

# Narrative Scaffolding: A Narrative-First Framework for Data-Driven Sensemaking

Oliver Huang  
 Department of Computer Science  
 University of Toronto  
 Toronto, Ontario, Canada  
 oliver@dgp.toronto.edu

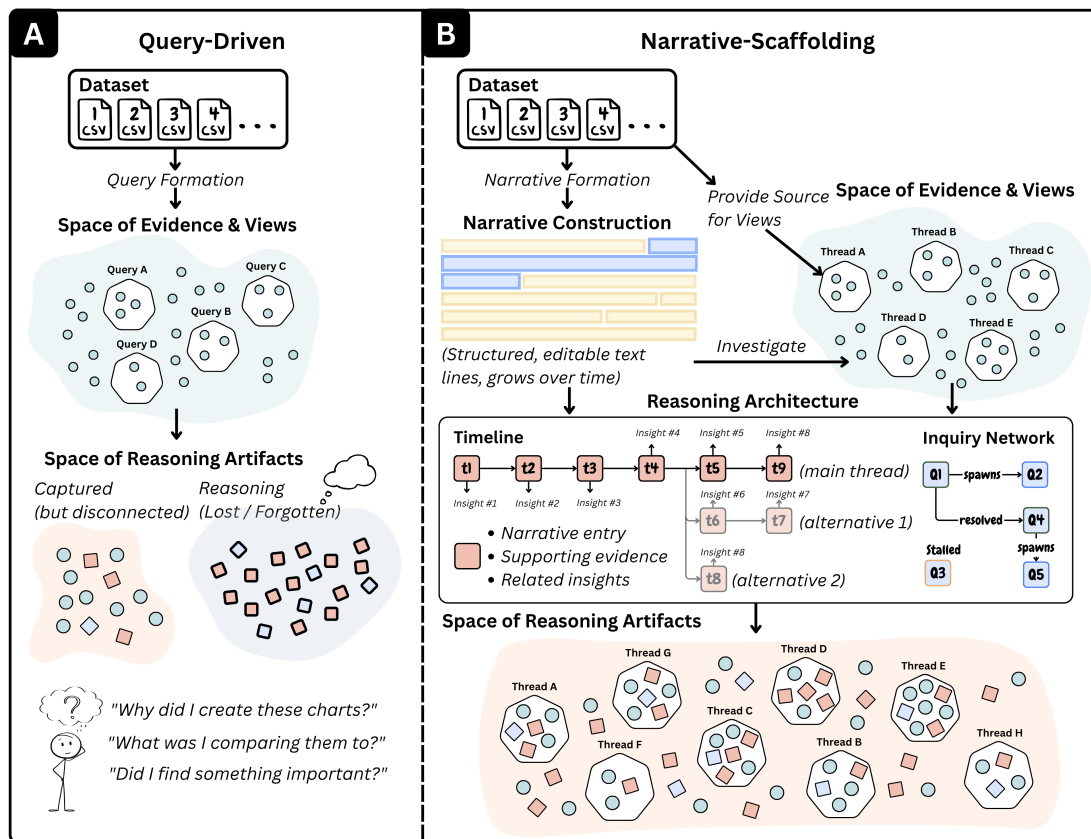
Muhammad Fatir  
 Department of Computer Science  
 University of Toronto  
 Toronto, Ontario, Canada  
 m.fatir@mail.utoronto.ca

Tianyu Luo  
 Department of Computer Science  
 University of Toronto  
 Toronto, Ontario, Canada  
 steventian.luo@mail.utoronto.ca

Sangho Suh  
 The Allen Institute for AI  
 Seattle, Washington, USA  
 sanghos@allenai.org

Hariharan Subramonyam  
 Stanford University  
 Stanford, California, USA  
 harihars@stanford.edu

Carolina Nobre  
 University of Toronto  
 Toronto, Ontario, Canada  
 cnobre@cs.toronto.edu



**Figure 1: Narrative Scaffolding is a framework for data-driven sensemaking that treats narrative construction as the primary interface for exploration and reasoning. Unlike traditional query-driven analysis (A), where users formulate queries to generate disconnected views and reasoning remains implicit, Narrative Scaffolding (B) externalizes evolving reasoning through structured narrative text. As users write, the system generates semantically aligned visualizations, captures insights with provenance, and tracks inquiry evolution. All reasoning artifacts (narrative entries, supporting evidence, and inquiry networks) are organized within a single, connected space, enabling users to revisit, branch, and refine their thinking throughout analysis.**

## Abstract

When exploring data, analysts construct narratives about what the data means by asking questions, generating visualizations, reflecting on patterns, and revising their interpretations as new insights emerge. Yet existing analysis tools treat narrative as an afterthought, breaking the link between reasoning, reflection, and the evolving story from exploration. Consequently, analysts lose the ability to see how their reasoning evolves, making it harder to reflect systematically or build coherent explanations. To address this gap, we propose Narrative Scaffolding (NS), a framework for narrative-driven exploration that positions narrative construction as the primary interface for exploration and reasoning. We implemented this framework in a system that externalizes iterative reasoning through narrative-first entry, semantically aligned view generation, and reflection support via insight provenance and inquiry tracking. In a within-subject study ( $N = 20$ ), we demonstrated that narrative scaffolding facilitates broader exploration, deeper reflection, and more defensible narratives. An evaluation with visualization literacy experts ( $N = 6$ ) confirmed that the system produced outputs aligned with narrative intent and facilitated intentional exploration.

## CCS Concepts

• **Human-centered computing** → **Visual analytics; Interactive systems and tools; Natural language interfaces.**

## Keywords

Narrative-Driven Exploration, Narrative Scaffolding, Cognitive Offloading, Visual Analytics, Data Storytelling, AI-Assisted Analytics, Conversational Data Analysis, Insight Provenance, Natural Language Interfaces, Analytical Reasoning

### ACM Reference Format:

Oliver Huang, Muhammad Fatir, Tianyu Luo, Sangho Suh, Hariharan Subramonyam, and Carolina Nobre. 2026. Narrative Scaffolding: A Narrative-First Framework for Data-Driven Sensemaking. In *31st International Conference on Intelligent User Interfaces (IUI '26)*, March 23–26, 2026, Paphos, Cyprus. ACM, New York, NY, USA, 29 pages. <https://doi.org/10.1145/3742413.3789062>

## 1 Introduction

Exploring data is fundamentally a sensemaking process that involves practicing *interpretive reasoning* throughout: in which analysts continuously form and revise assumptions in response to evidence [6, 23]. Analysts rarely begin from blank slates [39, 102].



This work is licensed under a Creative Commons Attribution 4.0 International License. *IUI '26, Paphos, Cyprus*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1984-4/26/03

<https://doi.org/10.1145/3742413.3789062>

Instead, they arrive with assumptions about the data (e.g., “*economic growth drives housing demand*”) that shape initial inquiry, such as causal hypotheses (e.g., “*higher income areas will show greater price stability*”), exploratory questions (e.g., “*which boroughs deviate from this pattern?*”), or domain-driven expectations (e.g., “*urban centers tend to recover faster after downturns*”) [18, 30, 74]. In other words, data exploration entails a cycle of *interpretive reasoning*: analyzing information to extract meaning, questioning or refining these assumptions, and developing evidence-based explanations that guide the course of exploration.

However, this cycle of *interpretive reasoning* often remains implicit [15, 65]. Analysts may hold assumptions in memory, pursue lines of inquiry without articulating their rationale, or shift interpretations without documenting the changes that led to them [51]. When reasoning stays internal, assumptions can harden while still being unverified. Analysts may explore too narrowly, following a single line of inquiry without considering alternatives, or they may selectively seek evidence that confirms initial hunches while overlooking contradictions [63, 96]. As exploration deepens, the cognitive burden increases: earlier insights are forgotten, the rationale behind interpretive shifts fades, and promising questions get abandoned without resolution. Without mechanisms to externalize and revisit their evolving thinking, analysts struggle to maintain awareness of how their interpretations have developed, what alternatives they’ve considered, or which questions remain unresolved.

Constructing narrative offers a structure for making this interpretive reasoning explicit. By structuring questions, interpretations, and evidence into narrative form, internal thought becomes tangible statements that can be inspected, revised, and linked to supporting evidence [53, 58, 77]. This externalization creates opportunities for reflection, allowing analysts to trace how interpretations evolved, compare competing explanations, and identify gaps in their reasoning [34, 82].

However, most analytic tools do not support analysts’ narrative as a scaffold for *interpretive reasoning*. They commonly treat narrative as a static output, summarizing finished actions or presenting isolated facts [41, 103]. Even with recent advances in AI-assisted data analysis that integrate semantics, natural language, and visualization to support more human-centered exploration [18, 24, 78, 80], existing systems remain largely reactive: generating visualizations on demand, assembling analysts’ narratives retrospectively from logs [52, 97, 105], and responding to explicit prompts rather than participating in meaning construction as it unfolds [17, 49, 54]. This keeps exploration (data interaction) separate from interpretation (meaning construction) [4, 93], leaving interpretive reasoning invisible and unsupported [60, 91]. While provenance systems document interaction histories, they capture what analysts do but not why.

