeping an arm’s length between our notion of *design intent* and learning objectives is that our intent goals might be related more cleanly to additional objectives, including affective [4] and rhetorical framing [3], [43] goals.

Our design space builds on theoretical foundations in communication theory, cognitive psychology, and semiotics [1], [18], [19]. Central to this foundation is the notion of syntax (i.e., how marks are visually interpreted based on spatial grouping and organization), semantics (i.e., what the configurations of marks represent, such as data values, axes, legends, labels, etc.), and pragmatics (i.e., how indirect information

is conveyed beyond semantic interpretation, which may be affected by the context of viewing, how much information is included, etc.) [19]. Graphical perception literature helpfully distills principles grounded in syntactic and semantic notions [1], [2], [20], elucidating how humans comprehend visual information (e.g., use of shape encoding conveys categorical semantics effectively, but not quantitative). However, we find that visual augmentations often extend beyond denotative semantics to drive connotative meaning and pragmatic intent. For example, the use of caskets instead of bars (Figure 2) may evoke emotion or serve rhetorical aims [3] (e.g., shifting the focus to John Snow’s heroic efforts, or commenting on the religious influences in medicine at that time).

We consider the manner and context in which augmentation tactics might be used, reflecting both denotative semantics and varied connotative and pragmatic goals. For example, color encodings may be used to spotlight specific content (denotative semantics), associate data with cultural concepts (Figure 1(B,C)) (connotative semantics), or relate groups of content together due to spatial proximity (Figure 1(B,C)) (pragmatics). The distinction between denotative and connotative semantics and pragmatics helped us partition and reason about different uses of tactics more effectively: observation aids often affect denotative semantics, while interpretation and introspection tactics support more complex connotative and pragmatic aims that consider the background of the reader, the complexity of the data, and in what context the visualization will be viewed.

Visual augmentations in communicative contexts other than data visualization have been explored, such as mathematical notation [46], constrained highlighting for text [59], and code debugging [60]. We find that our design space shares certain tactics with these domains (e.g., annotation, use of color to link concepts), suggesting that some augmentations might share communicative grounding and learning benefits more broadly. Our work considers additional tactics central to spatiovisual data representations (e.g., animation, sequencing, & layout), including their relationship with varied design goals.

IV. BRIDGING THE CHASM: A DESIGN SPACE OF VISUAL AUGMENTATIONS

To bridge visual augmentations and designer intent, we sought to (1) identify instances of visual augmentations throughout the visualization and perception literature, (2) characterize higher-level design goals, and (3) ground augmentations in these goals. This survey resulted in a design space with three tiers: objectives, strategies, and tactics (Table I). *Objectives* encompass core viewer-oriented goals by which authors seek to communicate the message of their visualization, such as directing attention to important observations or adding context to aid a reader’s interpretation. *Tactics* are lower-level implemented design techniques, such as animated transitions or stylized encoding choices, that can influence viewer attention, interpretation, and consideration. *Strategies* are approaches the author selects to connect the two, applying one or more tactics to achieve an objective (e.g., both animation and encoding may be used to emphasize an observation).

Type	Augmentation	Description
Objective	Observe	Direct attention to one or more elements
	Interpret	Manipulate semantic or pragmatic intent
	Introspect	Promote learning and reflection
Strategy	Sequence	Arrange content in a particular order
	Spotlight	Focus viewer attention on specific content
	Relate	Conceptually link disparate data or views
	Associate	Ground representations in known referents
	Contextualize	Supplement chart with additional information
	Reflect	Self-examine knowledge
Tactic	Repair	Self-correct knowledge
	Transform	Apply transformations to data
	Animate	Temporally sequence static frames
	Layout	Spatially arrange two or more views
	Annotate	Introduce notational graphical elements
	Encode	Map data attribute(s) to visual channel
	Symbolize	Symbolically represent object or idea
	Guide	Describe data semantics or thematic content
	Query	Probe reader’s knowledge directly
	Predict	Self-predict unseen data
	Explain	Self-explain observed data
	Evaluate	Critique elicited prior knowledge
	Simulate	Update visualized data, driven by reader input

TABLE I
OVERVIEW OF VISUAL AUGMENTATIONS.

A. Survey Methodology

We seeded our survey of visual augmentations with several key papers and books in communicative visualization, visual learning, and graphical perception, such as Segel and Heer’s design space of narrative visualization techniques [30], Tversky’s evaluation of the effectiveness of animation for learning [26], Adar and Lee-Robbins’ extended learning objectives for communicative visualization [5], [6], Tufte’s *Visual Explanations* [12], and McCloud’s *Understanding Comics* [61]. We then snowball-sampled papers across related areas, such as uncertainty visualization [47], interactive articles [16], [32], [34], augmented visual tools [8], existing design spaces [49], [62] and visualization taxonomies [9], [27], [63]. Additionally, for completeness, we widened our search by searching for key terms and phrases on Google Scholar, such as “communicative visualization”, “visual augmentations”, “visual learning”, etc. In total, we collected an initial set of 70 papers.

As with prior design spaces [30], we do not claim that our sample is exhaustive, though we optimized for a rich set of diverse themes within the scope of communicative visualization, prioritizing a diversity of goals (e.g., perceptual [9], cognitive [5], etc.), media forms and contexts (e.g., narrative journalism [30], comics [49], dashboards [31], etc.), and augmentation roles (e.g., annotation [7], animation [26], self-prediction [52], etc.). We also found that additional relevant papers contained either repetitive ideas or were described by taxonomies already present in our corpus (e.g., animation in learning contexts, types of communication devices in visualization, etc.). Still, we acknowledge some inevitable subjectivity in this process. We removed 18 papers from our augmentation coding process that: described graphical comprehension (2/18), visual learning (3/18), or attention (4/18), but not augmentations therein; supported objectives outside the cognitive domain such as affective (1/18), rhetorical (1/18), or social engineering (2/18); or provided textual augmentations not primarily focused on aiding visualization (5/18).




Fig. 4. Overview of augmentations by Objective (Observe, Interpret, Introspect), Strategy (Sequence, Spotlight, Relate, Associate, Contextualize, Reflect, Repair), and Tactic (Transform, Animate, Layout, Annotate, Encode, Symbolize, Guide, Query, Predict, Explain, Evaluate, Simulate). The objectives span from lower-level perceptual aims to higher-level cognitive aims: observation goals seek to direct attention by emphasizing and sequencing content; interpretation aids add supplemental information affecting inference; and introspection helps readers assess and update their understanding.

We open-coded the resulting set of 52 papers, deriving a shared set of 12 tactics, 7 strategies, and 3 objectives, with the relationships between these categories depicted in Figure 4. We describe this open-coding process in more detail, revisit our paper corpus to categorize each paper along our design space dimensions, and provide a selected set of additional examples in the Appendix.

B. Augmentation Tactics

We first describe tactics, which categorize common augmentation techniques and their implementation (e.g., interactive features, context of use, etc.). Tactics indicate *how* an augmentation achieves its design goal.

Transform Designers employ a variety of data transformations including derivations such as regressions [7], [8], [10], segmentation (drill-down) [10], or normalization [54] (Figure 5(A)). Filtering transforms [9], [27] may also structure visual information or highlight particular insights.

Animate Animation is widely used across a variety of contexts [13], [51], [57], [62], [64], [65], including facilitating transitions between statistical views [63], [66], [67], providing narrative structure [30], [64], [68], aiding learning of complex processes [26], [51], or to convey real-time change [26], [66]. Clever use of animation can provide tacit tutorials for interactions [30]. Judicious use of motion and timing can also help group, filter, or understand data relations [69] (Figure 5(B)).

Layout Layout refers to any arrangement of multiple views in space used to guide readers or facilitate examination. Faceting is heavily employed [2], [12], [34], [44], [48], [70], [71], [72] to aid visual comparison [16], [73] (Figure 5(C)). Space can also structure content in a particular sequence to guide readers [48], including the use of scrolling [33], [34], steppers [64], [65], or ordering to fit a particular narrative [30], [49], [61]. Content may also be superimposed to facilitate engagement and relate distinct views (e.g., before and after images [12]), or hide supplemental information that can be displayed on-demand via mouse hovering or toggle buttons [34].

