Premature Crisping: When Visualization Systems Make Vague Data Look Precise

RYAN SMITH, The University of British Columbia, Canada

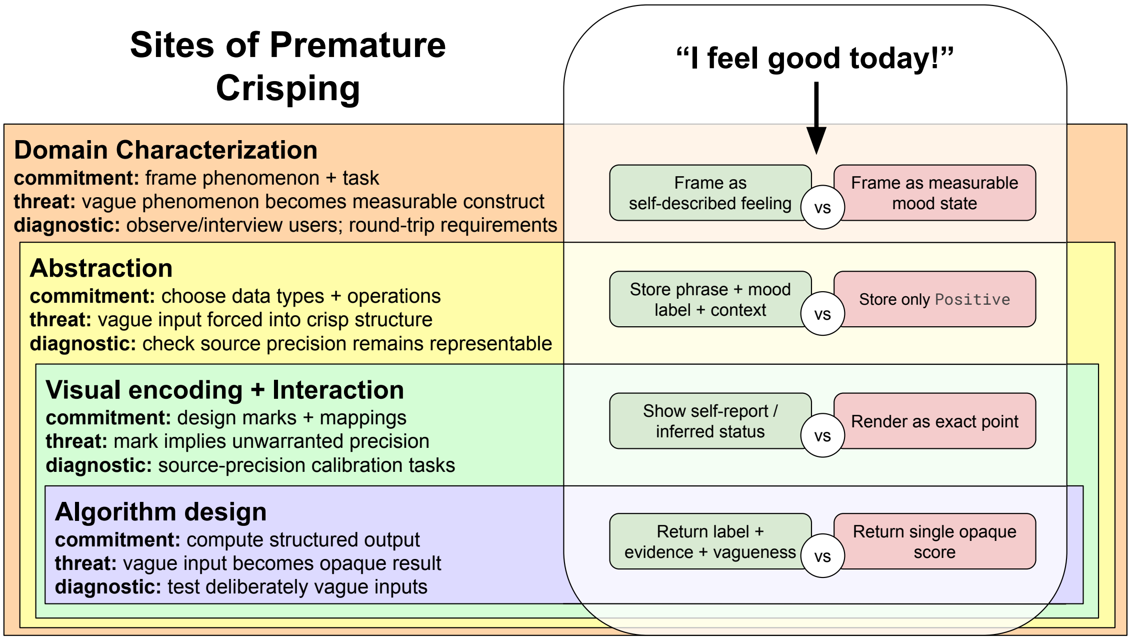

Fig. 1. Design sites where vague input can become prematurely crisp. Across domain characterization, abstraction, visual encoding and interaction, and algorithm design, designers decide how vague input becomes structured representation. Premature crisping occurs when those decisions silently impose precision and make it indistinguishable from precision provided by the source.

Visualization increasingly mediates between vague human expression and formal representation. Visualizing vague, language-derived data therefore poses a distinct design problem: systems must often impose structure without making their own interpretations appear source-provided. I identify premature crisping as a transversal validity threat in which semantic vagueness becomes crisp visual or computational structure while system-imposed precision becomes indistinguishable from source-provided precision. Using Munzner's nested model as an analytic scaffold, I locate the threat across domain characterization, abstraction, visual encoding and interaction, and algorithm design. I then develop an actor- and task-dependent account of interpretive recourse. An interpretive-recourse profile describes whether a particular actor can recognize that the system resolved a value, access the source vagueness, and contest or revise the stabilized interpretation. Through case analyses, I trace how premature crisping enters visualization pipelines and how available recourse changes across actors and tasks. I conclude with an audit for identifying where systems stabilize interpretations that the source did not provide.

Author's Contact Information: Ryan Smith, rsmith0914@gmail.com, The University of British Columbia, Vancouver, British Columbia, Canada.

Manuscript submitted to ACM

CCS Concepts: • **Human-centered computing → Visualization theory, concepts and paradigms**; **Visualization design and evaluation methods**.

Additional Key Words and Phrases: Vagueness, Visualization Design, Information Visualization, Information Loss, Provenance, Natural Language Input

ACM Reference Format:
Ryan Smith. 2027. Premature Crisping: When Visualization Systems Make Vague Data Look Precise. In *Proceedings of Proc. ACM CHI 2027 (SIGCHI '27)*. ACM, New York, NY, USA, 23 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

A patient describes their pain as starting "sometime in the morning." A timeline visualization needs a temporal position. The system picks one. In that small act of translation, a vague expression becomes a precise-looking data point. Viewers may then read the selected position as information the patient provided rather than precision the system imposed.

I call this failure **premature crisping**: the silent resolution of semantic vagueness into a representation that appears more precise than the source warrants. Premature crisping threatens visualization validity because systems routinely translate human expression into structured data and then into visual marks, interactions, layouts, and algorithmic outputs. A commitment at one layer constrains what later layers can preserve. Once a system converts a vague temporal phrase into an exact time, downstream components may propagate the inferred value as though the patient supplied it directly.

Premature crisping arises within the broader problem of visualizing vague, language-derived data. People routinely describe experiences through gradable, hedged, context-dependent, and partly underspecified language. A survey respondent may report that service was 'kind of slow," while a patient may describe symptoms as 'pretty bad lately." Such expressions convey meaning through comparison, context, granularity, and use [34, 35]. Their usefulness often depends on leaving some boundaries unsettled. The appropriate degree of precision also varies across settings, actors, and tasks. Visualization systems must therefore make vague expression computable while preserving the relationship between the source expression and the interpretation imposed during formalization.

The stakes extend beyond technical fidelity. Critical InfoVis emphasizes that structured representations privilege particular interpretations and narrow others [2, 16, 17]. When a system resolves vague expression, it helps determine which meanings survive formalization and which interpretations acquire authority.

Natural-language and conversational interfaces make the problem increasingly pressing because they invite free-form descriptions that systems must translate into structurem but the same pressure has long appeared in systems that operationalize human experience through scales, categories, and inferred computational states [37, 56, 61]. Conversational interfaces increase the volume and variety of vague expression entering computational pipelines, making the underlying interpretive commitments easier to overlook.

I use Munzner's nested model as an analytic scaffold for tracing premature crisping across four layers: domain problem and data characterization, operation and data type abstraction, visual encoding and interaction, and algorithm design [48]. Designers may first frame a vague phenomenon as measurable. The abstraction can then make that framing computationally necessary, while the encoding and algorithm make the resulting precision appear source-provided.

Premature crisping also concerns what relevant actors can do with provenance. A pipeline may retain the original expression while foregrounding only the value derived from it. For example, a database may preserve "sometime in the morning" while a decision-maker sees an exact point on a timeline. The trace remains available somewhere in the system, yet the interface presents the imposed interpretation as authoritative. Recording a transformation therefore provides limited protection when actors cannot recognize its effects or respond to them.

I capture that distinction through **interpretive recourse**: whether relevant actors can identify what the system added and respond when the resulting interpretation does not fit. For a given actor and task, an **interpretive-recourse profile** describes three components: **recognition**, **access**, and **agency**. A trace becomes useful when the actor has enough recourse for the task and the consequences of misinterpretation.

The primary contributions of this paper are:

- **Premature crisping as a visualization validity threat.** I identify how systems silently resolve semantic vagueness into unwarranted precision and make that precision difficult to distinguish from source-provided information.
- **An actor- and task-dependent account of interpretive recourse.** I introduce the **interpretive-recourse profile** to characterize whether a particular actor can recognize an imposed interpretation, access the vague source, and contest or revise the stabilized meaning.

The secondary contributions of this paper are:

- **A scoped problem definition for visualizing vague, language-derived data.** I characterize the design problem within which premature crisping occurs and clarify how required precision varies across actors, tasks, and settings.
- **An audit for premature crisping.** I provide a procedure for tracing where systems impose precision, documenting the resulting commitments, and assessing whether available recourse fits the relevant actor, task, and consequences.

2 Related Work

Premature crisping draws together several research lineages that address different parts of the problem. Semantic vagueness explains why language may remain meaningful without exact boundaries. Uncertainty and ambiguity visualization provide techniques for depicting non-exactness after a representational object has been defined. Research on categorization and provenance traces how transformations stabilize meaning and records how those transformations occurred. Subjective, humanistic, critical, and feminist visualization establish why source meaning and interpretive authority matter during formalization. Munzner's nested model then provides the scaffold for locating where premature crisping enters a visualization pipeline.

2.1 Semantic Vagueness

Vagueness is a long-studied property of natural language. Foundational work characterizes vague predicates through borderline cases, where an expression neither clearly applies nor clearly fails to apply [59]. The Sorites paradox illustrates the problem: removing one grain of sand seems insufficient to turn a heap into a non-heap, yet no individual removal establishes a clear boundary [29]. Formal theories explain the indeterminacy in different ways. Fuzzy set approaches model graded category membership [35]. Supervaluationist accounts preserve multiple admissible precisifications [21], while epistemicist accounts treat vague boundaries as sharp but unknowable [64]. Across these accounts, vague meaning exceeds ordinary measurement error.

Linguistic work shows how vagueness enters everyday expression. Gradable adjectives such as 'severe," 'recent," and 'high" depend on comparison classes rather than fixed thresholds [34]. Hedges such as 'kind of" and 'mostly" alter how strongly a predicate applies [35]. Pragmatic halos allow even apparently exact expressions to carry contextual tolerance, so ‘at 3:00” may permit a few minutes of flexibility [36]. Granularity creates uneven specificity within the same account. A person may locate one event to the day, another to the year, and a third only as “a while ago” [27].