They lack the narrative logic that connects observations into coherent interpretations [65, 73]. Without structured externalization, analysts often face difficulties engaging in systematic reflection [15, 94], which can lead to premature conclusions or leave assumptions unexamined [50, 63, 96]. When reasoning remains implicit, decisions lack the traceability needed for reproducibility and accountability [21, 73], which are crucial in fields like journalism [19] and intelligence analysis [75]. Ultimately, current approaches limit the ability to share, validate, and build upon insights [55, 101].

This work addresses these shortcomings by exploring a new interaction paradigm that scaffolds *interpretive reasoning* as it unfolds. By treating narrative construction as the primary interface of the sensemaking process, our approach makes evolving assumptions explicit, surfaces contradictions, and preserves reasoning paths that are otherwise lost. To achieve this, we draw on insights from theories in narrative psychology [58], external cognition [53, 77], and constructionist learning [3, 27], which collectively emphasize how externalized narratives can serve both as a cognitive medium for reasoning and as an artifact that preserves its evolution.

Taken together, grounded in these perspectives, we propose *Narrative Scaffolding*: a framework that repositions narrative construction as the central interface for data exploration and reasoning. Unlike traditional tools that log low-level interactions or assemble stories retrospectively, *Narrative Scaffolding* treats interpretation as a dynamic, evolving process that must be externalized and structured to support systematic reasoning. We instantiate this framework in a system that enables analysts to write narrative sentences as first-class analytic actions, with each sentence bidirectionally linked to visual evidence and tracked questions. The system generates semantically-aligned visualizations from natural language, captures insight provenance through a navigable timeline, tracks inquiry evolution through an interactive board, and detects opportunities for branching when reasoning diverges.

In summary, our contributions include:

- **Conceptual framework:** A *Narrative Scaffolding* framework derived from a formative study that positions narrative structure as the backbone of sensemaking, linking high-level goals, evolving hypotheses, insights, and visual evidence. It emphasizes ongoing reflection, integration of new insights, and preservation of reasoning paths throughout exploration.
- **System instantiation:** An interactive system that operationalizes this framework by making narrative construction the central interface through which analysts pose questions, connect evidence, and develop explanations.
- **Empirical evidence:** Results from a two-part evaluation: (1) a within-subject study (N=20) demonstrating that *Narrative Scaffolding* improves workflow efficiency, supports deeper reflection, encourages broader exploration, and produces more coherent and defensible narratives compared to reactive AI-assisted analysis; and (2) a system-level evaluation (N=6) examining the quality and reliability of AI-generated outputs across diverse analytical scenarios.

## 2 Related Work

Understanding how to support interpretive reasoning during data exploration requires examining both how people naturally make sense of information and how existing systems attempt to capture this process. We review prior work through three lenses: the cognitive role of narrative in reasoning, system approaches to capturing analytical provenance and interpretation, and how LLM integration shapes these processes through mixed-initiative interaction.

### 2.1 Narrative in Interpretive Reasoning

Narrative psychology posits that humans are fundamentally wired to make sense of information through constructing coherent narratives to connect events, observations, and meaning [58, 59]. This narrative mode of thought serves not merely as a communication device but as a cognitive tool for organizing experience and building understanding. Sensemaking research supports this view, showing how people iteratively generate and revise interpretations as they build mental models [43, 70], structuring their understanding around storylines that link goals, observations, and conclusions [14].

This narrative mode of reasoning manifests across diverse analytical contexts. In journalism, reporters develop story angles that guide their investigation and evidence gathering [1]. In scientific inquiry, researchers frame hypotheses narratively, iterating on explanatory models as new data emerge [13, 68]. In data analysis, analysts similarly develop and iterate on storylines for organizing complex information and guiding exploratory decisions [19, 34, 82]. Across these domains, narrative functions as an active reasoning scaffold rather than merely a final output format.

Recent empirical studies validate this narrative-driven reasoning in data exploration. Even without explicit tool support, users spontaneously construct storylines during analysis [82]. Observational research shows analysts following interpretive arcs, revisiting assumptions as evidence emerge [34], and beginning with specific narratives that guide their data exploration [18]. These findings confirm that reasoning unfolds through storyline construction during active exploration, not merely during final communication [84, 105].

Building on these theoretical foundations, our narrative scaffolding framework treats storyline construction as the primary reasoning interface throughout exploration, supporting narrative writing as an active, evolving process of sensemaking.

### 2.2 Capturing Analytic Reasoning

Supporting analytic reasoning requires capturing not just what analysts do, but why they do it and how their understanding evolves. Provenance tracking addresses this by recording the history of how analytical conclusions were reached, capturing the sequence of actions, decisions, and intermediate states that led to the final result [29]. In nonlinear, iterative exploration, provenance provides a structured record supporting reproducibility, error recovery, collaboration, and accountability [10, 21, 32, 55, 65, 73]. However, achieving these goals depends critically on what information can actually be captured.

Ragan et al. [73] proposed a foundational framework categorizing provenance into five types: data, visualization, interaction, insight, and rationale. While the first three can be captured through

system instrumentation, the latter two involve internal reasoning that is inherently more challenging to observe [15, 64, 101]. Current systems excel at the observable levels. Workflow-based systems like VisTrails [9] model analysis steps in structured pipelines, though they impose rigid workflows less suited to open-ended exploration. Interactive tracking systems [12, 22, 29] dynamically log actions and states, while recent approaches like ProvenanceWidgets [61] and ViStruct [32] embed tracking directly into interfaces. However, these systems primarily capture observable behaviors while providing limited support for representing underlying reasoning.

To capture insight-level provenance, systems like Calliope [81] support explicit argument construction by linking claims to data evidence, while CZSaw [36] enables analysts to annotate findings and attach rationale to visual states. However, these approaches require users to explicitly articulate insights at specific moments, creating friction between ongoing exploration [73]. Insights are often documented as static snapshots instead of representing the dynamic nature of evolving thought processes [65].

This leaves a fundamental gap in capturing interpretive reasoning as it unfolds. Provenance systems excel at documenting what users do but struggle with why they do it. This interpretive evolution (the narrative construction of meaning) is central to how analysts actually reason. When systems cannot capture or support this narrative evolution, reasoning processes remain implicit or fragmented, making it challenging to revisit rationale and trace interpretive evolution [64, 84, 94]. Our approach addresses this gap by enabling systems to capture insight-level provenance as reasoning naturally unfolds through writing.

### 2.3 LLM-Assisted Data-Driven Sensemaking

Complementary to narrative storyline construction and systematic provenance tracking, mixed-initiative analysis systems address interpretive reasoning through human-AI collaboration [18, 46, 52, 97]. These systems shift from passive tools to collaborative partners that generate views and respond to natural language queries during exploration.

Current approaches can be categorized into three main interaction paradigms. Conversational interfaces enable dialogue-based exploration where users pose questions and refine queries through natural language [20, 62, 67, 79, 89]. Structured reasoning systems support explicit argument construction through claim-evidence linking and branching logic paths [35, 81], while text-visualization integration approaches facilitate coherence between textual narratives and visual elements [46, 83, 88, 106]. A key affordance is that text serves as a more expressive interaction medium, enabling iterative communication between users and systems that goes beyond traditional point-and-click interfaces.

A significant focus has emerged on AI-assisted content generation during analysis. These systems leverage LLMs to automatically produce summaries, story fragments, or visual descriptions based on user actions and analytic traces [18, 52, 90, 97, 99]. The effectiveness of these content generation approaches depends on factors such as the sophistication of natural language processing, the quality of AI-generated insights, and the seamless integration of human input with system responses [48]. Our framework shifts

from generating narratives reactively to scaffolding them proactively, treating narrative construction as the primary interface for reasoning throughout exploration.

## 3 Formative Studies

To understand how interpretive reasoning unfolds during data exploration and where current tools create friction, we conducted a formative study that organized these issues into four interpretive challenges. We surveyed 48 analysts to validate the prevalence of challenges and conducted follow-up interviews with 5 participants to examine how reasoning strategies evolve. A preliminary version of these findings was reported in [31]. Below, we detail the procedure for survey and interview studies and describe the identified challenges that informed our design goals.

### 3.1 Procedure

**3.1.1 Formative Survey ( $N = 48$ ).** We surveyed 48 data practitioners ( $F_1-F_{48}$ ) who frequently engage in data exploration. Participants were recruited through Prolific [71] and compensated at the rate of \$16 USD/hr. The sample (26 male, 22 female, ages 21–55) included data analysts ( $n = 16$ ), managers ( $n = 7$ ), healthcare professionals ( $n = 6$ ), and sales consultants ( $n = 3$ ). This ensured diversity across professional domains, capturing a broad range of reasoning contexts. The survey combined Likert items with semi-structured open-ended prompts, probing how participants initiated exploration, adapted their thinking when encountering unexpected findings, managed analytical complexity, and revisited prior insights. (see supplementary materials for the full protocol)

Responses revealed common behavioral patterns: managing multiple exploration paths through workarounds, externalizing reasoning outside the analytic environment, and restarting from scratch when interpretations diverged. To further investigate how interpretive reasoning unfolds over time and what triggers specific analytical decisions, we conducted follow-up interviews with five participants.