Annotate Annotation primarily directs attention or adds useful context [7], [8], [12], [27], [32], [33], [44], [45], [61], [65], [73], [73], [74] (Figure 6(A)). Annotations help add / layer useful marks such as glyphs, threshold lines, and trends [10],

[66] to focus data characteristics. Occasionally, annotations may be dynamic for details-on-demand, which can be revealed through click or hover interactions [30], [68], [75]. Designers also annotate for context (e.g., using a question mark to indicate missing values [54] or uncertainty [47], [73]) or to guide readers through explanations [73] or sequences [74].

Encode Encoding refers to a broad set of augmentations that exploit visual channels or mappings of data to convey higher-level information [8], [9], [10], [27], [47], [54], [65], [74], [76] (Figure 6(B,C)). Common channels to emphasize particular information include brightness and size [66], color [9], [45], [66], [72], or the shapes of marks themselves [10]. Encodings, such as color, can also help order [10] or link [16], [33], [45], [65], [71], [72], [75] items within a view. We also noted that the layout *within* a single view (e.g., by modifying the position encoding via sorting or stacking [9], [23], [24], [77]) could be used to arrange items in a particular way to emphasize or order content [23], [24].

Symbolize Symbolize refers to the use of icons, pictures, or other referents to ground and concretize the meaning of a visualization, employed broadly across designs [12], [13], [54], [62], [74], [76], [78] (Figure 6(C)). For instance, Conlen et al. [33] utilized an interactive guitar to contextualize an article on tuning. Other uses might include annotated question marks to reference uncertainty [47], [73] or unit visualizations [34]. The degree to which symbols are used within a visualization can constitute chart junk [13], [14], although in proper contexts, pictorial illustrations can aid learning [6] and storytelling [30], [49], [61].

Guide Guides refer to a set of contextual aids that help communicate the visualization’s message, intent, or provenance. Examples include captions / titles [30], [79], footnotes [72], personalized information such as date, time, location [25], [72] (Figure 7(A)), the level-of-detail of text [13], [34], [80], [81], or progress bars to aid navigational context [30], [76]. Designers might also include links to external data or views [73] to indicate the source or methods of data collection.

Query Query mechanisms incorporate quizzes, potentially spaced out over time with recurrence, to test reader knowledge directly and encourage learning [34], [82] (Figure 7(B)).

Predict Prediction tactics engage readers in “guessing” the




Fig. 5. Example augmentation tactics to **observe**. (A) Applying a normalization transform to COVID-19 data, which alters the sequence of data [54]; (B) Animated sequence of world population demographics from young to old to draw attention to the gradual changes between each age group [69]; (C) overview+detail with faceted layout to sequence a horse’s trot, revealing that all four hooves break contact with the ground [34].

data, encouraging self-reflection (Figure 7(C)). These include graphical elicitation (i.e., drawing) [34], [52], interactive widgets and annotations to facilitate active reading / question generation [16], [32], [33], [34], [68], [71], [75], and games [36].

Explain Explanation tactics are mechanisms that motivate readers to self-reflect, potentially using inference to fill in missing gaps in their knowledge [52]. Examples include self-explanation summaries [34], [52] or self-drawing to explain complex processes [78].

Evaluate Evaluation techniques assess reader predictions, including graphical feedback [52], results of quizzes [82], and game win / loss [82]. Self-evaluation can also take place by hiding contextual information behind interactions, allowing readers to in effect use visualizations as flash cards [34].

Simulate Readers’ input with a visualization or interactive article can drive / simulate dynamic system behavior, enabling them to compare system output with tacit predictions about their input: [16], [32], [33], [34], [36], [68], [72], [75].

C. Designer Objectives

Augmentation tactics are used to implement broader design objectives. From our survey, we identified three broad categories of goals that guide design decisions, signaling *why* designers use varied augmentations. These objectives operationalize intent as core viewer-oriented mechanisms to effect differences in chart (1) attention and observation (i.e., what readers notice and in what order), (2) comprehension and inference (i.e., the message and intent that the reader discerns, potentially driven by attention aids), and (3) self-assessment (i.e., enabling the reader to explore and refine their comprehension).

We intentionally characterize design intent according to more general cognitive descriptors such as attention and understanding (rather than, for example, specific learning [5] or rhetorical [3] goals) to relax conceptual constraints, enabling us to relate and organize augmentation tactics more meaningfully and encode communication design as actionable techniques. To connect augmentation tactics with objectives, we surface design strategies that represent the manner in which divergent tactics are instrumented to achieve similar goals, grounding tactics in objectives as a result. Concretely, strategies represent *how* a tactic contributes to an objective.

Observe A variety of augmentations direct readers’ attention [24], [68], [73] to specific areas within a visualization. However, directing attention often serves the broader aim of

helping the reader *observe* specific or important aspects of the data. We categorize these augmentations as *Observation* aids, which help the designer more readily communicate insights for the reader to observe and takeaway [7], [73]. The following two strategies support this objective, commonly via tactics that facilitate important lower-level perceptual tasks, such as comparison, identification, or summarization [7], [9], [27], [44], [62], [83]:

1) **Sequence**: Sequencing refers to the general strategy of arranging content in a particular order. **Animation** can be effective tactic for readers to observe important aspects about the *manner* of change [26] (e.g., *how* temporal data changes over time). For instance, items might be unveiled (e.g., fading) in a particular order [62] to draw attention to specific changes between data values (Figure 5(B)). Spatial **layout** tactics, such as coordinated multiple views [73], faceting [2], [16], [72], or structured narratives [48], can draw readers’ attention to important differences that are more readily apparent between sequential views [10], [48], [72], [73]. A historical example of note is that sequenced frames helped people observe that all of a horse’s hooves break contact with the ground during a trot [34] (Figure 5(C)).

Further, position and color **encodings** (e.g., stacking / sorting [23] or sequential color hue [10]) can affect how readers group visual marks into separate arrangements to make comparisons [23], [24], [77] or observe salient trends more easily [27] based on Gestalt principles of perception [18]. In addition, certain **transformations** such as normalization can substantially alter a data distribution and, as a result, the order of a trend (Figure 5(A)).

While **annotations** are frequently used as attention or contextual aids [7], they can also explicitly sequence content. For example, annotated numbers or arrows can enumerate or reinforce existing sequences [7], [64], [74], or draw attention to less obvious sequences [5], [74] that convey important trends across time.

2) **Spotlight**: Spotlighting strategies focus or highlight specific information, most often as observation cues. Lone **encoding** variations and **annotations** emphasize and distinguish particular data [7], [66], [74]. This emphasis can help a reader recognize anomalies [66], values of significance [10], [74], missing values (e.g., election data [54]), or small / occluded items that orient a reader (e.g., mouse cursor [66]). For instance, Heer et al. [73] include an example of a reader using a crowdsourced system who draws a line (annotation) mimicking the data distribution, helping other readers recognize that the data distribution is being referenced. An annotated




Fig. 6. Example augmentation tactics to **interpret**. (A) Annotated trend line (dotted blue line), national average (dashed red line), and student performance relative to peers (green circle and red strike-through glyphs, with a legend provided in the bottom-left) to provide context about student's test scores [7]; (B) linked red / blue hue encoding to relate visualizations and text during an Arizona gubernatorial race and associate with political parties [76]; (C) use of repeated symbols to both relate consumer metrics after the 2008 recession (each row) and help associate the subject matter [3].

question mark is then included to further clarify that the distribution confuses the reader.

Linked or repeated **annotations** and **encodings** can emphasize groups of items to more readily make comparisons [9], [10], [45], [72] and highlight data trends [8], [44]. These augmentations also extend to different viewing contexts for the reader. For interactive content driven by user input, linked annotations (e.g., underlined text [16]) may facilitate observation of dynamic updates within the article. In addition, humans can dynamically **annotate** during presentations by gesturing with their hands [66], such as lassoeing with one's finger to attend to a certain grouping (circle annotation) or using spread fingers to indicate the range of data (interval annotation).