The precision required from an expression depends on the setting and purpose of communication. ‘Sometime in the morning” may provide enough detail for recounting a day while offering too little for coordinating medication. ‘Pretty bad lately” conveys comparison and temporal scope without determining a severity score or fixed interval. Visualization systems introduce an interpretive commitment when they translate such expressions into structured values. Premature crisping identifies cases where the resulting representation hides that commitment and presents its precision as source-provided.

2.2 Uncertainty, Ambiguity, and Visualization

Uncertainty visualization addresses a broad family of representational conditions. Established typologies distinguish uncertainty according to its source, form, and location within a data or visualization pipeline [39, 62]. In this paper, *statistical uncertainty* refers more narrowly to work that represents variability around an already-defined quantity or model. Common approaches communicate error, confidence, and plausible values through distributions, ensembles, intervals, and related encodings [7, 10]. Research has also traced how uncertainty enters through acquisition, modeling, transformation, and visual encoding [54].

Encoding choices shape how viewers interpret uncertain values. Prior work has examined alternatives to conventional error bars [15], representations of temporal uncertainty [24], and the practical pressures that lead authors to omit uncertainty from visualizations [28]. Cultural heritage systems face similar representational problems when records contain imprecise dates or locations. PolyCube communicates such uncertainty through coordinated views and a space-time cube [65]. Temporal-query and event-sequence systems also formalize intervals, absences, alignment, and event membership before users can inspect patterns [40, 45]. Across these examples, visualization depicts non-exactness after designers have established the object and attributes being represented.

Uncertainty also arises throughout data work. Analysts encounter incomplete evidence and unresolved judgments while discovering, cleaning, shaping, and communicating data [8, 46]. Digital-humanities workflows face similar challenges when records conflict or provide uneven evidence. Franke et al. treat confidence as a first-class data attribute so analysts can carry those judgments through processing and visual analysis [22]. Panagiotidou and Vande Moere broaden the representational account further by characterizing qualitative uncertainty through contextual and socio-technical conditions in data production [51]. This work expands the kinds of non-exactness that visualizations can represent and the stages at which it may arise.

Several constructive approaches externalize knowledge that standard data models omit. Data Hunches allows experts to annotate visualizations with personal knowledge about the data [38]. McCurdy et al. provide a framework for externalizing implicit error embedded in scientific practice [41]. Nowak and colleagues design for ambiguity in avalanche forecasting, where analysts combine incomplete evidence with professional judgment [49, 50]. These approaches preserve knowledge around an object that has already become representable.

Premature crisping examines the commitment that produces that object. It asks how vague expression acquired a structured form, where the system settled its meaning, and what recourse remains afterward. A system may support annotations or confidence judgments while still obscuring how it converted “pretty bad lately” into a score, interval, or category. Premature crisping therefore complements constructive approaches by diagnosing the resolution step that precedes their use.

Ambiguity-oriented visualization provides another representational stance. Close-reading systems and literary visualizations can retain several possible readings instead of collapsing them immediately into one result [30, 31, 42]. Visual text analysis research likewise traces how ambiguity and underspecification enter computational pipelines [26]. Premature crisping narrows that concern to the moment when a system stabilizes vague source meaning and renders the selected interpretation as settled.

2.3 Information Loss, Categorization, and Provenance

All representations transform their source material. Information theory treats loss as central to communication, compression, and approximation [58], while visualization changes data through abstraction, interaction, and encoding [63]. Premature crisping concerns a specific form of loss in which a system settles semantically vague meaning and leaves relevant actors unable to distinguish that settlement from the source.

Research on categorization shows how such commitments arise before visual encoding. Grounded theory and thematic analysis treat coding as interpretive work because researchers construct categories through engagement with their material [23]. Aeonium surfaces disagreement and ambiguity during collaborative coding [18], and machine-learning support for qualitative analysis can retain ambiguity during the analytic process [14]. Ground-truth research makes a parallel argument. Variation among human labels may express meaningful disagreement [53], including cases where disagreement persists under careful annotation protocols [52]. Labels also emerge through social and institutional negotiation [47]. A crisp category may therefore record the outcome of interpretation while concealing the alternatives that preceded it.

Visualization pipeline and provenance research offers mechanisms for tracing transformations. Reference models describe movement from source data through transformation and visual mapping [12]. Provenance systems record where data came from and why a result appears [11]. Wrangler exposes transformation histories as part of data work [33], and related research connects wrangling decisions to credibility and downstream trust [32]. Missing-data research offers a related warning: silent imputation or omission can make absent values appear observed [19, 60].

A transformation record supports interpretive recourse only when an actor can use it during the relevant task. A database may retain "sometime in the morning" while a timeline presents only an exact point. The pipeline remains traceable, yet the viewer may still attribute the selected time to the patient. Premature crisping extends provenance analysis by asking whether relevant actors can recognize the imposed interpretation, reach the vague source, and respond when the resolution does not fit.

2.4 Visualization Beyond Representation

Subjective and humanistic visualization explain why source meaning deserves attention during formalization. Personal storytelling can legitimately incorporate partial and situated perspectives [13]. Humanistic approaches similarly frame visualization as a medium through which people encounter and construct meaning [4]. From this perspective, formalization shapes the terms under which a source becomes legible. Preserving context matters because the representation influences how later viewers understand the person, experience, or account behind the data.

Critical and feminist approaches make the authority of representation more explicit. Critical InfoVis treats visualization choices as value-laden decisions [17]. Data feminism emphasizes pluralism and context while challenging reductive categories [16]. Feminist theory can also reshape the boundaries and outcomes of visualization research itself [2]. Together, these perspectives show how formal representation privileges some interpretations and narrows others.

Premature crisping identifies one mechanism through which that authority takes hold. When a system converts vague self-report or narrative into a stable label, the formalization determines which interpretation becomes legible as data. Later users may then encounter the selected interpretation as a property of the source rather than the outcome of a design decision. The distinction matters because the system can give one actor's interpretation operational authority over another's account.

HCI research on human experience provides concrete examples. Affective computing often operationalizes emotion, mood, stress, or engagement through discrete labels and dimensional models [56, 61]. Interactionist accounts describe emotion as emerging through social activity and context [5, 6]. Personal informatics systems also organize lived experience through stages and structured inputs, even though tracking practices remain uneven across people and situations [20, 37]. These operationalizations determine which aspects of experience remain available once a system turns them into data.

The mood-tracking case in Section 7 examines that process across several positions in the workflow. The person recording a mood, the viewer interpreting the history, and the organization maintaining the system may each hold different authority over the formalized meaning. Premature crisping makes that distribution visible by tracing how one interpretation becomes stabilized and whose recourse remains afterward.

2.5 The Nested Model as Analytic Anchor

Visualization design models provide vocabulary for locating commitments and their validity threats. Munzner's nested model organizes design and validation into four layers: domain problem and data characterization, operation and data type abstraction, visual encoding and interaction, and algorithm design [48]. Domain and abstraction choices constrain what later encodings, interactions, and algorithms can preserve.

Design-study methodology describes how researchers move from fuzzy domain situations toward tasks, abstractions, and implemented systems [57]. The design activity framework traces work across understanding, ideation, making, and deployment [43]. Criteria for rigor in visualization design study emphasize reflexivity and transparency around the interpretations that shape a system [44]. These frameworks already recognize visualization design as interpretive work whose early commitments shape later representations.

I use Munzner's nested model as an analytic scaffold. Its layers support a cross-layer account in which an early framing decision shapes the abstraction, limits the visual encoding, and determines what an algorithm can return. The model therefore helps trace how system-imposed precision enters a pipeline and becomes increasingly difficult to distinguish from precision supplied by the source.

3 Visualizing Vague, Language-Derived Data

Visualizing vague, language-derived data concerns systems that turn human expression into structure while meaning remains partly unsettled. People often communicate effectively without fixing exact boundaries. To make such expression computationally usable, a system must select a formal representation. That selection establishes a relationship between what a person expressed and the interpretation the system made operational.

Vagueness belongs to the source expression itself. 'Severe pain" takes meaning from a personal baseline and an implicit comparison class. 'Early last year" locates an event at a coarse temporal granularity. "Kind of slow" softens an evaluation without assigning it a degree. A system can translate each expression into a score, interval, or category, but the resulting value reflects an interpretation. The contextual conditions that shaped the original meaning remain relevant after formalization.

Vague expression passes through multiple representational decisions. Designers may deliberately preserve indeterminacy while concepts and representations remain provisional [1], yet later formalization can make the result appear more settled and obscure its relationship to the source.

As discussed in the related work, statistical uncertainty usually concerns variation around an already-defined quantity, whereas vague language may leave the quantity or its boundaries open. Qualitative and ambiguity-oriented approaches preserve context or competing readings once material has become representable. The present problem arises when a system must give an expression computational form while part of its meaning remains unresolved.

Whether a resolution fits depends on the setting. A coarse date may serve a patient recounting a life history while providing too little detail for a clinician reconstructing medication timing. Cultural conventions and institutional practices further shape what speakers and listeners consider adequately precise. The same expression may therefore support one task and require clarification in another. Its adequacy follows from the demands and expectations of the setting in which people use it.

Resolution often serves a necessary purpose. Systems need stable representations to compute and coordinate, and a task may justify selecting one interpretation. Once that interpretation guides later reasoning, it also acquires authority. The actors who depend on the result may therefore need enough recourse to understand how the value was resolved and challenge it when it no longer fits the task. The consequences of misinterpretation determine how much recourse the workflow should provide.