**3.1.2 Follow-up Interview ( $N = 5$ ) and Challenge Formation.** We recruited five participants from the survey ( $F'_1-F'_5$ ) who have experience in narrative-driven analytical workflows and typically using mixed-initiative analytic tools. Each 30–45 minute interview ( $M=39.2$ ) was compensated at a rate of \$20 USD/hr and followed a semi-structured protocol where participants walked through recent analysis sessions, describing how they initiated exploration, articulated evolving goals, and managed reasoning shifts in response to new evidence. (See supplementary materials for the full protocol.)

The interviews revealed rich insights which directly informed our design goals. Open-ended survey responses and interview transcripts were analyzed using inductive thematic analysis [66]. Table 1 summarizes the four identified challenges (C1–C4), observed behaviors (B1–B9), and corresponding design goals (DG1–DG4).

### 3.2 Challenges

**3.2.1 Challenge 1: Unsupported Transition from Intent to Evidence.** Prior work shows that analysts often begin exploration with tentative or underspecified intentions, yet most analytic systems demand premature precision or rigid query formulation [29, 70, 73]. In our survey, 39 of 48 participants reported that

**Table 1: Summary of user challenges, observed barriers, design goals, and system features addressing interpretive reasoning in data exploration.**

Challenges	Observed Barriers	Design Goals	System Features
C1: Unsupported transition from exploration intent to evidence (Sec 3.2.1)	B1: Exploration intentions lack detail	DG1: Support fluid transitions from vague inquiries to evidence-grounded insights	Semantic-aligned views
	B2: Generated views become disconnected		Bidirectional narrative construction
	B3: Reliance on external note tools		
C2: Unmanaged inquiry and reflective continuity (Sec 3.2.2)	B4: Interpretations not explicitly captured	DG2: Externalize inquiry intent and maintain continuity of reflection	Insight timeline
	B5: Limited reflective review		
	B6: Reasoning paths left unresolved		Inquiry board
C3: Disrupted divergence in exploratory reasoning (Sec 3.2.3)	B7: Restarts instead of branching	DG3: Support nonlinear exploration and alternative reasoning	Narrative branching
	B8: Hard to compare alternative paths		
C4: Unstructured notes lacking formalization (Sec 3.2.4)	B9: Difficult synthesis into analysis	DG4: Support structured and defensible narratives	Story Integration

their early exploration goals lacked detail (B1).  $F_{12}$  reflected that *"It gives me something too specific when I'm still just exploring"*. This mismatch limits how users pursue their curiosity or flexibly define intentions.

As exploration progressed, 41 of 48 reported that generated views became disconnected from their evolving questions (B2).  $F_{21}$  explained, *"I take screenshots when something surprises me, but later I forget why it mattered."* Without mechanisms to anchor interpretive context to evidence, visualizations quickly lose their meaning, making them less valuable to revisit.

To manage the situation, participants (36 of 48) offloaded their reasoning into external note-taking tools, such as slides or documents (B3).  $F_6$  noted, *"I'll write down an idea like 'maybe this drop is due to [factor],' but that note just lives separately from the data."* These workarounds highlight how reasoning fragments drift apart from evidence, creating temporal degradation. As  $F'_2$  reflected, *"I usually remember my conclusions, but not the reasoning that led me there."* This illustrates how outcomes are retained while the reasoning behind them fades, making it harder to revisit or refine past analyses.

**3.2.2 Challenge 2: Unmanaged Inquiry and Reflective Continuity.** As exploration unfolded, many interpretations that participants formed were never made explicit within their analytic environment. From the survey, 44 of 48 participants acknowledged that their reasoning often remained undocumented during the process (B4).  $F_{14}$  explained, *"I usually have an idea in my head, if I don't write it down, it just disappears as I move forward."* When interpretations are not explicitly captured, they often fade from memory as analysis progresses, which limits the continuity of inquiry and weakens opportunities for reflection. As  $F'_3$  explained, *"I sometimes can't get back to what I thought earlier, so I just move forward instead."*

Even when participants attempted to revisit their earlier reasoning, the process was limited and inconsistent. 32 of 48 participants

reported rarely reviewing past conclusions (B5).  $F'_5$  described, *"I sometimes go back to see if my conclusion still makes sense, but it's hard to remember the context of why I thought it."* Without preserved context, retrospective reflection becomes fragmented and difficult to sustain [33, 37].

Finally, reasoning paths often ended unexpectedly as participants shifted focus. 25 of 48 participants acknowledged abandoning questions without resolution (B6). As  $F'_2$  noted, *"I left one question midway, but I never went back to see if it was still valid."* This lack of continuity results in abandoned questions and forgotten assumptions, making it difficult to trace how reasoning evolved.

**3.2.3 Challenge 3: Disrupted Divergence in Exploratory Reasoning.** Most analytic tools record actions in a single linear sequence [21, 22], but reasoning often branches as analysts test alternatives [65, 70]. Participants (29 of 48) reported restarting or duplicating their workspace when their thinking shifted (B7).  $F_{21}$  explained, *"If I want to try a new idea, I just duplicate the current exploration path and restart another, changing the original feels risky."* Restarting in this way creates a break in continuity: earlier interpretations are severed from the new line of thought, forcing participants to *"manage multiple versions manually"* ( $F_{43}$ ).

Prior work has shown that a narrative-driven approach can inadvertently reinforce existing assumptions when alternative explanations are not explored [50, 63, 96]. Participants (33 of 48) emphasized the importance of considering alternative explanations to maintain a balanced approach to their reasoning. However, a further 23 of 48 described difficulty keeping track of multiple possible directions when their reasoning diverged (B8).  $F'_4$  reflected, *"When I try two different parallel explanations, I usually lose track of one of them."* Without support to manage alternatives side by side, users struggled to weigh competing possibilities, which could limit consideration of disconfirming evidence.

**3.2.4 Challenge 4: Unstructured Notes Lacking Analytical Formalization.** While participants extensively captured interpretive thoughts during the exploration, these insights remained fragmented and unstructured, making them difficult to reuse or communicate effectively. Most participants (46 of 48) reported difficulty reconstructing their analytical logic from informal notes and scattered observations (B9).  $F_{10}$  noted, "I sometimes conclude quickly, and later I can't remember why that was in the first place." Others described how their reasoning fragments lacked structure for meaningful reuse.  $F_1$  explained, "I leave comments in documents, but they never translate back into the final analysis itself."

This fragmentation creates a temporal sensemaking problem where insights lose their analytical context over time.  $F_5$  remarked, "I jot things down as I go, but those notes don't connect to show how my thinking evolved." The challenge lies not merely in communication to stakeholders, but in preserving the analytical coherence needed for iterative reasoning and insight building.

### 3.3 Informed Design Goals

Based on the outlined challenges (C1–C4) from our formative study, we developed four design goals (DG1–DG4) to address the observed barriers (B1–B9). These goals informed our system architecture and aim to scaffold *interpretive reasoning* by externalizing evolving insights, supporting reflection, and enabling flexible narrative exploration.

**DG1: Support fluid transitions from vague inquiries to evidence-grounded insights.** To address the difficulty of moving from tentative goals into evidence-grounded analysis (C1), the system should provide semantic-aligned view generation that transforms underspecified starting points into meaningful visualizations (B1). It should also support linked authoring of narrative evidence to keep interpretive insights embedded in visual context rather than separated into external notes (B2, B3), enabling analysts to revisit earlier thoughts with continuity.

**DG2: Externalize inquiry intent and maintain continuity of reflection.** To address undocumented interpretations and abandoned questions (C2), the system should provide insight-level provenance that records reasoning alongside actions, allowing users to retrace prior thoughts (B4, B5). An inquiry board should maintain continuity by keeping unresolved questions visible and available for later return (B6), reducing reliance on memory alone.

**DG3: Support nonlinear exploration and alternative reasoning.** To enable divergent reasoning without breaking continuity (C3), the system should support narrative branching and direct comparison of alternatives (B7, B8). By making multiple reasoning paths visible and manageable, the system helps analysts systematically consider competing explanations and broaden their exploration scope.

**DG4: Support structured and defensible narratives.** To bridge scattered notes and informal observations into coherent outputs (C4), the system should scaffold narrative construction by maintaining evidence linkages throughout exploration (B9). This ensures final outputs remain defensible and reduces the effort required to reorganize materials for communication.