While less common, creative uses of **animation** can function as a filtering / sliding window [66], zoom-in lens [57], or motion separator [62], reducing information overload and revealing data outliers more easily [66]. **Transformations** such as filtering [9], [27] and threshold annotations [7] can graphically aid observation of irregularities with data distributions [66]. Aggregation or drill-down **transformations** may also highlight fine-grained details more clearly, such as how individual data items contribute to the whole [10]. Finally, juxtaposed cut-outs of visualizations regions (**layout**) can emphasize [49] important content.

Interpret Even when augmentations are used to facilitate observation, readers may lack the appropriate background knowledge to understand the context and message of a visualization [74]. To address this issue, *Interpretation* aids are a class of augmentations designed to help the reader understand what the visualization means (i.e., the semantics), such as relationships between data or multiple charts, thematic content, or external information important for context. The following three design strategies aid reader interpretation:

1) **Relate**: Relate strategies connect and link content to more easily understand conceptual relationships or convey semantic similarity. Sequence tactics can directly support the relation strategy for readers to more easily comprehend difficult concepts or transitions. A primary example is the use of **animation** to follow difficult transitions or maintain orientation, helping readers understand the relationship and connection between disparate views [26], [51], [57], [63],

[64], [66], [67], [68], [69]. Further, the sequential **layout** of content can convey cause-and-effect [34], [74] relationships or narrative events [49], [61], [62], [68], [72] more effectively, providing benefits to learning [81]. Sequential content can also relate the overall structure and organization of a visualization or story, enabling readers to comprehend the intent or message more easily while also enhancing aesthetics and clarity [45]. For example, parallel layouts [12], consistent templates (e.g., general-to-specific transitions within a narrative story [30], [48]), and navigation bars [30], [64] can orient readers within the narrative and help them form connections between views.

More broadly, relate augmentations can help readers understand exact relationships, classifications, or meanings between parts of a visualization (e.g., groups of text or marks), affecting the message that is conveyed. For instance, **encodings** (e.g., partial coloring [10]), **annotations** (e.g., glyphs [7]), motion (e.g., marks with different **animated** movements [69] that exploit the Gestalt Law of Common Fate [84]), repeated use of **symbols** (Figure 6(C)), or **layout** (e.g., clustering) can link similar groups of items [16], [33], [45], [71], [72]. In addition, certain annotations like arrows or connectors may not only facilitate observation of specific chart data, but also link contextual descriptions to the data [74], [81] (Figure 6(A)).

Although linked **encodings** help delineate categories in standard graphical perception [2], they may also link the semantics (e.g., linked content can indicate reactive updates, interactive behavior [16], [34], [53], [71], or dynamic annotations for details-on-demand [34], [75]). For example, judicious use of underlined text (**annotation**) in interactive documents [16], [32], [33], [72] not only serves to draw attention to updated content, but also for readers to interpret such text as linked and interactive. This augmentation can then function as a spotlighting mechanism for the reader to search for similar content later on (as with text highlighting [59]).

2) **Associate**: Association strategies ground content with semantic or connotative referents, facilitating interpretation [13], [14], [33], [34], [62] if the reader has sufficient background knowledge [74]. For instance, a guitar **symbol** [33] readily associates a narrative with musical content, which can serve as a future memory aid [13], [14]. The timing of **animated** transitions [57] and **encoded** angles and colors of




Fig. 7. Example augmentation tactics to **interpret** (A) and **introspect** (B,C). (A) Personalized content as guides, which displays income mobility data for the state that the reader accesses the article from [25]; (B) use of quizzes (within an article about spaced repetition) to query reader knowledge and promote reflection [34]; (C) self-predicting the trend of CO₂ emissions to encourage reflection [34].

annotations [73] can connote emotion [57]. **Encoding** channels and **animation** can also communicate uncertainty, such as blurred motion [62] and disordered stacking [47]. Further, designers can employ colors signifying political parties [76] or errors [52] and use **symbolic** representations of missing data (e.g., question marks or disjointed lines [54]).

3) **Contextualize**: Context strategies provide additional information to deepen the reader’s understanding. **Annotations** are widely used as contextual aids to provide details-on-demand [30], [32], [34], [68], [75] or include external information necessary for interpretation [7]. Annotations can be used in tandem with spotlighting strategies to contextualize specific data [73]. **Guides** also contribute important context, such as captions and titles [30], [52], legends [13], footnotes [72], personalized text [25], [34], [72], [80], [81] (Figure 7(A)), progress bars [30], [36], [76], and external links [73].

Content tutorials provide useful context for the reader prior to engaging with main content. Initial **layouts** can provide tacit tutorials [30], [73], while **animating** specific interactions may help readers understand their functionality [57], [73]. Spatial **layouts** can also provide overview+detail context [49]. Further, sequential illustrations may be used to provide additional information not readily comprehended from text alone [49], [61], [78].

Introspect Observation and interpretation help designers tailor the message of their visualization, but do not motivate readers to examine or repair their own assumptions, biases, knowledge, or meta-knowledge (i.e., the process by which they acquired that knowledge [52]). *Introspection* aids are a class of augmentations that facilitate self-learning [52], hypothesis generation and testing to explore one’s understanding [16], and repair of faulty knowledge [52], [82]. The following two strategies can refine readers’ prior beliefs and knowledge:

1) **Reflect**: Reflect strategies prompt self-examination of one’s assumptions to apply their understanding and generate inferences. Graphical interventions drive readers to **predict** partially unseen data, encouraging them to construct external representations of their prior knowledge and engage more deeply with understanding the data [34], [52] (Figure 7(C)). For instance, Kim et al. [52] asked readers to predict the change in voting percentage of Hispanic voters between income brackets by drawing a line; percentages for White and Black voters were provided as context. Readers can also generate textual (self-) **explanations** [52], which may prompt

them to reflect on their prior knowledge to generate inferences that fill in gaps [38]. For explaining complex processes, self-illustration may be used instead [78].

Other types of visual augmentations that encourage active reading, such as explorable explanations [16], [32], [33], [34], [68], [75] (Figure 3), can prompt (self-) **prediction** in which readers generate hypotheses about the data that they test by modifying system parameters via dynamic (interactive) widgets. Assessment **queries** (e.g., quizzes as in Figure 6(B) or flashcards [34], [82]) can directly prompt reflection, and may use a sequencing strategy (e.g., interspersing quizzes throughout data narratives [34]) to promote learning [82].

2) **Repair**: Repair strategies provide feedback for readers to observe knowledge failures, encouraging self-learning and potentially critiquing meta-knowledge [52]. We observed three primary tactics to facilitate self-repair. **Evaluate** tactics *directly* assess the correctness and quality of readers’ elicited beliefs / knowledge from the reflection stage. For example, self-prediction techniques help readers repair their knowledge by providing the correct result [52]. Graphical feedback may utilize observation and interpretation aids to draw readers’ attention to the literal gap in their knowledge (i.e., magnitude of the error in their prediction [52]). Visualization games [36] and quizzes / flashcards [34], [36] can also evaluate reader input (e.g., game win / loss, correct result shown after quiz).

Interactive content, such as explorable explanations (Figure 3), may provide a similar feedback loop to evaluation by **simulating** dynamic content with interactive input, rather than assessing the correctness of that input. Specifically, readers may generate predictions and hypotheses, modify the system state via interactive widgets, and use dynamically executed results to infer relationships and help repair faulty knowledge and assumptions.

Prior work [52] has shown that (self-) **explanation** [52] can improve data recall, which may be explained by the mental model repair hypothesis [38]: *quality* self-explanations (which may be dependent on adequate prior knowledge) prompt readers to generate inferences that fill in gaps in their prior knowledge, and *repair* faulty knowledge during this process.

D. Augmentation Properties

1) **Concurrency**: We observed that certain tactics could achieve multiple strategies or objectives *simultaneously*, as shown by tactics being shared between strategies in Figure 4, though only for **observation** and **interpretation**. After

completing the coding process, we went through each tactic, reasoning about cases in which a tactic could support multiple strategies in tandem. We describe a few illustrative cases here.