The problem space therefore calls for diagnostics that trace how vague expression becomes structured and assess what remains available afterward. Premature crisping supplies such a diagnostic by identifying when imposed precision becomes difficult to distinguish from precision provided by the source. The following section traces that failure across the visualization design process.

4 Premature Crisping

Premature crisping is a transversal validity threat in visualization systems that transform vague human expression into structured visual representations. It occurs when a system silently resolves semantically vague input into a representation that appears more precise than the source warrants, making system-imposed precision difficult to distinguish from source-provided precision. I use Munzner's nested model as an analytic scaffold to trace the threat across four design sites: domain characterization, abstraction, visual encoding and interaction, and algorithm design (see Figure 1).

4.1 The Transversal Threat

All visualizations select, abstract, and transform source material. Premature crisping identifies a narrower failure in which a system resolves vagueness into crisp structure while obscuring the interpretive work that produced it. The failure is **premature** because the system commits to a precise interpretation before the source warrants such precision. The failure is **crisping** because the resulting value can look identical to data that was precise from the start.

For example, a patient who reports that symptoms have been "getting worse lately" has not provided a severity score, start date, or fixed temporal interval. The phrase expresses change, comparison, and partial temporal specification. A visualization system may need to derive a score, category, or interval from that phrase, but the validity threat arises when viewers cannot tell whether the resulting value came from the patient or from the system.

Premature crisping can occur even when a pipeline preserves some provenance. A database may retain the original phrase while a visualization foregrounds only a crisp score, category, or position. The source remains available somewhere in the system, but the rendered representation can still make one interpretation appear authoritative. Section 5 develops interpretive recourse as a way to assess whether relevant actors can recognize, access, and respond to imposed interpretations.

4.2 Cross-Layer Propagation

The nested structure matters because early commitments constrain later representations. Consider the phrase "sometime in the morning." A design team may characterize the domain problem as placing symptom events on a timeline. The abstraction may then require a single timestamp, a parser may map the phrase to 9:00 a.m., and the encoding may render a point at that position. Viewers may then read the point as a time the patient reported directly.

I call the sequence of commitments through which vague expression becomes more settled a **crisping pathway**. A crisping pathway traces where imposed precision first enters, how later sites preserve or amplify it, and how the resulting representation becomes consequential for downstream actors. The sites do not always operate in a fixed order. A pathway may begin with an early framing decision, a restrictive abstraction, an algorithmic transformation, or a visual encoding that makes an existing commitment appear more precise than the source warrants.

No single downstream decision fully explains the resulting precision. In the timeline example, the domain framing made exact placement desirable, the abstraction made exact placement necessary, the algorithm supplied a timestamp, and the encoding made the timestamp appear source-provided. The crisping pathway shows how a commitment at one site can constrain or amplify later decisions. A system can interrupt propagation by retaining a coarse temporal value or the original phrase upstream. It can then mark the position as inferred or let viewers inspect the resolution at the point of use. Premature crisping becomes harder to repair as imposed precision travels downstream and begins to look ordinary.

4.3 At Domain Characterization

Site and commitment. Domain characterization commits designers to a problem framing. Munzner's nested model identifies misunderstanding the domain problem as the primary threat at the layer. Premature crisping introduces a related threat: designers may understand the broad domain while framing a vague phenomenon as more settled, measurable, or categorical than the source warrants.

Crisping threat. Domain characterization often starts from interviews, observations, documents, or field notes. Such sources contain hedges, comparisons, partial temporal references, and uneven specificity. During analysis, designers move from transcripts to codes, from codes to themes, and from themes to requirements. Each step can make vague expression cleaner than the source.

A transcript line such as ‘things have been kind of off lately" might become a code for ‘symptom variability," then a requirement to "support symptom variability tracking." The requirement may guide design, but it no longer preserves the unsettled and contextual quality of the original account. The system now targets structured variability rather than the vague experience that motivated the requirement, and later design decisions inherit the narrower framing.

Diagnostic focus. At the domain layer, designers should ask whether the framing still reflects the vague phenomenon that motivated the system. Round-trip elicitation can expose early crisping: designers can present themes or requirements back to source users and ask whether the framing feels faithful. Designers can also preserve direct quotes alongside requirements and name vague phenomena explicitly.

4.4 At Abstraction

Site and commitment. Abstraction commits designers to data types and operations. Premature crisping occurs when the data structure cannot distinguish source-provided precision from system-imposed precision, lowering the ceiling on what downstream layers can preserve.

Crisping threat. Abstraction is often where language becomes data. Categorical attributes assume discrete membership, quantitative attributes assume comparable values, and temporal attributes often assume dates, intervals, or bins. Vague input can strain those assumptions because it may be gradable, mixed, multi-granular, reference-class-dependent, or partially specified.

A phrase such as ‘pretty bad” contains severity information, but the phrase is not equivalent to a number on a fixed scale. A response such as ‘the service was kind of slow, but the food was great” contains topic-specific sentiment and contrast rather than one clean label. A temporal phrase such as “sometime in the morning” contains temporal information without specifying an exact time. A preservation-oriented abstraction can store a value, label, or interval while retaining the source phrase, precision status, inference, and unresolved alternatives where relevant.

Diagnostic focus. At the abstraction layer, designers should ask whether downstream layers can distinguish what the source stated from what the system inferred. Structured representation review can expose failures before implementation. Users or domain experts can inspect abstracted records and identify whether the structure reflects the source expression.

4.5 At Visual Encoding and Interaction

Site and commitment. Visual encoding and interaction commit designers to how structured values appear and how viewers can inspect or revise them. Premature crisping at the combined layer concerns what the rendered mark implies and what interpretive access the interface realizes for the viewer.

Crisping threat. The main failure pattern is misleading precision: inferred, approximate, mixed, or unresolved values receive the same visual treatment as values stated precisely by the source. A timeline that snaps ‘early in the year” to January 1 makes a coarse expression look like an exact date. A sentiment dashboard that renders a hedged response as a clean bar segment can make an interpretation look directly stated. A mood chart that turns ‘bad today” into one point can make a complex day look commensurable with every other daily point.

Interaction can mitigate or intensify the threat. A system may render an inferred value as a point for layout reasons while allowing viewers to inspect the source phrase, see the inference, view alternatives, or revise the value. Conversely, a visually marked approximation may still become practically final when the interface hides source text, provenance, or revision mechanisms.

Diagnostic focus. At the encoding and interaction layer, designers should ask whether viewers can distinguish source statements from system inferences. Source-precision calibration tasks can test the distinction by asking viewers to classify marks as stated, inferred, approximate, mixed, unresolved, or missing. Designers should also test whether viewers can inspect source language and revise local resolutions.

4.6 At Algorithm Design

Site and commitment. Algorithm design commits designers to computational procedures. Munzner’s nested model emphasizes runtime and memory threats at the layer. Premature crisping adds a semantic threat: algorithms may resolve vagueness in ways that lower the ceiling on what users, designers, or downstream components can access.

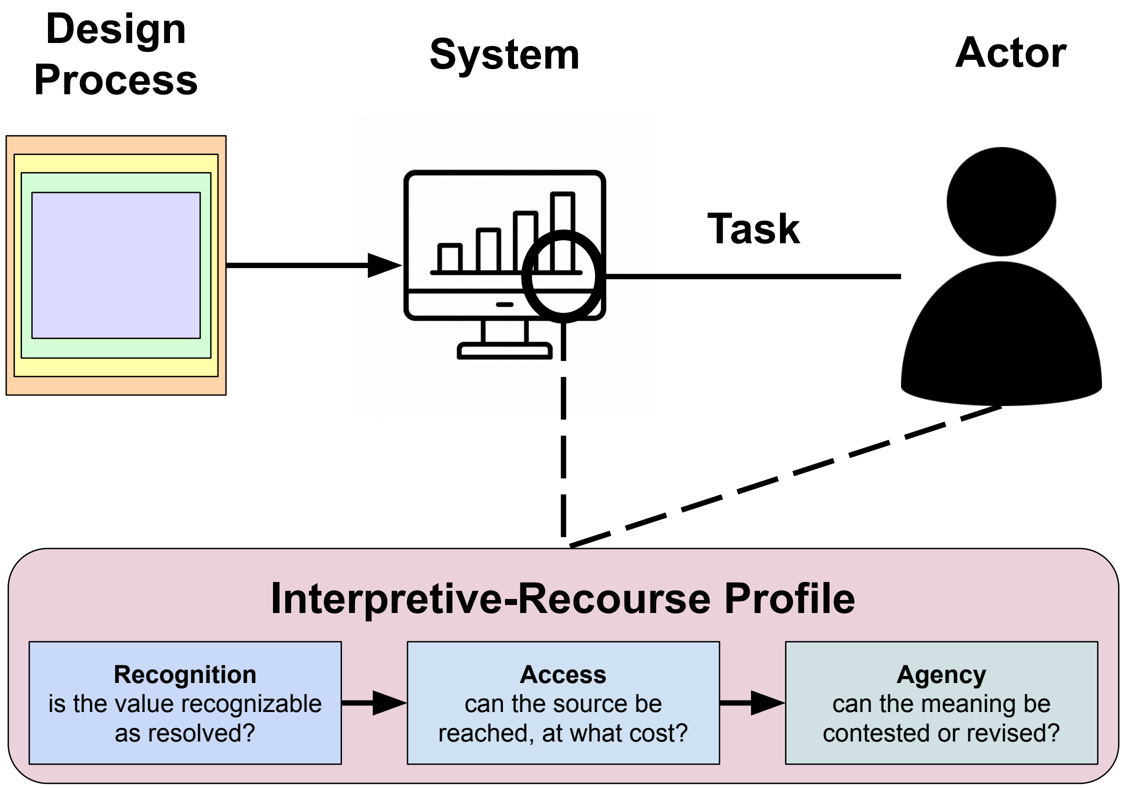

Fig. 2. The relationship between the design process and an actor- and task-dependent interpretive-recourse profile. Decisions made across the premature-crisping sites shape the resulting visualization system and constrain what recourse remains possible. When an actor uses the system to perform a task, the resulting profile describes whether the actor can recognize that a value was resolved, access the source vagueness, and contest or revise the stabilized interpretation.