## 4 Narrative Scaffolding Framework

### 4.1 Stages and Theoretical Grounding

To connect our established design goals (DG1–DG4) to actual analytic practice — and use it to ground the design of our system (presented in Sec 5), we propose a conceptual model that captures how narrative construction unfolds through distinct stages of *interpretive reasoning*. This model provides a structured view of the cognitive processes our system is designed to support. The model was adapted from two established perspectives on analytic reasoning: the Sensemaking Loop [70], which describes an iterative cycle between searching for information and organizing it into mental structures, and the Data–Frame model [42], which emphasizes how analysts construct, test, and revise interpretive frames in response to new evidence. We also draw on external cognition theory [77] and narrative psychology [58] to ground how each stage supports reasoning. Building on these foundations, as shown in Fig. 2, our adaptation repositions narrative construction as the primary reasoning interface, and maps four stages to our design goals.

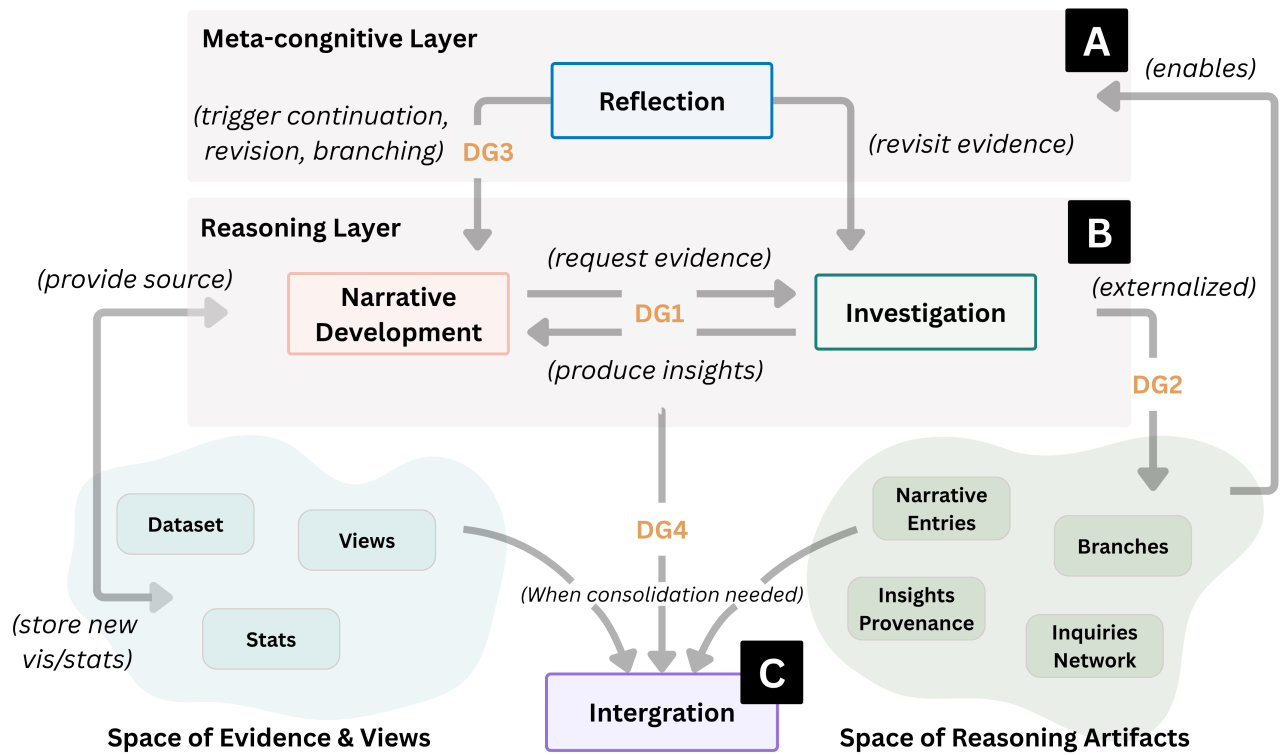
- **Narrative Development:** Externalizing assumptions as narrative entries that can be revised or branched.
- **Investigation:** Using narrative claims to generate evidence, and using evidence to reshape claims.
- **Reflection:** Revisiting explicit records of reasoning to check coherence, contradictions, and alternatives.
- **Integration:** Consolidating dispersed reasoning into provisional but defensible accounts.

Each stage operationalizes specific theoretical constructs. **Narrative Development** draws on external cognition theory: by externalizing assumptions as written text, analysts transform fleeting thoughts into manipulable artifacts that can be inspected and revised [53, 77]. **Investigation** operationalizes the Data–Frame model's emphasis on testing interpretive frames against evidence, where claims and data mutually reshape one another [42]. **Reflection** embodies the Sensemaking Loop's iterative cycling between bottom-up evidence gathering and top-down schema revision [70]. **Integration** draws on narrative psychology's insight that constructing coherent narratives is itself a mode of reasoning, not merely a post-hoc communication step [8, 58].

### 4.2 Framework Operation

Narrative Scaffolding operates as a cyclical reasoning loop structured across two layers. At the **reasoning layer** (Fig. 2, B), analysts move between **narrative development** and **investigation**. Narrative development externalizes assumptions and emerging claims as narrative entries, turning them into manipulable artifacts rather than fleeting thoughts (DG1). These entries can be revised or branched into alternatives, preserving multiple lines of reasoning rather than collapsing them into a single path (DG3). Narrative entries also drive investigation, generating semantically aligned visualizations that test or expand the claims. Investigation produces insights, which are themselves externalized back into the narrative, ensuring evidence and reasoning remain tightly coupled (DG1).

At the **meta-cognitive layer** (Fig. 2, A), **reflection** regulates this cycle. Because reasoning has been externalized through narrative entries, insight timelines, and inquiry boards, reflection is



**Figure 2: The Narrative Scaffolding framework.** Analysts begin with narrative development, where assumptions and questions are externalized in response to the dataset at hand. These narratives guide investigation, producing visualizations and insights grounded in the data. Insights feed back into narrative development as new or revised hypotheses. Reflection occurs throughout, as analysts revisit their externalized reasoning and evidence. Integration consolidates claims, branches, and supporting data into coherent accounts.

no longer reliant on memory. Instead, analysts can revisit earlier claims, inspect supporting evidence, and recognize contradictions or gaps (DG2). In doing so, reflection often sparks additional narrative development, introducing new hypotheses or branches that expand the inquiry (DG3).

**Integration** (Fig. 2, C) occurs when analysts choose to consolidate what has been externalized so far. It utilizes a comprehensive set of externalized reasoning, including active claims, alternative branches, and supporting evidence, to create coherent and defensible accounts (DG4).

Figure 2 highlights this dependency: narrative development externalizes thoughts; reflection acts on those externalizations; investigation attaches evidence to them; and integration consolidates the entire set into coherent narratives. Narrative construction thus enables reasoning to remain visible, traceable, and revisable at every stage. The following section describes how these framework stages are instantiated as concrete system features and interactions in the Narrative Scaffolding system.