Animation can be used as a moving time window, both **sequencing** and **spotlighting** content to more easily **observe** errors [66]. Dynamic widgets may be used as **interpretation** aids and (self-) **prediction** mechanisms (e.g., using toggle to display contextual annotations in a math formula to check understanding [34]). An **annotation** can both provide **context** (aiding **interpretation**) and **spotlight** specific portions (**observing** data characteristics) [7] (e.g., a single text annotation next to a peak in a distribution can both “point out” the peak and provide external context to explain its occurrence). In the same scenario, an arrow **annotation** could both point out the peak and **relate** it to the text annotation.

2) **Hierarchy**: Our survey revealed an implicit hierarchical dependency among the objectives (though not for strategies or tactics within each objective), as illustrated by the graduated arrows spanning from **observe** to **introspect** in Figure 4. The objective hierarchy represents the increasing cognitive complexity on the part of the viewer, with simpler objectives often invoked to satisfy more complex ones. For example, **observation** aids serve as “first-order” perceptual cues to highlight salient aspects of the data. However, **interpretation** aids may be needed to contextualize such insights (e.g., Figure 8(B,C)) or communicate the message more viscerally (e.g., Figure 8(F,G)). When **introspection** is desired, lower-level objectives can provide essential context and direct the reader’s attention to their mistakes (e.g., Figure 10(E,F)). This dependency is reflected in our design space coding (see Table A1 in Appendix): individual objectives, notably **observe** and **introspect**, are less common, while **observe** and **interpret** co-occur together frequently and with **introspect**.

One explanation for this hierarchy is concurrency: use of certain tactics, such as **encode** or **annotate**, may be used primarily as **interpretation** aids, but serve as lower-level **observation** aids by default when used appropriately (e.g., using a magma color gradient in Figure 8(D,E) to **associate** with temperature also **sequences** the data by default, though a different color gradient could have been used).

3) **Containment**: Increasing complexity within the objective hierarchy results in overlap between certain tactics and strategies from different tiers. Specifically, augmentations classed within an individual objective (e.g., **symbolize** for **interpret** or **reflect** for **introspect**) may *contain* a tactic / strategy from a lower-level objective within its implementation. For instance, a textual **guide** (e.g., captions or titles) also **annotates** the chart; **symbols** may **encode** data (e.g., Figure 2); and **sequencing** may be used as a sub-strategy to **relate** disparate views (e.g., parallel layout [12]).

Conceptually, containment functions similar to concurrency, such as when a textual description both **annotates** and **guides**, though at a level below tactics (i.e., implementation). However, we retain properties of containment in our design space to both ensure consistent levels of abstraction (i.e., connections only exist between different tiers) and represent the hierarchical nature of certain tactics and strategies themselves. For instance, **encode** and **symbolize** both connect

to **interpret** via **relate** and **associate** strategies, though do not share any connections through **observe**. This relationship accurately reflects the fact that, when containment occurs, any **sequencing** or **spotlighting** of data is an artifact of the **encoding** (i.e., *how* data values are mapped to **symbols**).

V. CROSSING THE CHASM:

GENERATING & ANALYZING COMMUNICATIVE GRAPHICS

Our design space helps express and ground communicative intent as a series of design goals, which can then be mapped via strategies to specific, actionable tactics (i.e., “crossing” the chasm from intent to implementation). This benefit not only supports formal specification of an augmented chart, but also contributes a principled framework which can assist with generating and searching over a space of different augmentations (e.g., by varying tactics / strategies for a chosen goal). Our design space also aids analysis and reasoning about affordances different charts may provide. For instance, differences in reader takeaways can ostensibly be mapped to structural variations in augmentation graphs (Figure 8(C,E,G)), surfacing communicative tradeoffs to designers (from the perspective of readers) that can help align design intent and reader learning. We demonstrate both generative and analytic properties in the following examples, along with other insights and observations that result from using different tactics.

A. Example: CO₂ Concentration

To illustrate the value of our design space, we selected an initial visualization of atmospheric CO₂ concentration from 1958-2020 (Figure 8(A)) absent augmentations (e.g., the trendline and axes are required to decode denotative chart semantics). We then explored several augmentations to generate the following design variations and reason about their differences: (Figure 8(B)) adding context to more clearly visualize and call attention to the exponential rate of change; (Figure 8(D)) cleverly cycling the trend to emphasize finer detail (i.e., zooming in), with sequential color hue to indicate warmer temperatures (originally authored by Gregor Aisch [85]); and (Figure 8(F)) using animation with ramping speed to convey CO₂ increase more viscerally (originally authored by The New York Times (NYT) [86]).

Although a non-linear increase is apparent in Figure 8(A), suppose that we want to help readers better *understand* how quickly CO₂ concentration is increasing (Figure 8(B,C)). We can describe this intent as *showing* the reader what a slower rate (e.g., linear) would look like (i.e., provide **context** for **interpretation**). Of the different tactics for **context**, **annotating** with external data (i.e., linear regression) best enables visual comparison between the linear and exponential trends to facilitate comprehension. If we also wanted readers to **observe** the increasing gap, we could **encode** (e.g., shade) the region between the two trend lines as a **spotlighting** technique. Given that the linear regression is only fit to an initial subset of data (1956-1970), we add a text **annotation** with this information as **context**, though **guides** (e.g., footnote, chart title, or caption) could achieve the same effect. We use an arrow **annotation** to **relate** this **contextual** description




Fig. 8. Application of our design space to a CO₂ emissions chart to generate new designs. (A) Original, non-augmented chart. (B, D, F) A variety of possible augmented versions with (C, E, G) associated augmentation graphs. Designers can map high-level intents to a set of actionable tactics via strategies of their choice, enabling search and exploration of augmented communicative visualizations.

with the regression, as well as **spotlight** when the fitted data ends. To further help readers **interpret** the regression's fit, we dot the remaining line as "imaginary" to **associate** it as predicted linear data. Finally, to direct the reader's attention (i.e., **observe**) to the linear trend – the primary focus of this chart – we **spotlight** it with a red hue (**encoding**). It is worth noting that **introspection** tactics, such as (self-) **prediction** or **simulating** different regression curves, could be used to encourage the reader to **reflect** on their prior expectations about CO₂ levels, given the importance and global urgency of the climate crisis (Figure 7(C)).

Expanding the trend in Figure 8(A) adds more detail, helping reveal both the exponential rate of change and the seasonal variation in level (Figure 8(D,E)). Given the inherent ordering of the temporal data, we can make clever use of **transformation** to cycle or "wrap" the data around the y-axis, helping the reader **observe** the over-decades increase. Although y-axis labels provide the time periods necessary for context, we can use **sequential encoding** to make the order of data more transparent. Many sequential color schemes might be appropriate, but the magma color **encoding** simultaneously exploits color-concept **associations** (i.e., bluish hues for colder temperatures and reddish ones for warmer) to convey the relationship (**interpret**) between increasing CO₂ levels and global temperature.

The variation in Figure 8(D) could be augmented further to help readers **interpret** this rate more viscerally as demonstrated in Figure 8(F,G). This visualization retains the cyclical **transform** from Figure 8(D) to help readers **observe**

annually increasing temperatures, while using **animation** to draw attention to this change over time. Given that periodicity may be less obvious compared with Figure 8(B,D), we can add a footnote (**guide**) to **contextualize** seasonal differences between May and August. We can also exploit the **speed** of the **animation** to both convey the rate of increase more directly and connote urgency (**association**). In particular, just as the actual CO₂ trend increases exponentially, so does the rate of animated motion that introduces each annual time slice in **sequence** (i.e., the animation starts out slowly, ramps up exponentially, then slows down towards the end to avoid an abrupt halt). As a result, the animation offers the reader a "roller coaster" experience to better **interpret** the alarming rate at which CO₂ levels are rising.