Crisping threat. Algorithms often require concrete outputs: sentiment models return labels or scores, temporal parsers return dates or intervals, layout algorithms assign positions, and clustering algorithms assign membership. Concrete output does not automatically produce premature crisping. The threat arises when an algorithm returns a single opaque output where the source supported a hedged, mixed, approximate, or unresolved interpretation.

Rule-based and model-mediated systems create different repair opportunities. A rule that maps 'early last year" to January 1 exposes a commitment that designers can modify or display. A model that converts 'kind of slow, but the food was great" into a positive sentiment label may hide the basis for the decision. In both cases, designers should ask whether relevant actors and downstream components can identify what the algorithm added.

Diagnostic focus. At the algorithm layer, designers should inspect whether outputs retain source phrases, precision status, alternatives, or provenance. Tests with deliberately vague inputs can expose crisping pressure, including hedged evaluations, mixed sentiment, coarse dates, unresolved alternatives, and reference-class-dependent terms. Accuracy against a labeled dataset is not sufficient because labels may already contain premature crisping.

5 Interpretive Recourse After Crisping

Visualization systems sometimes need to turn vague expression into structured representation. **Interpretive recourse** names whether people can still understand and respond to the vagueness that the system has structured. Interpretive recourse depends on the actor and task: a clinician, patient, designer, system maintainer, or downstream system may need different forms of standing relative to the same imposed interpretation.

For a given actor and task, an **interpretive-recourse profile** describes three components: **recognition**, **access**, and **agency** (see Figure 2).

The components form a soft escalation rather than a strict sequence. Recognition can prompt an actor to seek additional context, and access can support meaningful action. However, systems can separate the components in practice. A marked inference may provide recognition without source access, while a correction control may offer agency without enough context to support an informed revision.

Visualization research has previously decomposed interaction and tasks into distinct components [9, 66]. The interpretive-recourse profile addresses a different question: what standing does an actor have relative to an interpretation imposed by the system? The profile also differs from interpretability, which traditionally asks whether an actor can read structure from a display. Interpretive recourse asks whether an actor can identify what the system added and respond when the resulting interpretation does not fit.

5.1 Crisping Pathways and Interpretive-Recourse Profiles

A crisping pathway traces where a system adds or amplifies precision. Domain characterization, abstraction, algorithm design, and visual encoding and interaction can each shape how vague expression becomes structured representation. An interpretive-recourse profile describes what a particular actor can understand and do after encountering the result.

As argued in the previous section, upstream decisions constrain what downstream layers can preserve. If an abstraction discards the original phrase, a later encoding cannot restore it. If an algorithm stores only one date, an interface cannot explain the original temporal vagueness unless another record preserves it. Visual encoding and interaction then determine what an actor can recognize, access, or revise in practice.

The crisping pathway determines what recourse remains possible; the actor's role and task determine what recourse remains reachable.

5.2 Recognition

The first profile component is **recognition**: whether an actor can tell that the system resolved a value rather than receiving it directly from the source.

For example, a timeline may place an Event described as occurring "early last year" at a representative date. A clinician should be able to distinguish the representative position from a date supplied by the patient. A visual cue, annotation, or label may signal that the placement reflects an inference.

Recognition does not guarantee access. A timeline may mark a point as inferred without showing the original phrase or explaining the rule that produced the position. Access also does not guarantee recognition. A timeline may display the phrase "early last year" while rendering an exact-looking point that viewers read as a patient-provided date. The second case matters because the source vagueness remains available but fails to shape interpretation.

5.3 Access

The second profile component is **access**: whether an actor can reach source vagueness and how much effort the pathway requires. The appropriate form of access depends on the actor, task, and consequences of misinterpretation.

Visible access makes source vagueness available in the primary representation. For example, a timeline may display the phrase “early last year” directly within an Event mark.

Inspectable access lets an actor reach source vagueness through ordinary interaction. For example, a clinician may hover over a timeline point to view the original phrase and see that the position reflects a layout decision.

Auditable access preserves source vagueness through records or transformation histories outside the primary workflow. For example, a designer may inspect how the system converted “during high school” into a timeline span.

Privileged access preserves source vagueness only through pathways outside the actor’s ordinary permissions or workflow. For example, a system maintainer may query a database or inspect parser outputs that ordinary users cannot access.

Inaccessible resolution leaves no pathway available to the relevant actor. For example, a downstream system may receive only a representative date without a machine-readable precision indicator.

The gradient describes access relative to an actor and workflow. The same log may provide auditable access for a designer while remaining inaccessible to a clinician during an appointment. Effort also matters; a tooltip may fit a workflow that requires occasional inspection, while a separate database query may not fit a short clinical appointment.

Access also has costs. Visible source phrases, annotations, tooltips, logs, and revision controls can add interface clutter, cognitive load, and technical complexity. Designers should calibrate access to the actor and task rather than expose every available trace.

5.4 Agency

The final profile component is **agency**: whether an actor can contest or revise a stabilized interpretation.

Contestability means that an actor can question or reject the interpretation. A patient should be able to indicate that an inferred placement misrepresents the intended meaning of “early last year.”

Revisability means that an authorized actor can update the structured representation without discarding the original source. A patient may revise an inferred date while preserving the original phrase and surrounding timeline. A system maintainer may revise the transformation rule itself.

Different actors may need different forms of agency. A clinician scanning a timeline may only need to recognize an inference. A patient reviewing the authored story may need to contest or revise a placement. A system maintainer may need to change the underlying rule. A downstream system may need metadata that prevents an inferred value from silently becoming an exact input to later computation.

Provenance can support interpretive recourse, but provenance does not guarantee it. A transformation history may let a designer recognize and access the phrase “early last year” while leaving a clinician unable to recognize that a plotted point reflects an inference. A system may also preserve recognition and access while withholding any way to contest or revise the result. Interpretive recourse therefore depends on both perception and action: the actor must be able to identify what the system added and respond when the stabilized interpretation does not fit.

6 Auditing for Premature Crisping

The interpretive-recourse account supports a structured audit of where a visualization pipeline introduces precision that the source did not provide. The audit does not eliminate interpretation or maximize recourse in every interface. Instead, it identifies mismatches between the precision a representation implies and the recourse available to relevant actors.

Designers can begin with either a vague source expression or a crisp output that may conceal one, then work through four steps. First, they trace the crisping pathway by locating where vague expression becomes more settled across domain characterization, abstraction, visual encoding and interaction, and algorithm design. Second, they record what the source expressed, what interpretation the system added, and how the added interpretation affects downstream use. Third, they construct an interpretive-recourse profile for each relevant actor and task by assessing recognition, access, and agency. Finally, they judge whether the available recourse fits the workflow and consequences of misinterpretation.

The audit produces a contextual judgment rather than a binary score. A representative timeline position may preserve sufficient recourse when a clinician can recognize the placement as inferred and inspect the source phrase with little effort. The same position may constitute premature crisping when an unmarked point appears to represent an exact patient-provided date. In another case, the original phrase may remain visible while the plotted position still appears source-provided. The audit helps distinguish those failures rather than treating source preservation alone as sufficient.

A single vague expression may produce multiple audit entries as it moves through a pipeline. One interpretation may enter through abstraction, another through an algorithmic transformation, and another through visual encoding. Recording the interpretations separately helps designers locate where imposed precision first enters and where later decisions amplify it. Section 7 applies the audit across three systems, and Table 1 summarizes the resulting judgments.

7 Case Analyses

I apply the audit to three systems that structure vague human expression: Qualtrics Text iQ [55], Daylio [25], and HealthTale [3]. The cases illustrate different crisping pathways. Qualtrics Text iQ transforms open-ended feedback after expression, as survey text becomes topics, sentiment labels, scores, and dashboard elements. Daylio asks users to categorize affective experience during input so the system can generate longitudinal summaries. HealthTale provides a preservation-oriented contrast: the system retains source language and flexible temporal specifications while still introducing layout and interpretation risks.

I analyze each system through its available interface and supporting materials. The vantage point provides one actor standing among several: I can inspect what the interface surfaces and reason backward about crisping pathways, but I cannot always determine how internal transformations operate. Where the interface does not reveal an internal mechanism, the analysis records the limit rather than inferring undocumented behavior. For each case, I trace where precision may enter across the four premature-crisping sites, then construct an interpretive-recourse profile for each relevant actor and task. The analyses illustrate how to use the audit rather than evaluate the systems as successes or failures. Table 1 summarizes the resulting judgments.