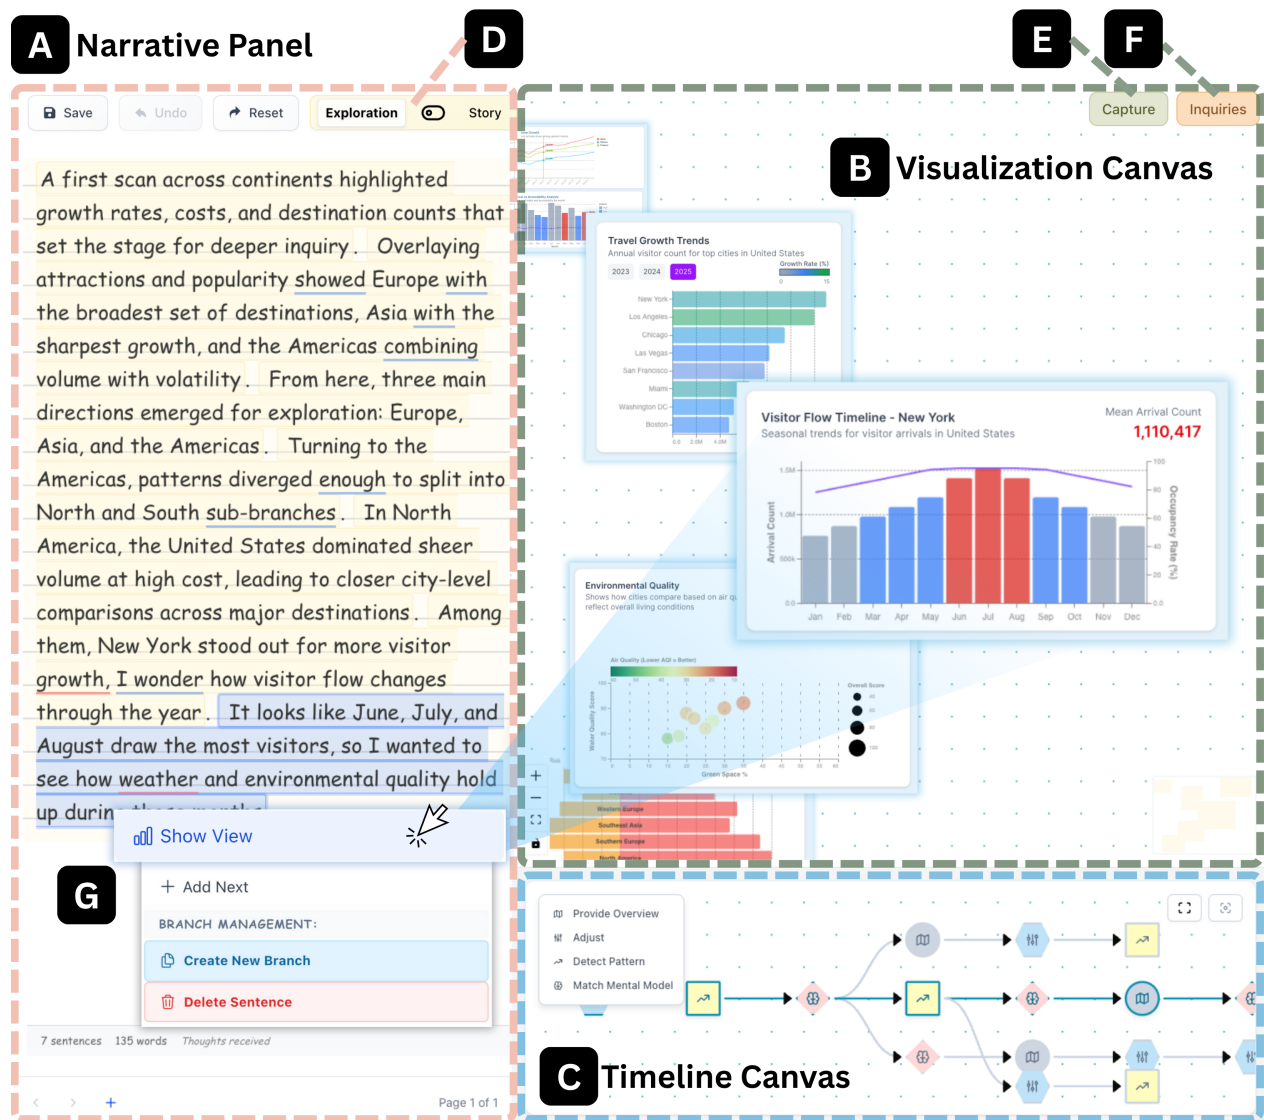
## 5 Narrative Scaffolding System

This section details how our system instantiates the conceptual model (Fig. 2), which has four core stages: narrative development, investigation, reflection, and integration. Specifically, we describe

the mechanisms driving each stage: how the system supports narrative construction and branching (5.1), enables investigation through narrative–evidence coupling (5.2), scaffolds multi-level reflection (5.3), and integrates insights into coherent, data-driven stories (5.4).

### 5.1 Narrative Development: Reasoning Through Construction & Branching

Narrative construction is the core interface of our system, where reasoning is made explicit and evolves throughout exploration. Narrative text is maintained as a tree-structured sequence of sentences, rather than a fixed linear document, allowing analysts to expand and adapt their stories as new insights emerge. Users shape this structure through four lightweight operations: appending, inserting, updating, and deleting sentences (Figure 4). These operations treat each narrative entry as a manipulable object, enabling claims, questions, and interpretations to be progressively articulated and revised. When alternative explanations arise, users can initiate **narrative branching** from any sentence (Figure 3, G), generating a parallel path that inherits all prior context. This design preserves multiple interpretations without collapsing exploration into a single account, supporting both divergence and continuity in reasoning.



**Figure 3: The Narrative Scaffolding interface integrates three canvases. (A, left) The Narrative Panel is the primary authoring space where analysts articulate, revise, and branch reasoning. (B, top right) The Visualization Canvas displays semantically aligned views that respond to narrative claims and capture user interactions as evidence. (C, bottom right) The Timeline Canvas records how insights evolve over time, preserving provenance as a navigable sequence of reasoning steps. Together, these canvases provide a shared workspace where narrative and evidence remain continuously linked.**

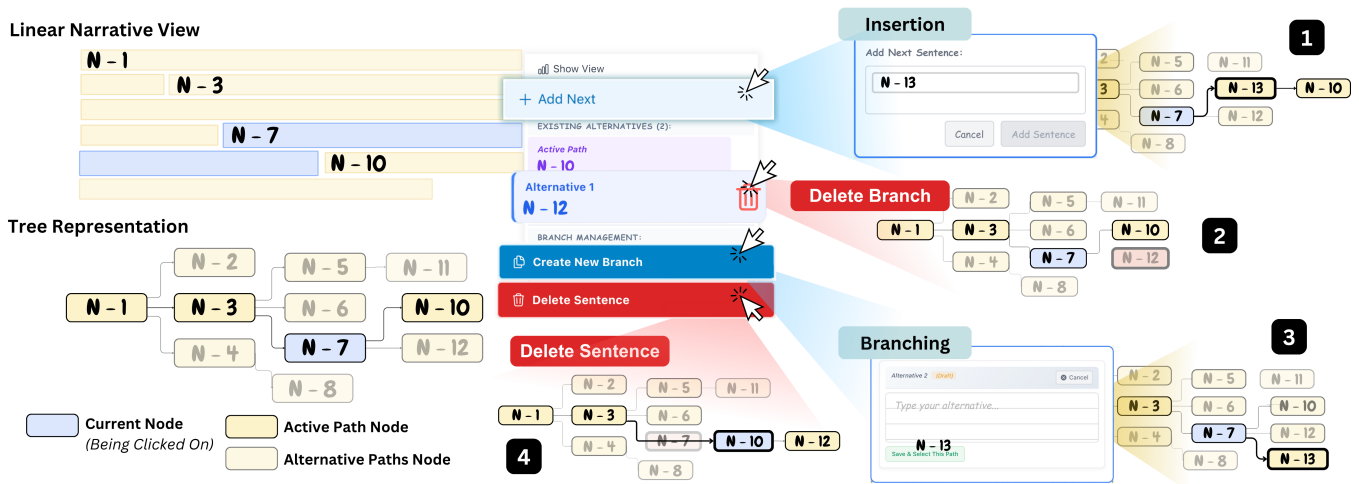
## 5.2 Investigation: Narrative–Evidence Coupling

Investigation is the central stage of narrative scaffolding, where the narrative’s intent is transformed into evidence, and that evidence subsequently generates new narrative directions. At this stage, user-authored claims generate semantically aligned visualizations and statistical summaries, while interactions with these views can be captured and integrated back into the narrative as insights. These artifacts keep reasoning linked to their supporting data, and at the same time, provide material for analysts to explore alternatives and formulate new intent.

### 5.2.1 Semantic-Aligned View Generation.

A core function of investigation is transforming underspecified narrative intent into meaningful visualizations. To address this, we developed a semantic alignment pipeline that matches narrative statements to relevant visualizations through a joint embedding space of narrative propositions and chart specifications (detailed in Appendix Figure 14).

The system preprocesses dataset schemas to generate narrative proposition templates covering common analytical patterns (ranking, temporal change, outlier detection) and pairs them with Vega visualization specifications. Both are normalized into semantic tags



**Figure 4: Core narrative editing operations in Narrative Scaffolding.** Analysts construct reasoning as editable narrative trees, using the linear narrative view and the branching tree representation. The system supports lightweight operations for (1) inserting new sentences, (3) creating branches for alternative reasoning paths, and (2) deleting sentences or (4) branches while maintaining coherent structure.

and embedded alongside pre-generated text-to-SQL queries (see Appendix A.1 for implementation details).

When users write a narrative and select *Show View* (Figure 3, G), the system matches their text against this joint space, retrieves relevant propositions, and determines whether a single chart or dashboard best addresses the claim. For instance, "Asia is cheaper but more crowded" triggers a dashboard with cost and crowding metrics, while "Porto stands out for affordability" generates a focused comparison chart. The system executes corresponding SQL queries and renders interactive visualizations, enabling meaningful view generation from vague inputs while maintaining semantic consistency.

**5.2.2 Interaction-to-Narrative Capture.** Exploration does not only flow from narrative to evidence, interactions with evidence can also be reformulated back into narrative form. As users interact with charts through actions such as clicks, hovers, or filters, the system records these operations as part of the investigative context. The five most recent interactions, together with the currently active narrative claims, are passed to the LLM to generate candidate narrative statements. When users click *Capture* (Figure 3, E) in the visualization canvas, these candidate insights summarize what the analyst is probing in the data and are presented for capture into the narrative pane. Once accepted, they are appended as narrative sentences and simultaneously logged in the *Insight Timeline*, preserving both the text and its supporting evidence as artifacts.

### 5.3 Reflection: Tracing Reasoning Evolution

Reflection provides moments to revisit, evaluate, and reframe emerging narratives. In the system, this process is structured across three layers: expression, interpretation, and reasoning. These layers ensure that users can re-engage with their analysis, whether by recalling a single claim, examining a sequence of insights, or assessing the overall trajectory of inquiry.