While sequential color **encoding** is also retained to help readers' **interpret sequence**, it is not required given the use of animation and labeled year at the center of the visualization. Further, the color hue in Figure 8(F) may negatively affect **interpretation** given the mismatch between color-concept **associations**. Specifically, increasing levels of CO₂ are encoded with darkening blue hues, despite this scheme's typical **association** with cooler temperatures (as in Figure 8(D)). While this augmentation is in line with standard graphical perception principles, our design space helpfully reveals how this encoding might conflict with higher-level **interpretation** objectives. One possible explanation for this discrepancy is that a red hue **encoding** is used at the end of the animation to instead **spotlight** CO₂ concentration surpassing 420 parts per million (enabling **association** with dangerous levels), which is




Fig. 9. Application of our design space to analyze augmentation differences between designs of the same income inequality data across US metro areas. (A) First augmented chart with (B) associated augmentation graph that prioritizes recognizing global trends; (D) second augmented chart with (E) associated augmentation graph that prioritizes recognizing change over time between two metro areas. Potential differences between reader takeaways (C, F) can be linked to augmentation differences, aiding design reasoning to align designer intent with reader takeaways.

referenced later in the article to compare with levels before the Industrial Revolution [86]. However, other encodings (e.g., bolded font) or annotations (e.g., arrow call-out) might be used instead.

B. Example: Income Inequality

To demonstrate the analytic affordances of our design space, we consider a visualization authored by Mike Bostock that snapshots the change in income inequality across major United States metros in 1980 versus 2015 [87] (Figure 9(A)). This visualization was adapted from a NYT publication [88], which uses animation instead of single static frame (Figure 9(D)).

The static visualization (Figure 9(A, B)) uses arrow and sequential color hue encoding as sequencing strategies for readers to more readily observe differences between 1980 and 2015 income inequality. Selective annotations for certain metro areas serve as both spotlighting emphasis to observe the significant rise in inequality among major cities and contextual aids to help readers interpret the data. Finally, the visualization uses red hue (encoding) as an association for more severe inequality, given that the overwhelming majority of metro areas worsened in this respect. This visualization also makes subtle use of augmentation concurrency with both tactics by using color encoding to achieve multiple strategies simultaneously (associate and sequence) across different objectives (interpret and observe), as well as annotations (to both spotlight and contextualize).

The NYT visualization (Figure 9(D, E)) uses animated transitions to sequence data points between 1980 and 2015 frames and spotlights (bold font) Binghamton and New York to draw readers' attention (observe) to the fact that New York's inequality worsened significantly more than Binghamton's, which remained relatively stagnant. Contextual annotations for other major metro areas are also provided to aid interpretation. Guides (titles, captions, links, footnotes with data sources) also provide context.

When comparing potential differences between reader takeaways from these two visualizations (Figure 9(C, F)), we can reason from structural divergences in the augmentation graphs (Figure 9(B, E)) to expose such differences more concretely and transparently. The graphical relationships between objectives and strategies are structurally the same, aside from using an association strategy to aid communication in static form (Figure 9(B)). As a result, we can expect that differences in takeaways likely result from which tactics are employed. For instance, readers of both visualizations might notice that the majority of US metro areas had a significant increase in inequality due to shared use of sequencing. The NYT's animated representation reveals that Binghamton's inequality not only increased substantially less than New York's, but also substantially less relative to most metro areas. However, Bostock's visualization may bias readers to focus only on the significant increase in inequality among major metro areas (the annotated points). Finally, the NYT's spotlighting of New York and Binghamton (bold font encoding) may bias




Fig. 10. Application of our design space to different visualizations that encourage active learning. Augmented visualizations (A, C, E) are shown with associated augmentation graphs (B, D, F). Each visualization uses the same strategies for introspection (*reflect* and *repair* via self-prediction), but with varying tactics for other objectives (details-on-demand (A), game-playing (C) and graphical elicitation (E)), suggesting that learning differences may be attributed to supporting tactics and strategies which aid observation and interpretation, or how readers are prompted to engage in learning.

readers to primarily focus on the wide variance between these two areas, while Bostock’s global **encoding** representation emphasizes the wide variance across all metro areas as a result of significant mark overlap.

C. Example: Self-prediction

As demonstrated with prior examples, a central benefit to our design space is the ability to reason more precisely and descriptively about differences between augmented visualizations. On the other hand, we can also identify similarities between charts that differ substantially in either form or presentation. For example, the following three visualizations encourage active learning with interactive charts spanning different contexts: (Figure 10(A)) self-directed tool for learning the telltale signs in machine-generated images [34]; (Figure 10(C)) a visualization game for guessing binge drinking percentage rates across different US counties [36]; and (Figure 10(E)) an interactive visualization for predicting rates of Hispanic voters by income bracket, with other voting statistics provided

as context [56]. In Figure 10(A), readers may hover over **annotated** regions to view descriptions of machine-generated artifacts; multiple examples are provided in a faceted **layout** at the bottom right corner for the reader to explore. To guess health statistics in Figure 10(C), readers can change the slider threshold to (re-) **encode** county data above / below the new slider value. Readers (self-) **predict** data in Figure 10(E) by drawing a line with their mouse.

Despite dissimilar presentation (e.g., interactive features), data domains, and viewing context (e.g., interactive article that includes additional text and material), our design space (Figure 10(B,D,F)) reveals both important similarities and differences. In particular, each of the visualizations use the same tactics and strategies to **introspect** on the data. In each case, the reader is encouraged to **reflect** by **predicting** something about the data that is not displayed in the default view: machine-generated artifact (Figure 10(A)); health statistic (Figure 10(C)); and voting trend (Figure 10(E)). To **repair** the reader’s prior knowledge, including assumptions

or (mis)understanding of the data, reader **predictions** are **evaluated** via feedback. Such feedback includes **annotated** explanations (Figure 10(B)), a **guide** that displays the running total of the reader’s score (Figure 10(D)), and **annotated** text and arrows with red color **encoding** to **spotlight** (Figure 10(F)) any **prediction** errors.

Despite sharing **introspection** aids, the “lower-level” tactics and strategies which facilitate **observation** and **interpretation** diverge in certain respects (e.g., only Figure 10(A) **sequences** deliberately). As a result, these “lower-level” cues may affect self-learning directly, “guiding” the reader through the process of understanding their mistakes. For example, Figure 10(A,B) utilizes **layout** techniques (i.e., details-on-demand, which hide content behind interaction, and faceting) to provide additional **observations** and help **relate** the types of machine-generated errors to facilitate learning. However, by **spotlighting** these mistakes in the default view (via circle **annotations**), is the reader able to recognize machine-generated images in the future when augmentations are absent? By comparison, the other two examples (Figure 10(C,D)) and Figure 10(E,F)) also augment with **context** and **relation** strategies, via **encoding** channels. Figure 10(C) also provides readers with a dynamic threshold (**transform**) that (re-) **encodes** data to **observe** geographical correlations more readily, which could benefit learning.

Using the descriptive and systematic reasoning affordances of the design space, this analysis suggests an interesting possibility. Do learning benefits generalize to similar **introspection** aids used in different contexts? For instance, self-prediction improves recall and data comprehension with graphical elicitation techniques [52]; does this hold for other **prediction** techniques (e.g., details-on-demand across multiple instances in Figure 10(A) and guessing games in Figure 10(C))? If not, can learning differences be attributed to divergent **observation** and **interpret** aids? Future work should explore this relationship more deeply to identify the exact process of self-learning under these conditions, which we discuss in section VI.

VI. BEYOND THE CHASM: DISCUSSION & FUTURE WORK

A. Contribution & Limitations

Our primary contribution is a design space for augmentations in communicative visualization, which consists of *tactics* (the implemented design features), *strategies* (the approaches the designer selects to achieve an objective), and *objectives* (how the reader is affected perceptually or cognitively). This design space formally bridges the gap between intent and design, enabling designers to map their goals via strategies to a concrete set of actionable tactics.

We frame communicative / designer intent as explicit cognitive behaviors on behalf of the viewer. This framing synthesizes perceptual and cognitive visualization goals from prior work to provide a useful level of abstraction for deconstructing the different roles and uses of varied augmentations. Unlike prior objectives, our viewer-oriented goals helpfully relate an expressive range of augmentation tactics (e.g., animation, annotation, self-prediction, etc.) across different contexts (e.g., narrative visualization, data comics, dashboards). As a result,

our design space is able to capture more nuanced relations and dependencies (e.g., concurrency and hierarchy) at multiple levels (tactics, strategies, objectives) that help analyze and evaluate communicative aims more precisely. Designers can reflect on the way they are using augmentations, how these decisions reinforce their high-level goals, and communicative tradeoffs between different designs. Our design space may offer support for exploring alternatives (Figure 8), either with respect to different approaches to achieve the same goals, or modifying the augmentation to support alternative goals, which we aim to validate in future work.