7.1 Qualtrics Text iQ

Qualtrics Text iQ supports analysis of open-ended feedback by assigning topics, performing sentiment analysis, and reporting results through widgets and dashboards [55] (see Figure 3). Consider a restaurant feedback response: “The

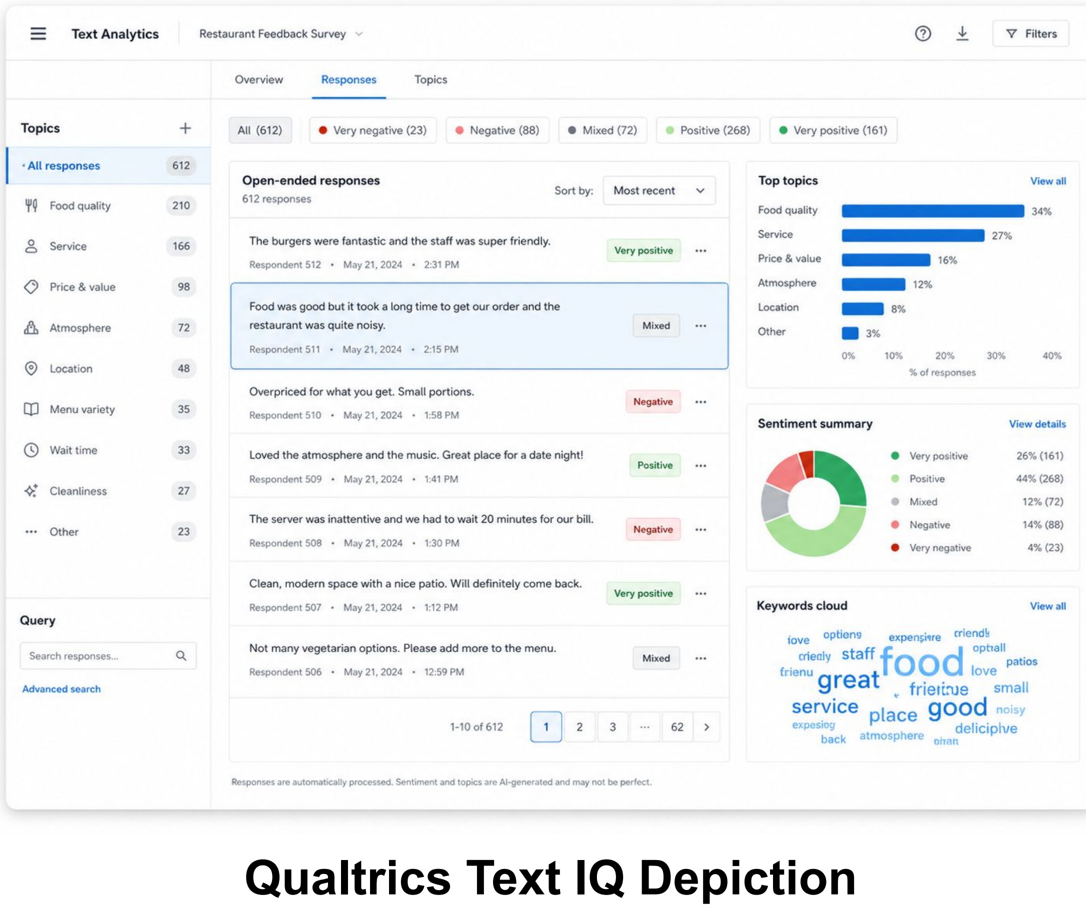

Fig. 3. An abstraction of a Qualtrics Text iQ-style workflow in which a system transforms open-ended responses into topics, sentiment labels, scores, filters, charts, and dashboard summaries.

service was kind of slow, but the food was great." The response contains a hedge, a contrast, and topic-specific evaluations. A Text iQ-style workflow transforms the response into topics, sentiment labels, scores, and dashboard elements.

Crisping pathway. At the *domain-characterization* site, the workflow treats open-ended feedback as material that organizations can categorize, compare, and summarize. The framing supports large-scale analysis, but it also treats sentiment as a stable property of the response. The respondent did not provide one sentiment value. The sentence contains a softened criticism of the service and praise for the food.

At the *abstraction* site, the workflow converts the sentence into structured attributes such as topics, sentiment labels, and scores. Text iQ assigns sentiment to the response as a whole. A label such as *Mixed* preserves more complexity than a single-valence label, but the abstraction still compresses the hedge, contrast, and relationship between the evaluations. At the *algorithm-design* site, aggregation amplifies the initial resolution by counting, ranking, or comparing responses using assigned values. At the *visual-encoding and interaction* site, dashboards present the attributes through bars, colors, filters, and aggregate summaries. Viewers may read the resulting value as a property of the response rather than as an interpretation produced by the workflow.

Interpretive-recourse profile and fit. The central imposed commitment is the conversion of a mixed and hedged response into response-level sentiment attributes that support aggregation. The relevant actor is an analyst reviewing aggregate feedback to identify patterns or make organizational decisions. The interface provides *recognition* by identifying sentiment values as system-generated rather than respondent-provided. The analyst also has *access* to source text through drill-down from the dashboard to individual responses. The effort grows with the number of responses and the distance between an aggregate summary and the underlying text. The analyst also has *agency*, as the interface allows analysts to revise assigned sentiment values. Respondents also do not see the resulting analysis or revise their responses after submission. Interpretive recourse may fit exploratory analysis when analysts use sentiment widgets as entry points for closer reading. The same interface may support premature crisping when organizations act on aggregate scores without revisiting the mixed and hedged expressions that produced them.

7.2 Daylio

Daylio lets users record moods, activities, notes, and photos then review the records through charts, statistics, calendars, correlations, goals, and year-in-pixels views [25] (see Figure 4). Daylio does not convert extended free text into sentiment analytics. Instead, the app asks users to structure affective experience at input. Consider a user who thinks, 'I feel bad today." The word 'bad" may refer to sadness, fatigue, anger, anxiety, physical discomfort, or a mixture of experiences. Selecting a mood category turns the vague feeling into a repeated entry for later visualization.

Crisping pathway. Daylio introduces its central commitment at the *domain-characterization* and *abstraction* sites. The app frames mood tracking as a low-friction practice of recording feelings, activities, and longitudinal patterns. The framing makes mood categorizable. The interface then asks the user to choose a discrete mood that can stand in for a complex day. Optional notes, photos, audio, and custom tags can preserve context, but the selected mood remains the primary structured value that the interface carries into later summaries.

At the *algorithm-design* site, statistics and correlations further stabilize the selection by linking mood categories to activities or longitudinal trends. At the *visual-encoding and interaction* site, charts, calendars, and year-in-pixels views aggregate selections over time. A colored dot or plotted point can make one category appear to summarize an entire day and look directly comparable to every other daily value. The later outputs depend on an interpretive commitment introduced at input: one chosen category represents the day.

Interpretive-recourse profile and fit. The relevant actor is often the same person at two moments: the user records the experience during input and later interprets summaries during reflection. Because the user selects the mood category, the interface provides *recognition* that the value reflects a prior act of categorization. Summary views also link back to individual entries, giving the user *access* to notes, activities, photos, audio, and custom tags. The effort remains low for opening one entry but grows when reflection requires reviewing many entries behind an aggregate pattern. Users have some *agency*: they can revise saved mood categories and contextual information. However, the app does not preserve a revision history, and users cannot redefine the core mood categories beyond adding custom tags. A compact dot may fit broad pattern scanning while becoming misleading when a user treats the summary as a complete record of how each day felt.

Fig. 4. An abstraction of a Daylio-style workflow in which a user selects a mood, adds activities or notes, and later views the entry through charts, calendars, or summaries.

7.3 HealthTale

HealthTale elicits free-form health narratives and transforms them into structured timeline representations for initial clinical encounters [3] (see Figure 5). HealthTale provides a preservation-oriented contrast because the system carries forward patient-authored language, temporal imprecision, and narrative context. Consider a patient who says: “I hurt my knee sometime in high school, then I had a shoulder injury in 2020, and I am still dealing with back pain now.” The account combines an approximate life period, a calendar year, and an ongoing condition.

Crisping pathway. HealthTale begins with a preservation-oriented *domain characterization*: patient-authored health stories contain uneven temporal detail, contextual meaning, and lived experience that should not be forced into a complete clinical chronology. At the *abstraction* site, the system stores Events with structured attributes and patient-authored text. Events can remain unspecified, partially specified, or fully specified. The abstraction can therefore distinguish ‘sometime in high school,” ‘in 2020,” and “still ongoing” rather than flattening the expressions into one temporal format.

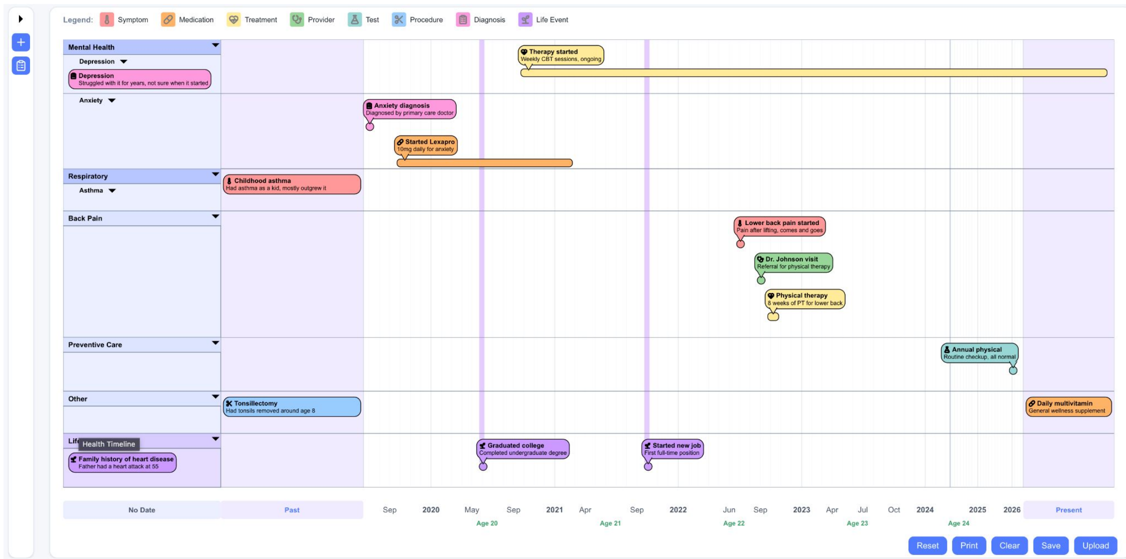

HealthTale

Fig. 5. A HealthTale visualization of a patient-authored health story. The system transforms free-form patient narrative into structured Events with flexible temporal specifications, patient-authored text, and timeline-based organization.