### Sentence-Level Reflection



**Figure 5: Sentence-level reflection allows analysts to revisit specific narrative statements in context, linking individual sentences to visual evidence for localized reasoning.**

**5.3.1 Sentence-to-View Linking.** The narrative panel supports direct recall between narrative and evidence (Figure 5). Every visualization generated during the investigation is preserved on the canvas and linked to the sentence that introduced it. This linkage is bidirectional: users can recall a visualization from its associated sentence, or select a visualization to highlight the claim it supports. At this level, reflection occurs on a local, sentence-by-sentence basis, providing in-the-moment awareness (i.e., "Am I making a claim? Have I justified it?"). This bidirectional recall ensures that evidence and expression remain tightly connected throughout the analysis.

**5.3.2 Insight Provenance Timeline.** The insight timeline captures how insights emerge and evolve as users write, recording

each update as a node (Figure 6). We adapted the Data–Frame taxonomy [42, 104] to classify insights into four relationship types describing how claims shift the trajectory of reasoning: **provide overview**, **adjust**, **detect pattern**, and **match mental model**. The timeline is organized as a tree structure, preserving both the current path and alternative branches. At this level, reflection focuses on how a user’s thinking shifts over time, helping users answer questions such as “How did I arrive at this claim?” or “Has my hypothesis drifted?”. By visualizing the trajectory of reasoning as a tree, this design provides a navigable record that allows users to revisit and recover the precise state of thought at any point in their exploration.

**5.3.3 Inquiry Board.** The inquiry board captures how inquiries unfold over time (Figure 7). Inquiries are derived directly from the narrative: when users write hypotheses or intentions to explore, these statements are formalized as issues, while claims provide potential answers. Because exploration is ongoing, many inquiries may remain unanswered or give rise to new ones as paths branch. The board categorizes all current inquiries into three statuses: **open**, **resolved**, or **stalled**, and arranges them in an IBIS graph [11, 45] to show how issues relate to one another. At this level, reflection shifts from local claims to the broader organization of reasoning. The inquiry board enables users to revisit unresolved issues, re-engage with abandoned threads, and assess whether current claims have sufficient support, preventing forgotten questions from fading silently out of the reasoning process.

Together, these layers scaffold narratives through reflection at multiple levels: from local articulation of claims, to tracing how insights evolve, to organizing inquiries into a higher-level reasoning structure.

## 5.4 Integration: Synthesis as Data-Driven Story

**Integration** reorganizing nonlinear exploration into a single, communicable narrative grounded in the user’s reasoning process. Triggered by switching to the *Story Mode* (Figure 3, D), the system compiles all narrative branches and identifies key conclusions, decision points, and supported insights across the entire insight timeline. The LLM is prompted with the entire narrative graph structure, including associated inquiry paths and metadata that link each claim to evidence and its corresponding inquiry context. Only insights that are statistically grounded are included to ensure coherence and relevance. The system generates a clean, linear narrative where each paragraph contains synthesized insights linked back to their supporting visualizations.

## 5.5 Usage Scenario and System Walk-through

Imagine Amy, a data journalist, receives an assignment to analyze several travel datasets for spring destination recommendations.

**Narrative Development & Investigation.** She begins with intention and narrative development, drafting an opening claim in the narrative pane about balancing affordability, environmental quality, and crowding. This externalized sentence guides her exploration, and when she clicks *Show View*, the system generates aligned dashboards: cost distributions by region, sustainability rankings, and popularity metrics, that ground her claim in evidence.

As Amy examines the visualizations, she poses a new question about European sustainability premiums compared to Asia, and the

system responds with targeted comparisons. Clicking *Capture* saves her observation (“Asia is 21.5% cheaper, but 62.8% more crowded”) into the narrative, while also attaching the supporting views. The *Insight Timeline* records a new node, ensuring both her claim and the underlying data remain visible in the evolving reasoning process. Her assumptions and observations are captured by the system as externalized artifacts that remain persistently linked to the visual evidence that generated them.

**Reflection.** Drilling down from regional to city-level patterns, Amy captures insight “Porto stands out for budget appeal among mid-size cities”, saved to the narrative with linked cost distributions and a timeline node. As Amy continues building her narrative around multiple destinations, she realizes she can’t recall whether Porto’s low cost reflected seasonal pricing or year-round affordability. To verify this detail, she needs to revisit earlier reasoning. She reopens her Porto claim with *Show View*, inspects the attached visuals, and notices a weakness in environmental quality that contradicts her initial assumption. She edits the sentence to qualify the recommendation (“Porto offers budget appeal but lags in environmental quality compared to northern alternatives”), and the system logs this revision as a new insight on the timeline. Since Amy’s reasoning is externalized and persistently linked to evidence, reflection becomes systematic and retrieval-based.

As exploration continues into Asian cities, Amy encounters a clash between emerging patterns: while exploring Bangkok, her “cheap but crowded” storyline conflicts with strong cultural diversity metrics. Rather than overwrite the earlier path, she right-clicks the Bangkok paragraph and selects *Create Branch*, forking two parallel narratives: one on budget cultural immersion and another on premium sustainability. The *Insight Timeline* records the divergence, preserving both branches with their linked evidence so each can develop independently without losing coherence. Narrative branching allows her to externalize and pursue divergent reasoning paths while maintaining traceability to the decisions.

With multiple threads in motion, Amy steps back for strategic oversight. She opens the *Inquiry Board* to review resolved, active and stalled questions. Jumping from inquiry cards to their linked narrative sections, she abandons weaker paths and focuses on analyses with more substantial evidence, keeping exploration aligned with the evolving story. Inquiry-level reflection allows Amy to monitor the status and progression of her analytical questions, assess which lines of reasoning remain productive, and strategically redirect effort toward more promising directions.

**Integration.** When several threads feel coherent enough to communicate, Amy switches to *Story Mode*. The system assembles her active claims, branches, and embedded views into a structured draft. The resulting article aligns each claim with supporting evidence and preserves transparency to alternative branches explored during analysis. Story Mode enables narrative synthesis into a cohesive output where every claim remains connected to its supporting visualizations and reasoning context, producing a defensible and transparent account of her analytical process.

## 5.6 Implementation Details

Narrative Scaffolding is implemented as a web application with a Python Flask backend and React Next.js frontend. We employ a

### Insight-Level Reflection

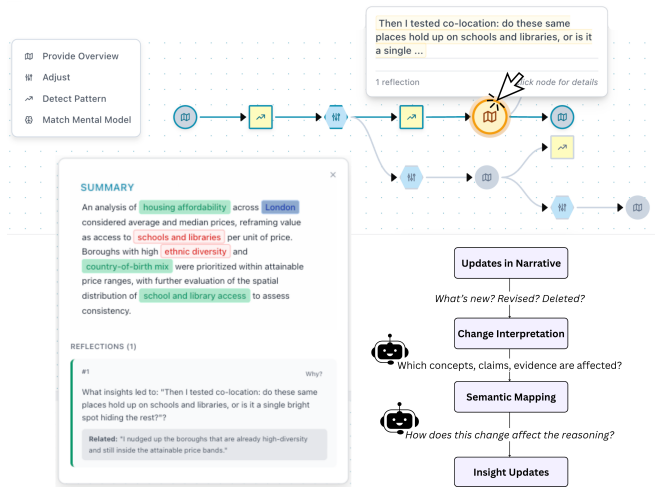


Figure 6: Insight-level reflection captures how edits in the narrative propagate to interpretive changes in the reasoning structure.

chain-of-thought [100] prompting strategy to make LLM reasoning more explicit and context-aware. GPT-5 was used to handle higher-level narrative generation tasks that require longer context windows and complex reasoning for storyline development, ensuring faithful mappings between user intent and chart configurations. GPT-4o-mini was used to manage lightweight operations where responsiveness is critical, such as insight extraction, inquiry summarization, and reflective nudges. This approach was chosen based on empirical testing that balanced consistency, speed, and task-specific performance.

Data visualizations utilize Vega-Lite for chart rendering and Leaflet.js for interactive maps. Text-to-SQL functionality leverages VannaAI [95] with OpenAI embedding and few-shot training for query consistency. The system is designed to be model-agnostic, allowing for easy substitution of LLMs as capabilities evolve. Complete prompt templates and configuration details are provided in supplementary materials.

### 6 Evaluations

Our evaluation spanned user and system evaluations. We aimed to assess the usability and effectiveness of Narrative Scaffolding (referred to as NS in the following sections), as well as its reliability in handling underspecified narrative input from users.

To assess whether NS achieves the design goals identified in our formative study (sec 3), we structured our evaluation around five research questions. Each of the first four questions corresponds to a design goal, examining how the system’s features translate into observable benefits during analysis. The fifth question examines system-level reliability, which also relates to DG1’s emphasis on handling vague input.