We present the hierarchical design space as a color-coded graph and apply it to a rich set of examples to illustrate its descriptive and analytic power (Figure 8, Figure 9, Figure 10). By grounding objectives in tactics, we provide a principled framework to help generate and search over a space of augmentations (Figure 8). Designers can formulate their high-level intents as a series of viewer-oriented objectives: what do they want the viewer to pay attention to, interpret, or reflect on? Each objective can then be mapped to a set of strategies and tactics, varied based on the data and context of use. For example, **annotate** and **transform** tactics can both **spotlight** data, though the designer’s choice depends on *what* they wish for the viewer to **observe** (e.g., a specific point along a trendline Figure 8(B) vs. trend increase over time Figure 8(D)). This process is further aided by analyzing structural divergences between augmentation graphs (Figure 9), encouraging designers to reflect on potential shifts in viewer behavior and reason about communication tradeoffs. Specifically, by adding, removing, or modifying an augmentation, what information and takeaways will the viewer likely prioritize over others?

Appropriate use of augmentations can also be navigated by first selecting individual tactics, then noting potential concurrent and hierarchical dependencies. For instance, our critique of different color **encodings** in Figure 8(D,F) stems from the concurrent relationship between **sequence** and **associate**. Figure 8(D) exploits the magma color **association** with rising temperatures, while Figure 8(F)’s blue gradient conflicts with these semantics. Our design space hierarchy facilitates this line of reasoning: tactics used to first **observe**, such as **encode** or **annotate** (Figure 9(A)), can be matched next against concurrent strategies for **interpret** to identify unintended effects. In such cases, the designer can select a different implementation for the same tactic (e.g., color scheme), an alternative tactic that supports the same strategy or objective, or use the proper implementation based on concurrent strategies and objectives.

A major advantage of our design space is that it provides a consistent lexicon for describing complex communicative visualizations that may appear very different (e.g., different tactics) but have similar goals (Figure 10). For instance, our design space helped relate explorable explanations [16] with other uses of similar augmentations from visualization games [36] or graphical self-prediction [52], despite significant differences in form and presentation. These connections could suggest that similar methodologies and assessments might be used to investigate actual learning benefits (which, for example, remains an open question for interactive articles [32]).

Despite promise, some augmentations, such as self-reflection and repair (Figure 10), are rarely used in general practice, likely due to the design and/or engineering knowledge necessary to implement them. Future (semi-) automated authoring systems may provide some relief on this front, leveraging our design space to evaluate or generate designs. Going forward, we plan to also continue evaluating our design space through discussions with expert designers themselves and more tightly integrating with learning objectives.

B. Scope of Augmentations

We expect our design goals to be sufficiently high-level to capture more abstract intents, while also providing enough grounding to help designers identify actionable strategies and tactics that can facilitate design iteration and reason about readers’ takeaways. However, future work is needed to test our design space with various designers across different domains, which may reveal gaps between how they express and organize their intents and how the literature captures communicative intent. Naturally, the intent may shift with varying communication contexts (e.g., didactic, rhetorical / persuasive, affective, etc.). In these situations, we expect that our design space may function as “middle-ware” in which the augmentation objectives uniquely map to the higher-level “context” objectives. This connective tissue might even serve as an additional resource for mapping between “context” objectives (e.g., how do learning objectives map to rhetorical aims via visual augmentations), or even connecting with lower-level tasks throughout perceptual [9] and analysis [27] taxonomies. As a result, designers may have more flexibility in generating designs that can be adapted across multiple contexts according to their intent.

Given the existence of augmentations, we expect there to be anti-augmentations (akin to anti-learning objectives as noted by Adar and Lee-Robbins [5]), which serve to distract or bias readers or impede learning. For example, removing clutter (e.g., gridlines, unhelpful annotations, etc.) [45]. Anti-augmentations may also entail removing particular types of content to aid observation (e.g., filtering) or interpretation, which has been observed with decluttering visualizations [45] and, in certain cases, animation [13].

C. Relation to Learning Objectives

We began this work particularly inspired by Adar and Lee-Robbins’ extended learning objectives taxonomy for visualization [5]. We view our characterization of augmentations and the resulting design space as a step towards evaluating and reasoning about learning objectives, allowing designers to use the strategies and tactics we have surveyed to achieve different learning outcomes. For instance, Adar et al. [5] operationalize learning objectives as statements of the form: **The viewer will [verb] [noun]**. This construction delineates cognitive constructs (verbs) (e.g., RECALL, CLASSIFY, IMPLEMENT, etc.) and knowledge dimensions (nouns) (e.g., factual knowledge like terminology or conceptual knowledge like categories). We might leverage **observation** mechanisms (i.e., **spotlighting** or

sequencing) to convey or focus alternative knowledge representations (nouns); **interpretation** and **introspection** aids can then instrument the desired learning processes (verbs). However, the exact connection between our design space and learning objectives, as supported by empirical evaluation, remains as future work, in part due to the difficulties and training that can be required for designers to formulate their intents as learning objectives [5].

D. Extensions to Augmentations & Authoring Tools

We contribute a systematic outline of the visual augmentation space that can serve as a springboard for future empirical work and communication tools. We hope to use our design space to ground empirical assessment and treatments of various augmentations, including their relationships with domain knowledge, other learning objectives (e.g., rhetorical or affective), and augmentations in different contexts (e.g., math formulas [46]), to assess ecological validity. Although trade-offs and opportunities for learning outcomes must be evaluated, visual augmentations may provide other benefits such as engagement or joy [27].

Several exciting directions for future work revolve around using our design space as the basis for a design language, potentially even a computational domain-specific language (DSL). For instance, do the tiers of tactics, strategies, and objectives serve as useful “building blocks” or composable primitives for a grammar of communicative graphics? The generative and analytic capabilities of our design space might help scaffold future authoring systems, enabling designers to create communicative visualizations more rapidly. Such capabilities may also be used to encode knowledge in DSLs for different domains, or more generally leverage this knowledge as input to AI systems (e.g., using language models to author data stories [89]) to identify possible augmentations during the design process.

The space of possible augmentations is expansive, encompassing non-visual channels (e.g., narration or text [81], [90], [91]), assumptions about readers (e.g., visual literacy and background knowledge), different device types (e.g., mobile vs. desktop [92]), and context of use (collaboration, dashboards, teaching, business presentations, etc.). Our design space examines a particular subset of augmentations (visual) within a specific context (data communication & learning), which we hope to help scaffold future work across other modalities (e.g., multimedia voice narration [81]), types of learning (e.g., affective [4] or social [56]), and contexts.

REFERENCES

- [1] J. Bertin, *Semiology of graphics*. University of Wisconsin Press, 1983.
- [2] J. Mackinlay, “Automating the design of graphical presentations of relational information,” *ACM Transactions on Graphics*, vol. 5, no. 2, pp. 110–141, 1986.
- [3] J. Hullman and N. Diakopoulos, “Visualization rhetoric: Framing effects in narrative visualization,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 12, pp. 2231–2240, 2011.
- [4] E. Lee-Robbins and E. Adar, “Affective learning objectives for communicative visualizations,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 1, pp. 1–11, 2022.