Crisping pressure enters through the relationship between the *algorithm-design* and *visual-encoding and interaction* sites. HealthTale places Events with no explicit date, broad past or current timing, and fully specified dates in separate regions. Patient-authored text remains visible within each Event mark, preserving some source vagueness directly. The structured record also retains machine-readable temporal precision. However, when the layout assigns a representative position to a coarsely specified Event, the rendered mark may imply greater precision than the patient provided. For example, the layout may place an Event specified only to a year or another coarse interval at a representative midpoint. The interface does not visually distinguish the representative position from an exact patient-provided date.

Interpretive-recourse profile and fit. The relevant actors include the patient reviewing the authored story and the clinician scanning a static timeline PDF during an appointment. HealthTale illustrates why *recognition* and *access* need to remain separate components. Event marks provide visible *access* to patient-authored wording, and separate no-date, past, current, and timeline regions preserve some temporal distinctions. However, representative positions do not receive a distinct visual treatment or explicit label. A plotted position may therefore look exact even when the Event text preserves a coarser temporal expression. Viewers can have high access to source text while lacking clear *recognition* that the position reflects an inference.

Patients have some *agency* during authoring because they can revise inferred dates. Revisions do not preserve the original temporal value separately from the patient-authored text. Clinicians receive a static PDF and cannot revise placements during use. The interpretive-recourse fit therefore depends on whether patient-authored wording provides enough context for viewers to recognize representative positions without reverse-engineering the layout.

Table 1. Case comparisons using the premature-crisping audit. The table records interpretive-recourse profiles based on the inspected interfaces and available supporting materials.

Case	**Crisping pathway**	**Relevant actor and task**	**Recognition**	**Access**	**Agency**	**Fit**
Where does precision enter or become amplified?	*What commitment becomes consequential?*	*Who uses the result, and for what task?*	*Can the actor tell that the value was resolved rather than source-stated?*	*Can the actor reach the source vagueness, and at what cost?*	*Can the actor contest or revise the interpretation?*	*Does the available recourse fit the workflow and consequences?*
Qualtrics Text iQ	Domain characterization, abstraction, algorithm design, and visual encoding progressively stabilize a hedged and mixed response as one response-level sentiment value and aggregate dashboard marks.	Analyst reviewing aggregate feedback to identify patterns or support organizational decisions.	Present. The interface identifies sentiment values as system-generated rather than respondent-provided.	Inspectable. Analysts can drill down from the dashboard to the original response. Effort grows with the number of responses and the distance between aggregate summaries and source text.	Limited. Analysts cannot revise assigned sentiment values. Respondents do not see the resulting analysis or revise their responses after submission.	Fits exploratory analysis better when analysts treat summaries as entry points for closer reading. May support premature crisping when decisions rely only on aggregate scores.
Daylio	Domain characterization and abstraction stabilize affective experience as one selected mood category. Algorithmic summaries and visual encodings carry the category forward as a repeated unit.	User recording an experience at input, then reviewing patterns during later reflection.	Present. Because the user selects the mood category, the value reflects an explicit act of categorization. Summary views still risk making the category appear to fully represent the day.	Inspectable. Users can open individual entries from summary views and review notes, activities, photos, audio, and custom tags. Effort remains low for one entry but grows when a pattern requires review across many entries.	Moderate. Users can revise saved moods and contextual information, but the app does not preserve revision history. Users can add custom tags but cannot redefine the core mood categories.	Fits broad pattern scanning better than precise retrospective interpretation of how each day felt.
HealthTale	Algorithm design assigns representative positions to coarsely specified Events, and visual encoding can make the positions appear more exact than the source warrants.	Patient reviewing an authored story; clinician scanning a static timeline PDF during an appointment.	Low. Representative positions do not receive a distinct visual treatment or explicit label, so coarse dates may look exact.	High for source text. Event marks display patient-authored wording, and the structured record retains machine-readable temporal precision. However, viewers cannot inspect the structured temporal value separately unless the wording appears directly in the Event text.	Uneven. Patients can revise inferred dates during authoring, but revisions do not preserve the original temporal value separately from the source text. Clinicians cannot revise the static PDF during use.	Fits only when patient-authored wording gives viewers enough context to distinguish representative positions from patient-provided dates. HealthTale illustrates how high access can coexist with low recognition.

7.4 Case Analyses Summary

Together, the three cases apply the audit to different crisping pathways. Qualtrics Text iQ introduces structured interpretation after open-ended expression and may amplify the interpretation through aggregate dashboards. Daylio asks users to resolve vague affective experience during input and later treats the selected category as a repeated visual unit. HealthTale preserves more source-level vagueness but still risks introducing misleading precision through layout. Across the cases, the relevant question is whether the interpretive-recourse profile for each actor and task provides enough recognition, access, and agency for the consequences of misinterpretation.

8 Discussion

Premature crisping identifies a specific validity threat within the broader problem of visualizing vague, language-derived data. Systems often need to stabilize meaning before they can store, compute, or display it. The diagnostic traces where that stabilization introduces precision and asks whether relevant actors retain enough interpretive recourse to understand and respond to the result. In this section, I clarify the diagnostic's scope, consider how recourse can guide design, and identify directions for future work.

8.1 Scope of the Diagnostic

The diagnostic applies most clearly when a source expression contains meaningful semantic vagueness and a system formalizes that expression into a representation that appears more settled than the source warrants. The relevant commitment may enter at any design layer. The audit then evaluates whether the resulting recourse fits the actor, task, setting, and consequences of misinterpretation.

Visualization relies on many forms of information reduction. Aggregation, filtering, and abstraction can all remove detail while preserving the intended meaning of the source. Premature crisping concerns a narrower transformation in which the system settles an interpretation and makes the resulting precision difficult to distinguish from source-provided information.

Uncertainty typologies describe many sources and forms of non-exactness [39, 62]. Measurement uncertainty, missing data, probabilistic models, and simulation each raise their own visualization concerns [7, 10, 19, 60]. Premature crisping focuses on semantic vagueness. A confidence interval describes uncertainty around an established quantity. An expression such as “pretty bad lately” leaves its comparison standard and temporal scope partly open [27, 34, 35].

The diagnostic also addresses an earlier stage than many constructive approaches to qualitative uncertainty. Data Hunches lets experts attach personal knowledge to structured data [38], while implicit-error frameworks externalize contextual knowledge that formal models omit [41]. Ambiguity-aware visual analytics and confidence-as-data approaches preserve judgment around an established object of analysis [22, 49]. Premature crisping asks how that object acquired its structure. It examines whether the system silently resolved vague expression and whether later actors can identify or revise that resolution. The diagnostic therefore complements techniques that preserve contextual knowledge after formalization.

8.2 Designing for Interpretive Recourse

Interpretive recourse provides a calibrated design target. People encounter imposed interpretations from different positions in a workflow, so the same representation may support one actor while constraining another. A clinician may need quick recognition of an inferred value. A patient may care more about whether the value remains revisable. A maintainer may need access to the transformation rule that produced it.

The profile also reveals how a system distributes interpretive authority. A source user may recognize that a value was inferred and still lack any way to change it. In that arrangement, another actor or institution controls the formalized meaning. The audit makes the allocation visible by showing whose interpretation remains contestable and whose becomes operationally final.

Vagueness remains situated throughout this process. Cultural conventions and institutional routines shape how much precision an expression conveys and how much a task demands. “Early” may carry an accepted range in one setting and require clarification in another. Actor- and task-dependent profiles provide a way to examine those differences without assuming one universal threshold for adequate recourse.

Crisping pathways show where recourse becomes possible. Early design choices determine whether the system retains source language or records that a value was inferred. Later interface decisions determine whether an actor can use that information during the task. A clinical system, for example, may preserve the original phrase in its database while presenting only the resolved value during an appointment. The pipeline retains a trace, but the workflow offers little practical recourse.

Supporting recourse also changes the interface. A visible marker can make inference status clear, yet repeated markers may compete with the information users came to inspect. Source access can help someone resolve a consequential ambiguity, though opening that context for every value can interrupt routine use. Revision creates another tension because changing a resolved value may affect later summaries or other users' interpretations. The design challenge lies in deciding when added context becomes useful and when it becomes burdensome.

Designers can respond by staging recourse according to need. A low-stakes overview may require only a subtle cue that more context exists. A consequential decision may warrant direct access to the source and a clear path for correction. The audit helps designers make that calibration explicit by comparing the precision a representation implies with the recourse the workflow actually provides.

The diagnostic also points toward a constructive counterpart for visualizing vague, language-derived data. Future research can examine how systems preserve coarse expressions and contextual meaning while still supporting computation. Interfaces might reveal additional context only when a task demands it, allowing routine reading to remain efficient while keeping deeper inspection available. Premature crisping offers a way to evaluate such designs by asking whether they preserve an appropriate relationship between source meaning and imposed structure.