RQ1 [Workflow Fluency] How does NS support fluid transitions from vague inquiries to evidence-grounded insights

### Inquiry-Level Reflection

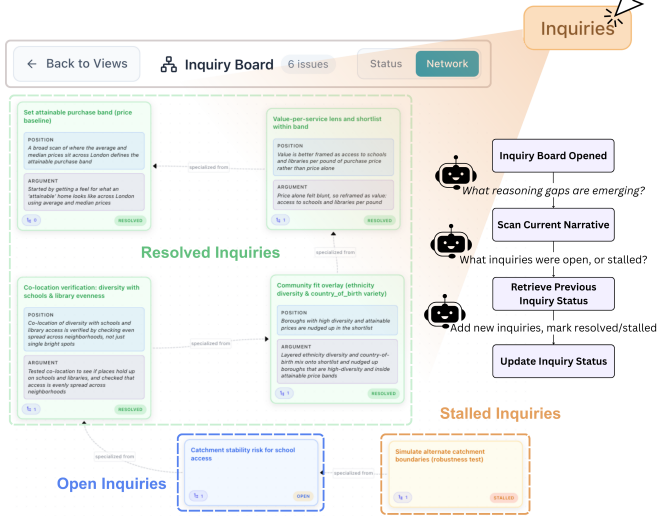


Figure 7: Inquiry-level reflection synthesizes across insights, reviewing and linking how questions evolve and progress over time.

(DG1), and what impact does this have on the integration of exploration and narrative construction?

RQ2 [Reflection Depth] How does NS help users externalize their reasoning and maintain continuity of reflection (DG2) as their analysis evolves over time?

RQ3 [Exploration Divergence] How does NS support nonlinear exploration and alternative reasoning (DG3), enabling users to explore broader analytical space and integrate more evidence?

RQ4 [Reasoning Traceability] How does NS support traceable reasoning (DG4), making the analytical process behind insights more visible and defensible?

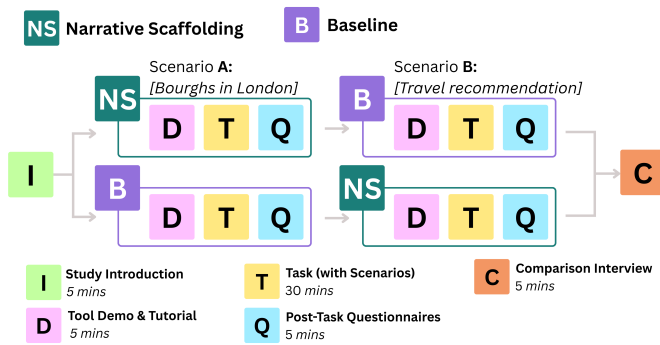
RQ5 [Output Reliability] How reliably does NS produce suitable analytical outputs when prompted with vague, ambiguous, or underspecified narratives (DG1)?

The first four research questions focus on assessing the impact on users’ workflows and output, while the fifth research question investigated the system’s capability. The institutional ethics review board approved both studies.

#### 6.1 Study 1: User Evaluation (RQ1-RQ4)

We conducted a within-subject evaluation (N=20) to examine whether NS enables more fluid workflows (RQ1), deeper reflection (RQ2), broader exploration (RQ3), and more defensible narratives (RQ4) compared to a baseline analytics paradigm, where participants explored data primarily through views and composed narratives only after analysis.

6.1.1 Baseline and Apparatus. In addition to our NS system, we implemented a Baseline for comparison. The Baseline is a web-based analytics workbench that replicates typical Power BI workflows [57], supporting natural-language querying and visualization



**Figure 8: Overview of the study procedure. Participants completed both Narrative Scaffolding (NS) and Baseline (B) conditions in counterbalanced order.**

generation within a traditional dashboard interface. Participants could issue queries, apply filters, generate charts, and export findings using conventional point-and-click interactions supplemented by conversational AI assistance. To ensure a fair comparison focused on workflow design rather than technical capabilities, both systems operated on the same data schema, text-to-SQL translation, and semantic view generation pipeline (5.2.1). For the final integration stage, we exported tabular data from generated visuals via the Power BI REST API to generate narrative summaries of trends, comparisons, and anomalies [56].


This web-based approach allowed participants to join sessions remotely via video conferencing, while also enabling precise logging of user interactions and timing data. The only difference between NS and Baseline lay in the interaction paradigm, isolating effects of interaction structure and reasoning support from differences in analytic functionality.

**6.1.2 Method.** We recruited 20 participants (11 male, 9 female; age 23–66,  $M = 32$ ) through Prolific, using a screening process to ensure prior experience with professional visualization tools. All participants reported using Power BI or Tableau at least once per week. The sample included 7 academic researchers, 7 managers, 5 data analysts, and 1 healthcare professional.


**[Tasks]:** Participants completed two tasks, each requiring them to write a data narrative recommending options based on multi-factor datasets with eight factors each. Task A involved recommending boroughs for long-term housing value, safety, and community well-being using borough-level data (housing prices, ethnicity, schools, libraries, renting prices, income, crime, gyms, restaurants). Task B involved recommending travel destinations using destination-level data (cost of stay, safety, reviews, diversity, environmental quality, attraction coverage, popularity, crowd dynamics). Both datasets were designed to be balanced in dimensionality and variance to ensure performance differences could be attributed to the systems rather than task complexity.

**[Procedure]:** We adopted a within-subject design where each participant used both NS and Baseline, with tool order counterbalanced to mitigate learning effects (Figure 8). Sessions began with

**SUS Positive Question Distribution**



**SUS Negative Question Distribution**



**Figure 9: SUS questionnaire results. Positive items (top) received strong agreement, while negative items (bottom) were largely disagreed with.**

informed consent, demographics collection, and system introduction. For each system, participants received a tutorial, completed the assigned task while thinking aloud [16] (explicitly verbalizing reflection, exploration, and insights), and completed a post-task survey evaluating usability, workload, and experience. After both systems, participants completed a final comparison survey reflecting on workflow differences. Each session lasted approximately 90 minutes, and participants received \$20 USD compensation upon completion. Study instructions and survey instruments are provided in the supplementary materials.

**6.1.3 Measures and Analysis.** We collected both perception and interaction data to evaluate the two systems comprehensively. We analyzed post-task survey responses that included 7-point Likert scale questions on workflow efficiency, exploration, reflection, and explainability. To assess usability and workload, we incorporated the SUS [7] and NASA-TLX [28], alongside self-assessments of the clarity and defensibility of participants’ final narratives. Interaction logs captured differences in analysis patterns, including query diversity, revisions, and visualizations created, with temporal measures normalized across sessions. In addition, we recorded think-aloud protocols during task sessions to capture participants’ reasoning processes in situ. These transcripts were then thematically analyzed to account for the reflection and exploration process. We compared survey responses between the Baseline and NS conditions using the Wilcoxon signed-rank test, given the ordinal nature of Likert-scale data. For continuous temporal measures, we applied Welch’s t-test when the normality assumptions were met and the Mann-Whitney U test otherwise. To control for false discovery due to multiple comparisons, we applied the Benjamini-Hochberg procedure [5]. We report Cohen’s d effect sizes for all parametric comparisons to indicate the magnitude of practical differences between conditions.

**6.2 Study 2: System Evaluation (RQ5)**

We evaluated system robustness to vague narrative input, addressing a key requirement for narrative-first interfaces: the ability to translate underspecified user intentions into actionable analytical starting points.

**6.2.1 Evaluation Material.** To assess system robustness to vague input, we constructed a pool of 100 prompts. We first selected 43

vague prompts collected from the user study that exhibited natural underspecification in areas such as timespan, entity, or aggregation level. We used an LLM to expand the pool by generating additional prompts that followed similar templates but varied in domain and phrasing. To ensure ecological validity, prompts were required to be underspecified in key dimensions (i.e, missing temporal scope, unclear entity reference, or ambiguous level of detail).

**6.2.2 Expert Rating.** We conducted an expert evaluation to assess how well the system outputs addressed vague prompts. Six experts were recruited, with eligibility limited to individuals reporting a minimum of five years of experience in visualization literacy ( $M=9.5$  years). These experts were selected for their ability to evaluate whether visualization choices appropriately match analytical intent and support iterative exploration. Each expert was assigned 50 prompts. Each of the 100 prompts was independently rated by three different experts, resulting in a total of 300 ratings across the pool.

Each visualization was rated on two dimensions using 7-point Likert scales: (i) how well the chart plan aligned with the intent of the vague prompt, and (ii) the extent to which the chart plan supported further exploration. We averaged the three ratings per visualization to obtain a single score on each dimension.

## 7 Results

We present an analysis of survey responses and participant interactions ( $P_1 - P_{20}$ ) with both conditions. Our results show that NS enabled participants to generate insights more efficiently, engage in deeper and more continuous reflection, explore broader analytical alternatives, and produce more defensible narratives compared to the baseline condition. NS received an average System Usability Score (SUS) of 79.8 (Figure 9), demonstrated high usability (*Excellent* according to [2] and passing the cutoff score, 68, for production by a large margin).  $P_{11}$  described the system as feeling "like a reliable companion" during analysis.