- [5] E. Adar and E. Lee, "Communicative visualizations as a learning problem," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 946–956, 2020.
- [6] E. Lee-Robbins, S. He, and E. Adar, "Learning objectives, insights, and assessments: How specification formats impact design," *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 1, pp. 676–685, 2021.
- [7] M. D. Rahman, G. J. Quadri, B. Doppalapudi, D. A. Szafir, and P. Rosen, "A qualitative analysis of common practices in annotations: A taxonomy and design space," *IEEE Transactions on Visualization and Computer Graphics*, 2024.
- [8] A. Srinivasan, S. M. Drucker, A. Endert, and J. Stasko, "Augmenting visualizations with interactive data facts to facilitate interpretation and communication," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 672–681, 2018.
- [9] G. J. Quadri and P. Rosen, "A survey of perception-based visualization studies by task," *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 12, pp. 5026–5048, 2021.
- [10] G. Guo, J. Stasko, and A. Endert, "What we augment when we augment visualizations: A design elicitation study of how we visually express data relationships," in *Proceedings of the 2024 International Conference on Advanced Visual Interfaces*, 2024, pp. 1–5.
- [11] "Trump broke the democrats' thermostat," <https://www.ft.com/content/73a1836d-0faa-4c84-b973-554e2ca3a227>, 2024.
- [12] E. R. Tufte and D. Robins, *Visual explanations*. Graphics Cheshire, CT, 1997.
- [13] J. Hullman, E. Adar, and P. Shah, "Benefitting infovis with visual difficulties," *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 12, pp. 2213–2222, 2011.
- [14] 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*, 2010, pp. 2573–2582.
- [15] B. Victor, "Up and down the ladder of abstraction," 2011.
- [16] —, "Explorable explanations," 2011.
- [17] W. Playfair, *The commercial and political atlas: Representing, by means of stained copper-plate charts, the progress of the commerce, revenues, expenditure and debts of england during the whole of the eighteenth century*. T. Burton, 1801. [Online]. Available: <https://books.google.com/books?id=dgRdAAAACAAJ>
- [18] S. Pinker, "A theory of graph comprehension," 1990.
- [19] S. M. Kosslyn, "Understanding charts and graphs," *Applied Cognitive Psychology*, vol. 3, no. 3, pp. 185–225, 1989.
- [20] W. S. Cleveland and R. McGill, "Graphical perception: Theory, experimentation, and application to the development of graphical methods," *Journal of the American Statistical Association*, vol. 79, no. 387, pp. 531–554, 1984.
- [21] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer, "Vega-lite: A grammar of interactive graphics," *IEEE transactions on visualization and computer graphics*, vol. 23, no. 1, pp. 341–350, 2016.
- [22] M. Bostock, V. Ogievetsky, and J. Heer, "D³ data-driven documents," *IEEE transactions on visualization and computer graphics*, vol. 17, no. 12, pp. 2301–2309, 2011.
- [23] R. Fygenon, S. Franconeri, and E. Bertini, "The arrangement of marks impacts afforded messages: Ordering, partitioning, spacing, and coloring in bar charts," *IEEE Transactions on Visualization and Computer Graphics*, vol. 30, no. 1, pp. 1008–1018, 2023.
- [24] C. Xiong, V. Setlur, B. Bach, E. Koh, K. Lin, and S. Franconeri, "Visual arrangements of bar charts influence comparisons in viewer takeaways," *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 1, pp. 955–965, 2021.
- [25] E. Adar, C. Gearig, A. Balasubramanian, and J. Hullman, "Persalog: Personalization of news article content," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017, pp. 3188–3200.
- [26] B. Tversky, J. B. Morrison, and M. Betrancourt, "Animation: can it facilitate?" *International Journal of Human-Computer Studies*, vol. 57, no. 4, pp. 247–262, 2002.
- [27] M. Brehmer and T. Munzner, "A multi-level typology of abstract visualization tasks," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2376–2385, 2013.
- [28] J. Heer, F. B. Viégas, and M. Wattenberg, "Voyagers and voyeurs: supporting asynchronous collaborative information visualization," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2007, pp. 1029–1038.
- [29] S. Nobarany, M. Haraty, and B. Fisher, "Facilitating the reuse process in distributed collaboration: a distributed cognition approach," in *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, 2012, pp. 1223–1232.
- [30] E. Segel and J. Heer, "Narrative visualization: Telling stories with data," *IEEE transactions on Visualization and Computer Graphics*, vol. 16, no. 6, pp. 1139–1148, 2010.
- [31] A. Sarikaya, M. Correll, L. Bartram, M. Tory, and D. Fisher, "What do we talk about when we talk about dashboards?" *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 682–692, 2018.
- [32] M. Conlen and J. Heer, "Idyll: A markup language for authoring and publishing interactive articles on the web," in *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, 2018, pp. 977–989.
- [33] M. Conlen, A. Kale, and J. Heer, "Capture & analysis of active reading behaviors for interactive articles on the web," in *Computer Graphics Forum*, vol. 38, no. 3. Wiley Online Library, 2019, pp. 687–698.
- [34] F. Hohman, M. Conlen, J. Heer, and D. H. P. Chau, "Communicating with interactive articles," *Distill*, vol. 5, no. 9, p. e28, 2020.
- [35] J. Heer, M. Conlen, V. Devireddy, T. Nguyen, and J. Horowitz, "Living papers: A language toolkit for augmented scholarly communication," in *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, 2023, pp. 1–13.
- [36] N. Diakopoulos, F. Kivran-Swaine, and M. Naaman, "Playable data: Characterizing the design space of game-y infographics," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2011, pp. 1717–1726.
- [37] E. Tufte, "Data-ink maximization and graphical design," *Oikos*, pp. 130–144, 1990.
- [38] M. T. Chi, "Self-explaining expository texts: The dual processes of generating inferences and repairing mental models," in *Advances in Instructional Psychology, Volume 5*. Routledge, 2013, pp. 161–238.
- [39] S. K. Card, J. Mackinlay, and B. Shneiderman, *Readings in information visualization: using vision to think*. Morgan Kaufmann, 1999.
- [40] P. Saraiya, C. North, and K. Duca, "An insight-based methodology for evaluating bioinformatics visualizations," *IEEE Transactions on Visualization and Computer Graphics*, vol. 11, no. 4, pp. 443–456, 2005.
- [41] P. Saraiya, C. North, V. Lam, and K. A. Duca, "An insight-based longitudinal study of visual analytics," *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 6, pp. 1511–1522, 2006.
- [42] C. Plaisant, J.-D. Fekete, and G. Grinstein, "Promoting insight-based evaluation of visualizations: From contest to benchmark repository," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 1, pp. 120–134, 2007.
- [43] C. Stokes, C. X. Bearfield, and M. A. Hearst, "The role of text in visualizations: How annotations shape perceptions of bias and influence predictions," *IEEE Transactions on Visualization and Computer Graphics*, vol. 30, no. 10, pp. 6787–6800, 2023.
- [44] G. J. Quadri, A. Z. Wang, Z. Wang, J. Adorno, P. Rosen, and D. A. Szafir, "Do you see what i see? a qualitative study eliciting high-level visualization comprehension," in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–26.
- [45] K. Ajani, E. Lee, C. Xiong, C. N. Knaflic, W. Kemper, and S. Franconeri, "Declutter and focus: Empirically evaluating design guidelines for effective data communication," *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 10, pp. 3351–3364, 2021.
- [46] A. Head, A. Xie, and M. A. Hearst, "Math augmentation: How authors enhance the readability of formulas using novel visual design practices," in *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 2022, pp. 1–18.
- [47] L. Padilla, M. Kay, and J. Hullman, "Uncertainty visualization," in *Balakrishnan N., Colton T., Everitt B., Piegorsch W., Ruggeri F., Teugels J. L. (Eds.). Wiley*, 2022.
- [48] J. Hullman, S. Drucker, N. H. Riche, B. Lee, D. Fisher, and E. Adar, "A deeper understanding of sequence in narrative visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2406–2415, 2013.
- [49] B. Bach, Z. Wang, M. Farinella, D. Murray-Rust, and N. Henry Riche, "Design patterns for data comics," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–12.
- [50] N. E. Adams, "Bloom's taxonomy of cognitive learning objectives," *Journal of the Medical Library Association*, vol. 103, no. 3, p. 152, 2015.
- [51] R. Ploetzner, S. Berney, and M. Bétrancourt, "A review of learning demands in instructional animations: The educational effectiveness of