8.3 Limitations and Future Work

This paper develops a conceptual diagnostic and procedural audit. Empirical work should test whether different analysts identify similar crisping pathways and whether the audit reveals risks that existing evaluation methods overlook. Studies should also examine whether the resulting findings lead to useful design changes.

The case analyses cover open-ended language, self-categorized input, and a preservation-oriented timeline system. Other settings may produce different forms of vagueness or distribute authority differently. The analyses also rely on an external-observer perspective. Public interfaces and documentation expose some transformations while concealing internal data models and system rules. Applying the audit with source users, designers, maintainers, and decision-makers would provide a fuller account of each pathway.

Judgments about adequate precision and recourse will vary across contexts. The audit makes those judgments explicit, but domain knowledge and institutional responsibility still shape the outcome. Future work should develop evaluation procedures that bring those perspectives into the audit while allowing disagreement to remain visible.

Many systems also require stable values. Premature crisping evaluates how systems reach and present those values. The threat arises when the system gives one interpretation greater authority than the source warrants and leaves affected actors with too little recourse.

9 Conclusion

This paper positions **premature crisping** as a specific diagnostic for the broader problem of visualizing vague, language-derived data. Premature crisping occurs when a system resolves semantic vagueness into a date, score, category, position, or computational output and makes that imposed precision difficult to distinguish from precision supplied by the source.

Using Munzner's nested model, I trace the threat across domain characterization, abstraction, visual encoding and interaction, and algorithm design. I then develop **interpretive recourse** and an accompanying audit for assessing whether relevant actors can recognize a resolution, reach the vague source, and respond when the stabilized meaning does not fit. Provenance supports that work when actors can use it at the point of need.

As visualization systems increasingly mediate between human expression and formal representation, the movement from language to structure deserves direct scrutiny. Systems will continue to resolve vague expressions because computation and coordination often require stable values. Those resolutions should retain an account of how they arose and remain open to review when the task makes that history consequential.

The issue ultimately concerns interpretive authority. Formalization determines which meanings become data and which interpretations guide action. Interpretive recourse reveals whether the people affected by those choices can still recognize and respond to them.

Acknowledgments

To Robert, for the bagels and explaining CMYK and color spaces.

References

[1] Jimmy Abualdenien and André Borrmann. 2020. Vagueness Visualization in Building Models across Different Design Stages. *Advanced Engineering Informatics* 45 (2020), 101107. doi:10.1016/j.aei.2020.101107

[2] Derya Akbaba, Lauren Klein, and Miriah Meyer. 2025. Entanglements for Visualization: Changing Research Outcomes through Feminist Theory. *IEEE Transactions on Visualization and Computer Graphics* 31, 1 (2025), 1279–1289. doi:10.1109/TVCG.2024.3456171

[3] Anonymous. 2026. HealthTale: A Patient-Centric Health Story Visualization Tool. (2026). Submitted for publication to IEEE Trans. Vis. Comput. Graph. (TVCG/IEEE VIS).

[4] Annie Bares, Daniel F Keefe, and Francesca Samsel. 2020. Close reading for visualization evaluation. *IEEE Computer Graphics and Applications* 40, 4 (2020), 84–95. doi:10.1109/MCG.2020.2993889

[5] Kirsten Boehner, Rogério DePaula, Paul Dourish, and Phoebe Sengers. 2005. Affect: from information to interaction. In *Proceedings of the 4th decennial conference on Critical computing: between sense and sensibility*. 59–68. doi:10.1145/1094562.1094570

[6] Kirsten Boehner, Rogério DePaula, Paul Dourish, and Phoebe Sengers. 2007. How emotion is made and measured. *International Journal of Human-Computer Studies* 65, 4 (2007), 275–291. doi:10.1016/j.ijhcs.2006.11.016

[7] Georges-Pierre Bonneau, Hans-Christian Hege, Chris R. Johnson, Manuel M. Oliveira, Kristin Potter, Penny Rheingans, and Thomas Schultz. 2014. Overview and State-of-the-Art of Uncertainty Visualization. In *Scientific Visualization: Uncertainty, Multifield, Biomedical, and Scalable Visualization*. Springer-Verlag, 3–27. doi:10.1007/978-1-4471-6497-5_1

[8] Nadia Boukhelifa, Marc-Emmanuel Perrin, Samuel Huron, and James Eagan. 2017. How Data Workers Cope with Uncertainty: A Task Characterisation Study. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, 3645–3656. doi:10.1145/3025453.3025738

[9] Matthew Brehmer and Tamara Munzner. 2013. A multi-level typology of abstract visualization tasks. *IEEE transactions on visualization and computer graphics* 19, 12 (2013), 2376–2385. doi:10.1109/TVCG.2013.124

[10] Ken Brodlie, Rodrigo Allendes Osorio, and Adriano Lopes. 2012. A Review of Uncertainty in Data Visualization. In *Expanding the Frontiers of Visual Analytics and Visualization*. Springer-Verlag, 81–109. doi:10.1007/978-1-4471-2804-5_6

[11] Peter Buneman, Sanjeev Khanna, and Tan Wang-Chiew. 2001. Why and Where: A Characterization of Data Provenance. In *International Conference on Database Theory*. Springer Berlin Heidelberg, Berlin, Heidelberg, 316–330. doi:10.1007/3-540-44503-X_20

[12] Stuart K. Card, Jock D. Mackinlay, and Ben Shneiderman. 1999. Information Visualization. In *Readings in Information Visualization: Using Vision to Think*, Stuart K. Card, Jock D. Mackinlay, and Ben Shneiderman (Eds.). Morgan Kaufmann, San Francisco, CA, 1–34.

[13] Sheelagh Carpendale, Alice Thudt, Charles Perin, and Wesley Willett. 2017. Subjectivity in Personal Storytelling with Visualization. *Inf. Des. J.* 23, 1 (2017), 48–64. doi:10.1075/idj.23.1.07thu

[14] Nan-Chen Chen, Margaret Drouhard, Rafal Kocielnik, Jina Suh, and Cecilia R Aragon. 2018. Using machine learning to support qualitative coding in social science: Shifting the focus to ambiguity. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 8, 2 (2018), 1–20. doi:10.1145/3185515

[15] Michael Correll and Michael Gleicher. 2014. Error Bars Considered Harmful: Exploring Alternate Encodings for Mean and Error. *IEEE Trans. Visualization & Computer Graphics (TVCG)* 20, 12 (2014), 2142–2151. doi:10.1109/TVCG.2014.2346298

[16] Catherine D'ignazio and Lauren F Klein. 2023. *Data feminism*. MIT press.

[17] Marian Dörk, Patrick Feng, Christopher Collins, and Sheelagh Carpendale. 2013. Critical InfoVis: exploring the politics of visualization. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. 2189–2198. doi:10.1145/2468356.2468739

[18] Margaret Drouhard, Nan-Chen Chen, Jina Suh, Rafal Kocielnik, Vanessa Pena-Araya, Keting Cen, Xiangyi Zheng, and Cecilia R Aragon. 2017. Aeonium: Visual analytics to support collaborative qualitative coding. In *2017 IEEE Pacific visualization symposium (PacificVis)*. IEEE, 220–229. doi:10.1109/PACIFICVIS.2017.8031598

[19] Cyntrica Eaton, Catherine Plaisant, and Terence Drizd. 2005. Visualizing missing data: Graph interpretation user study. In *IFIP Conference on Human-Computer Interaction*. Springer, 861–872. doi:10.1007/11555261_68

[20] Daniel A Epstein, An Ping, James Fogarty, and Sean A Munson. 2015. A lived informatics model of personal informatics. In *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing*. 731–742. doi:10.1145/2750858.2804250

[21] Kit Fine. 1975. Vagueness, Truth and Logic. *Synthese* 30, 3–4 (1975), 265–300. doi:10.1007/BF00485047

[22] Max Franke, Ralph Barczok, Steffen Koch, and Dorothea Weltecke. 2019. Confidence as First-Class Attribute in Digital Humanities Data. In *Proceedings of the 4th IEEE VIS Workshop on Visualization for the Digital Humanities (VIS4DH)*.

[23] Barney G. Glaser and Anselm L. Strauss. 1998. Grounded Theory. In *Strategien Qualitativer Forschung*. Huber, Bern, 339–385. doi:10.1007/978-3-030-85124-8_9

[24] Yashvir Grewal, Sarah Goodwin, and Tim Dwyer. 2021. Visualising Temporal Uncertainty: A Taxonomy and Call for Systematic Evaluation. In *IEEE Symp. Pacific Visualization (PacificVis)*. 41–45. doi:10.1109/PacificVis52677.2021.00013

[25] Habitics. 2026. Daylio: Journal, Diary and Mood Tracker. https://daylio.net/. Accessed 2026-05-27.

[26] Pantea Haghighatkhah, Mennatallah El-Assady, Jean-Daniel Fekete, Narges Mahyar, Carita Paradis, Vasiliki Simaki, and Bettina Speckmann. 2022. Characterizing Uncertainty in the Visual Text Analysis Pipeline. In *IEEE VIS Workshop on Visualization for the Digital Humanities (VIS4DH)*. 25–30. doi:10.1109/VIS4DH57440.2022.00010

[27] Jerry R. Hobbs. 1990. Granularity. In *Readings in Qualitative Reasoning about Physical Systems*. Elsevier, 542–545. doi:10.1016/B978-1-4832-1447-4.50053-5

[28] Jessica Hullman. 2020. Why Authors Don't Visualize Uncertainty. *IEEE Trans. Visualization & Computer Graphics (TVCG)* 26, 1 (2020), 130–139. doi:10.1109/TVCG.2019.2934287

[29] Dominic Hyde. 2011. The Sorites Paradox. In *Vagueness: A Guide*. Springer-Verlag, 1–17. doi:10.1007/978-94-007-0375-9_1

[30] Stefan Jänicke, Greta Franzini, Muhammad Faisal Cheema, and Gerik Scheuermann. 2017. Visual Text Analysis in Digital Humanities. *Computer Graphics Forum (CGF)* 36, 6 (2017), 226–250. doi:10.1111/cgf.12873

[31] Stefan Jänicke, Greta Franzini, Muhammad Faisal Cheema, Gerik Scheuermann, et al. 2015. On Close and Distant Reading in Digital Humanities: A Survey and Future Challenges. *Computer Graphics Forum (CGF)* 34, 3 (2015), 83–103.

[32] Sean Kandel, Jeffrey Heer, Catherine Plaisant, Jessie Kennedy, Frank Van Ham, Nathalie Henry Riche, Chris Weaver, Bongshin Lee, Dominik Brodbeck, and Paolo Buono. 2011. Research Directions in Data Wrangling: Visualizations and Transformations for Usable and Credible Data. *Information Visualization* 10, 4 (2011), 271–288. doi:10.1177/1473871611415994

[33] Sean Kandel, Andreas Paepcke, Joseph Hellerstein, and Jeffrey Heer. 2011. Wrangler: Interactive Visual Specification of Data Transformation Scripts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 3363–3372. doi:10.1145/1978942.1979444

[34] Christopher Kennedy. 2007. Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives. *Linguist. Philos.* 30, 1 (2007), 1–45. doi:10.1007/s10988-006-9008-0

[35] George Lakoff. 1973. Hedges: A Study in Meaning Criteria and the Logic of Fuzzy Concepts. *J. Philos. Logic* 2, 4 (1973), 458–508. doi:10.1007/BF00262952

[36] Peter Lasersohn. 1999. Pragmatic Halos. *Language* 75, 3 (1999), 522–551. doi:10.2307/417059

[37] Ian Li, Anind Dey, and Jodi Forlizzi. 2010. A stage-based model of personal informatics systems. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 557–566. doi:10.1145/1753326.1753409

[38] Haihan Lin, Derya Akbaba, Miriah Meyer, and Alexander Lex. 2023. Data Hunches: Incorporating Personal Knowledge into Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2023), 504–514. doi:10.1109/TVCG.2022.3209451

[39] Alan M. MacEachren, Anthony Robinson, Susan Hopper, Steven Gardner, Robert Murray, Mark Gahegan, and Elisabeth Hetzler. 2005. Visualizing Geospatial Information Uncertainty: What We Know and What We Need to Know. *Cartography and Geographic Information Science* 32, 3 (2005), 139–160. doi:10.1559/1523040054738936

[40] Sana Malik, Fan Du, Megan Monroe, Eberechukwu Onukwugha, Catherine Plaisant, and Ben Shneiderman. 2015. Cohort Comparison of Event Sequences with Balanced Integration of Visual Analytics and Statistics. In *Proceedings of the 20th International Conference on Intelligent User Interfaces (IUI '15)*. ACM, 38–49. doi:10.1145/2678025.2701407

[41] Nina McCurdy, Julie Gerdes, and Miriah Meyer. 2019. A Framework for Externalizing Implicit Error Using Visualization. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 925–935. doi:10.1109/TVCG.2018.2864913

[42] Nina McCurdy, Julie Lein, Katharine Coles, and Miriah Meyer. 2016. Poemage: Visualizing the Sonic Topology of a Poem. *IEEE Trans. Visualization & Computer Graphics (TVCG)* 22, 1 (2016), 439–448. doi:10.1109/TVCG.2015.2467811

[43] Sean McKenna, Dominika Mazur, James Agutter, and Miriah Meyer. 2014. Design activity framework for visualization design. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2191–2200. doi:10.1109/TVCG.2014.2346331

[44] Miriah Meyer and Jason Dykes. 2019. Criteria for rigor in visualization design study. *IEEE transactions on visualization and computer graphics* 26, 1 (2019), 87–97. doi:10.1109/TVCG.2019.2934539

[45] Megan Monroe, Rongjian Lan, Juan Morales del Olmo, Ben Shneiderman, Catherine Plaisant, and Jeff Millstein. 2013. The Challenges of Specifying Intervals and Absences in Temporal Queries: A Graphical Language Approach. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, 2349–2358. doi:10.1145/2470654.2481325

[46] Michael Muller, Ingrid Lange, Dakuo Wang, David Piorkowski, Jason Tsay, Q. Vera Liao, Casey Dugan, and Thomas Erickson. 2019. How Data Science Workers Work with Data: Discovery, Capture, Curation, Design, Creation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, 1–15. doi:10.1145/3290605.3300356

[47] Michael Muller, Christine T Wolf, Josh Andres, Michael Desmond, Narendra Nath Joshi, Zahra Ashktorab, Aabhas Sharma, Kristina Brimijoin, Qian Pan, Evelyn Duesterwald, et al. 2021. Designing ground truth and the social life of labels. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–16. doi:10.1145/3411764.3445402

[48] Tamara Munzner. 2009. A nested model for visualization design and validation. *IEEE transactions on visualization and computer graphics* 15, 6 (2009), 921–928.

[49] Stan Nowak and Lyn Bartram. 2024. Designing for Ambiguity in Visual Analytics: Lessons from Risk Assessment and Prediction. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2024), 924–933. doi:10.1109/TVCG.2023.3326571

[50] Stan Nowak, Lyn Bartram, and Pascal Haegeli. 2020. Designing for Ambiguity: Visual Analytics in Avalanche Forecasting. In *2020 IEEE Visualization Conference (VIS)*. IEEE, 81–85. doi:10.1109/VIS47514.2020.00023

[51] Georgia Panagiotidou and Andrew Vande Moere. 2022. Communicating Qualitative Uncertainty in Data Visualization: Two Cases from Within the Digital Humanities. *Inf. Des. J.* 27, 1 (2022), 52–63. doi:10.1075/idj.22014.pan

[52] Ellie Pavlick and Tom Kwiatkowski. 2019. Inherent disagreements in human textual inferences. *Transactions of the Association for Computational Linguistics* 7 (2019), 677–694. doi:10.1162/tacl_a_00293

[53] Barbara Plank. 2022. The "problem" of human label variation: On ground truth in data, modeling and evaluation. In *Proceedings of the 2022 conference on empirical methods in natural language processing*. 10671–10682. doi:10.18653/v1/2022.emnlp-main.731

[54] Kristin Potter, Paul Rosen, and Chris R. Johnson. 2011. From Quantification to Visualization: A Taxonomy of Uncertainty Visualization Approaches. In *IFIP Working Conf. Uncertainty Quantification*. 226–249. doi:10.1007/978-3-642-32677-6_15

[55] Qualtrics. 2026. Text iQ Documentation. https://www.qualtrics.com/support/survey-platform/data-and-analysis-module/text-iq/. Accessed 2026-05-27.

[56] James A Russell. 1980. A circumplex model of affect. *Journal of personality and social psychology* 39, 6 (1980), 1161. doi:10.1037/h0077714

[57] Michael Sedlmair, Miriah Meyer, and Tamara Munzner. 2012. Design study methodology: Reflections from the trenches and the stacks. *IEEE transactions on visualization and computer graphics* 18, 12 (2012), 2431–2440. doi:10.1109/TVCG.2012.213

[58] Claude E. Shannon. 1948. A Mathematical Theory of Communication. *Bell System Technical Journal* 27, 3 (jul 1948), 379–423. doi:10.1002/j.1538-7305.1948.tb01338.x

[59] Stephanie Solt. 2015. Vagueness and Imprecision: Empirical Foundations. *Annu. Rev. Linguist.* 1, 1 (2015), 107–127. doi:10.1146/annurev-linguist-030514-125150

[60] Hayeong Song and Danielle Albers Szafir. 2018. Where's my data? evaluating visualizations with missing data. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 914–924. doi:10.1109/TVCG.2018.2864914

[61] Jianhua Tao and Tieniu Tan. 2005. Affective computing: A review. In *International Conference on Affective computing and intelligent interaction*. Springer, 981–995. doi:10.1007/11573548_125

[62] Judi Thomson, Elizabeth Hetzler, Alan MacEachren, Mark Gahegan, and Misha Pavel. 2005. A Typology for Visualizing Uncertainty. In *Visualization and Data Analysis 2005*, Vol. 5669. SPIE, 146–157. doi:10.1117/12.587254

[63] Barbara Tversky. 2013. Visualizing Thought. In *Handbook of Human Centric Visualization*. Springer New York, New York, NY, 3–40. doi:10.1007/978-1-4614-7485-2_1

[64] Timothy Williamson. 1992. Vagueness and Ignorance. *Proc. Aristotelian Soc., Suppl. Vol.* 66 (1992), 145–177.

[65] F. Windhager, V. A. Filipov, S. Salisu, and E. Mayr. 2018. Visualizing Uncertainty in Cultural Heritage Collections. In *EuroVis Workshop on Reproducibility, Verification, and Validation in Visualization*. The Eurographics Association, 7–11. doi:10.2312/eurorv3.20181142

[66] Ji Soo Yi, Youn ah Kang, John Stasko, and Julie A Jacko. 2007. Toward a deeper understanding of the role of interaction in information visualization. *IEEE transactions on visualization and computer graphics* 13, 6 (2007), 1224–1231. doi:10.1109/TVCG.2007.70515

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009