- animations unfolds if the features of change need to be learned,” *Journal of Computer Assisted Learning*, vol. 36, no. 6, pp. 838–860, 2020.
- [52] Y.-S. Kim, K. Reinecke, and J. Hullman, “Explaining the gap: Visualizing one’s predictions improves recall and comprehension of data,” in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017, pp. 1375–1386.
- [53] J. Boy, F. Detienne, and J.-D. Fekete, “Storytelling in information visualizations: Does it engage users to explore data?” in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2015, pp. 1449–1458.
- [54] L. W. Ge, M. Easterday, M. Kay, E. Dimara, P. Cheng, and S. L. Franconeri, “V-framer: Visualization framework for mitigating reasoning errors in public policy,” in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–15.
- [55] T. M. Haladyna, *Developing and validating multiple-choice test items*. Routledge, 2004.
- [56] Y.-S. Kim, K. Reinecke, and J. Hullman, “Data through others’ eyes: The impact of visualizing others’ expectations on visualization interpretation,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, no. 1, pp. 760–769, 2017.
- [57] F. Chevalier, N. H. Riche, C. Plaisant, A. Chalbi, and C. Hurter, “Animations 25 years later: New roles and opportunities,” in *Proceedings of the 2016 International Working Conference on Advanced Visual Interfaces*, 2016, pp. 280–287.
- [58] C. E. Shannon, “A mathematical theory of communication,” *The Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948.
- [59] N. Joshi and D. Vogel, “Constrained highlighting in a document reader can improve reading comprehension,” in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–10.
- [60] K. Chandra, K. M. Collins, W. Crichton, T. Chen, T.-M. Li, A. Weller, R. Nigam, J. Tenenbaum, and J. Ragan-Kelley, “Watchat: Explaining perplexing programs by debugging mental models,” *arXiv preprint arXiv:2403.05334*, 2024.
- [61] S. McCloud, *Understanding comics*. Kitchen Sink Press, Inc., 1993.
- [62] Y. Shi, X. Lan, J. Li, Z. Li, and N. Cao, “Communicating with motion: A design space for animated visual narratives in data videos,” in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–13.
- [63] J. Heer and G. Robertson, “Animated transitions in statistical data graphics,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1240–1247, 2007.
- [64] S. McKenna, N. Henry Riche, B. Lee, J. Boy, and M. Meyer, “Visual narrative flow: Exploring factors shaping data visualization story reading experiences,” in *Computer Graphics Forum*, vol. 36, no. 3. Wiley Online Library, 2017, pp. 377–387.
- [65] N. Sultanum, F. Chevalier, Z. Bylinskii, and Z. Liu, “Leveraging text-chart links to support authoring of data-driven articles with vizflow,” in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–17.
- [66] W. C. Hill and J. D. Hollan, “Deixis and the future of visualization excellence,” in *IEEE Visualization*, vol. 91, 1991, pp. 314–320.
- [67] Y. Kim and J. Heer, “Gemini: A grammar and recommender system for animated transitions in statistical graphics,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 485–494, 2020.
- [68] A. Satyanarayan and J. Heer, “Authoring narrative visualizations with ellipsis,” in *Computer Graphics Forum*, vol. 33, no. 3. Wiley Online Library, 2014, pp. 361–370.
- [69] T. Ge, B. Lee, and Y. Wang, “Cast: Authoring data-driven chart animations,” in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–15.
- [70] Y. Kim, K. Wongsuphasawat, J. Hullman, and J. Heer, “Graphscape: A model for automated reasoning about visualization similarity and sequencing,” in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017, pp. 2628–2638.
- [71] B. Victor, “Scientific communication as sequential art,” 2011.
- [72] —, “Magic ink: Information software and the graphical interface,” 2006.
- [73] J. Heer and B. Shneiderman, “Interactive dynamics for visual analysis: A taxonomy of tools that support the fluent and flexible use of visualizations,” *Queue*, vol. 10, no. 2, pp. 30–55, 2012.
- [74] R. C. Clark and C. Lyons, *Graphics for learning: Proven guidelines for planning, designing, and evaluating visuals in training materials*. John Wiley & Sons, 2010.
- [75] W. Crichton, “A new medium for communicating research on programming languages,” in *Workshop on Human Aspects of Types and Reasoning Assistants*, 2021.
- [76] M. Cai and M. Kay, “Watching the election sausage get made: How data journalists visualize the vote counting process in us elections,” in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–19.
- [77] P. Shah, R. E. Mayer, and M. Hegarty, “Graphs as aids to knowledge construction: Signaling techniques for guiding the process of graph comprehension,” *Journal of Educational Psychology*, vol. 91, no. 4, p. 690, 1999.
- [78] W. H. Levie and R. Lentz, “Effects of text illustrations: A review of research,” *Educational Technology Research and Development*, vol. 30, no. 4, pp. 195–232, 1982.
- [79] L. Yang, C. Xiong, J. K. Wong, A. Wu, and H. Qu, “Explaining with examples: Lessons learned from crowdsourced introductory description of information visualizations,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 3, pp. 1638–1650, 2023.
- [80] D. I. Cordova and M. R. Lepper, “Intrinsic motivation and the process of learning: Beneficial effects of contextualization, personalization, and choice,” *Journal of Educational Psychology*, vol. 88, no. 4, p. 715, 1996.
- [81] R. E. Mayer, “Cognitive theory of multimedia learning,” *The Cambridge Handbook of Multimedia Learning*, vol. 41, no. 1, pp. 31–48, 2005.
- [82] A. Matuschak and M. Nielsen, “How can we develop transformative tools for thought?” *Numinous Productions (blog)*, 2019.
- [83] J. S. Ancker, N. C. Benda, and B. J. Zikmund-Fisher, “Do you want to promote recall, perceptions, or behavior? the best data visualization depends on the communication goal,” *Journal of the American Medical Informatics Association*, vol. 31, no. 2, pp. 525–530, 2024.
- [84] K. Koffka, “Perception: an introduction to the gestalt-theorie,” *Psychological Bulletin*, vol. 19, no. 10, p. 531, 1922.
- [85] “Six decades of carbon dioxide concentration in the atmosphere,” https://www.datawrapper.de/_/OHgEm/.
- [86] “Carbon dioxide levels have passed a new milestone,” <https://www.nytimes.com/interactive/2024/04/20/upshot/carbon-dioxide-growth.html>, 2024.
- [87] “Inequality in american cities,” <https://observablehq.com/@d3/inequality-in-american-cities>.
- [88] “Watch 4 decades of inequality drive american cities apart,” <https://www.nytimes.com/2019/12/02/upshot/wealth-poverty-divide-american-cities.html>, 2019.
- [89] H. W. Wang, L. Birnbaum, and V. Setlur, “Jupyter: Operationalizing a design space for actionable data analysis and storytelling with llms,” in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–24.
- [90] C. Stokes, V. Setlur, B. Cogley, A. Satyanarayan, and M. A. Hearst, “Striking a balance: Reader takeaways and preferences when integrating text and charts,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 1, pp. 1233–1243, 2022.
- [91] A. Srinivasan, V. Setlur, and A. Satyanarayan, “Pluto: Authoring semantically aligned text and charts for data-driven communication,” in *Proceedings of the 30th International Conference on Intelligent User Interfaces*. Association for Computing Machinery, 2025, p. 1123–1140.
- [92] J. Hoffswell, W. Li, and Z. Liu, “Techniques for flexible responsive visualization design,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–13.

Luke S. Snyder is a Computer Science & Engineering Ph.D. student at the University of Washington, where he conducts research on data visualization and human-computer interaction. He holds a B.S. degree from the University of Arkansas and M.S. degrees from Purdue University and the University of Washington, all in Computer Science.

Maureen C. Stone joined the University of Washington as an affiliate professor after retiring from Tableau Research (2012–2022). While best known for her expertise in digital color and its application, she has a broad interest in data and information visualization, computer-human interaction and design. She holds B.S. and M.S. degrees in Computer Engineering from U of Illinois, and a M.S. in Computer Science from Caltech.

Jeffrey Heer is the Jerre D. Noe Endowed Professor of Computer Science & Engineering at the University of Washington, where he co-directs the Interactive Data Lab and conducts research on data visualization, data science, and human-computer interaction. He holds B.S., M.S., and Ph.D. degrees in Computer Science from UC Berkeley.