Figure 10 provides a timeline view of participant activity across conditions. It highlights how NS sessions were denser and more varied, with frequent reflection and branching events, whereas Baseline sessions consisted largely of investigation alone. This overview illustrates the systemic advantage of NS, which we unpack across workflow fluency (RQ1), reflection (RQ2), exploration divergence (RQ3), and explainability (RQ4).

### 7.1 [RQ1]: Workflow Fluency

**7.1.1 Participants Generated More Insights Through Deeper Engagement.** Participants took slightly more time to complete tasks with NS ( $M = 31.59$ ,  $SD = 0.92$  min) compared to Baseline ( $M = 29.84$ ,  $SD = 0.87$  min). This result is unsurprising, as our goal was to support deeper engagement with data through NS rather than to accelerate task completion.

However, this similar time investment resulted in substantially greater analytical productivity. NS participants generated significantly more insights ( $M = 15.65$ ,  $SD = 6.95$ ,  $p < 0.001$ ) compared to the Baseline ( $M = 6.85$ ,  $SD = 1.69$ ). This increased insight volume translated into dramatically improved efficiency. The time per insight was significantly lower in NS ( $M = 2.31$  min/insight,  $SD = 0.76$ ,  $p < 0.001$ ) compared to Baseline ( $M = 4.62$  min/insight,  $SD = 1.22$ ),

representing faster insight articulation. This result aligns with our design goal of enabling more efficient insight generation during exploratory analysis.

**7.1.2 Participants Experienced Reduced Cognitive Load.** In addition to generating significantly more insights, NS also reduced perceived cognitive load (Figure 11, A). We used NASA-TLX to measure the perceived workload associated with each system. Compared to Baseline, NS had significantly lower mental demand ( $M = 3.00 < 4.45$ ,  $p = .015$ ), required significantly less effort ( $M = 2.95 < 4.60$ ,  $p < .001$ ), resulted in substantially less frustration ( $M = 3.10 < 4.75$ ,  $p = .001$ ), and led to significantly better perceived performance ( $M = 2.25 < 3.60$ ,  $p = .024$ ), where lower scores indicate better performance. This suggests that NS made the analytical process feel more natural and manageable despite the increased analytical output.

**7.1.3 Participants Found Direction and Generated Structured Output More Easily.** NS participants rated significantly higher ability to find direction without specific questions in mind (Figure 11, B) (NS:  $M = 5.45$ ,  $SD = 0.69$ ; Baseline:  $M = 2.55$ ,  $SD = 0.93$ ,  $p < .001$ ). Participants appreciated how NS helped them navigate uncertainty in early exploration.  $P_6$  noted the system could "quickly see what's important and suggest possible directions," while  $P_{12}$  described how the previews eliminated the challenge of "starting at abstract ideas trying to figure out where to start."

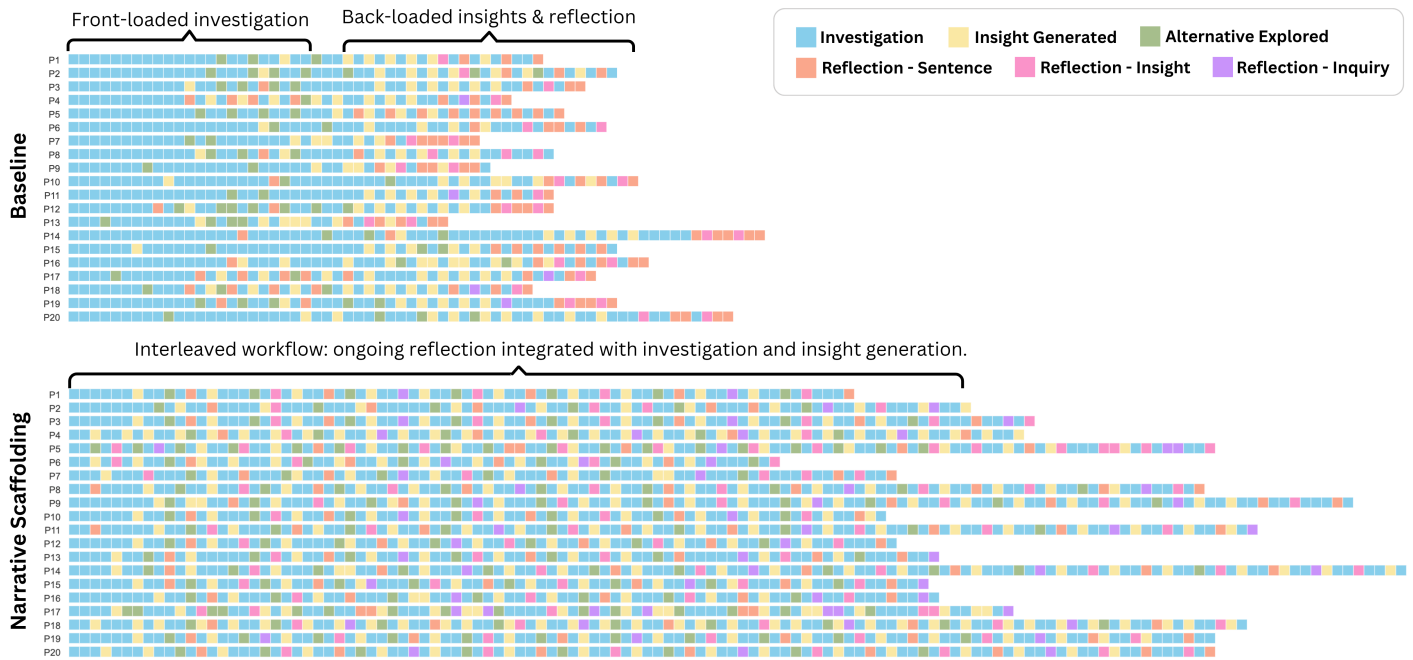
NS also significantly improved participants' ability to turn exploration into structured output (Figure 11, B) (NS:  $M = 5.45$ ,  $SD = 0.89$ ; Baseline:  $M = 3.50$ ,  $SD = 1.06$ ,  $p = .002$ ).  $P_7$  found that "turning exploration into a report feels almost effortless," and  $P_{17}$  highlighted how the system "transforms explored branches into a structured story with one click." These ratings align with the behavioral efficiency measures, indicating that participants both experienced and recognized the workflow improvements provided by NS, supporting DG1's goal of fluid transitions from vague inquiries to evidence-grounded insights.

### 7.2 [RQ2]: Reflection Depth

**7.2.1 Participants Engaged in More Frequent and Deeper Reflection.** Across three levels of reflection: sentence-level, insight-level, and inquiry-level, NS participants engaged in significantly more reflections ( $M = 14.80$ ,  $SD = 3.69$ ,  $p < 0.001$ ) compared to Baseline ( $M = 6.50$ ,  $SD = 1.91$ ). The improvement varied substantially by reflection type: sentence-level showed a 1.2× increase, insight-level increased 3.7×, and inquiry-level demonstrated a 13.0× improvement in NS compared to Baseline.

The temporal distribution of reflective activity also differed between conditions (Figure 13). NS participants distributed reflection events more evenly across the session timeline, while Baseline participants concentrated most reflection events in the final portions of their sessions. As  $P_1$  noted: "With [NS] it felt natural to pause along the way, look back at what I'd written."

**7.2.2 Participants Reported Higher Reflection Quality and Usefulness.** While NS generated more reflection events, we examined whether these were actually useful for reasoning improvement. We measured reflection events that were followed within 30 seconds by an edit to the related narratives. Baseline participants showed higher reflection-edit rates (58.8% led to immediate edits) compared



**Figure 10: Events timeline visualization comparing analytical workflows across Baseline and Narrative Scaffolding conditions. Each horizontal bar represents one participant’s session with color-coded activities. Baseline (top) exhibits a front-loaded investigation pattern with insights and reflection concentrated toward the end of sessions, reflecting a separation between exploration and reasoning articulation. Narrative Scaffolding (bottom) demonstrates an interleaved workflow where investigation, insight generation, and reflection occur continuously throughout sessions, with substantially more alternative exploration and inquiry-level reflection.**

to NS participants (35.2%). However, this lower rate does not necessarily indicate a reduction in quality, as not all helpful reflections require immediate edits.

To assess usefulness, we examined self-reported measures of reflection quality (Figure 12, C). NS participants felt significantly better supported in recognizing when to revisit reasoning (NS:  $M = 5.85, SD = 1.23$ ; Baseline:  $M = 3.00, SD = 1.17, p < .001$ ) and reflecting on trade-offs between options (NS:  $M = 5.65, SD = 1.39$ ; Baseline:  $M = 3.60, SD = 1.27, p < .001$ ). Most importantly, participants reported that revisiting previous steps significantly improved their final output quality (NS:  $M = 6.05, SD = 0.89$ ; Baseline:  $M = 2.65, SD = 1.18, p < .001$ ).

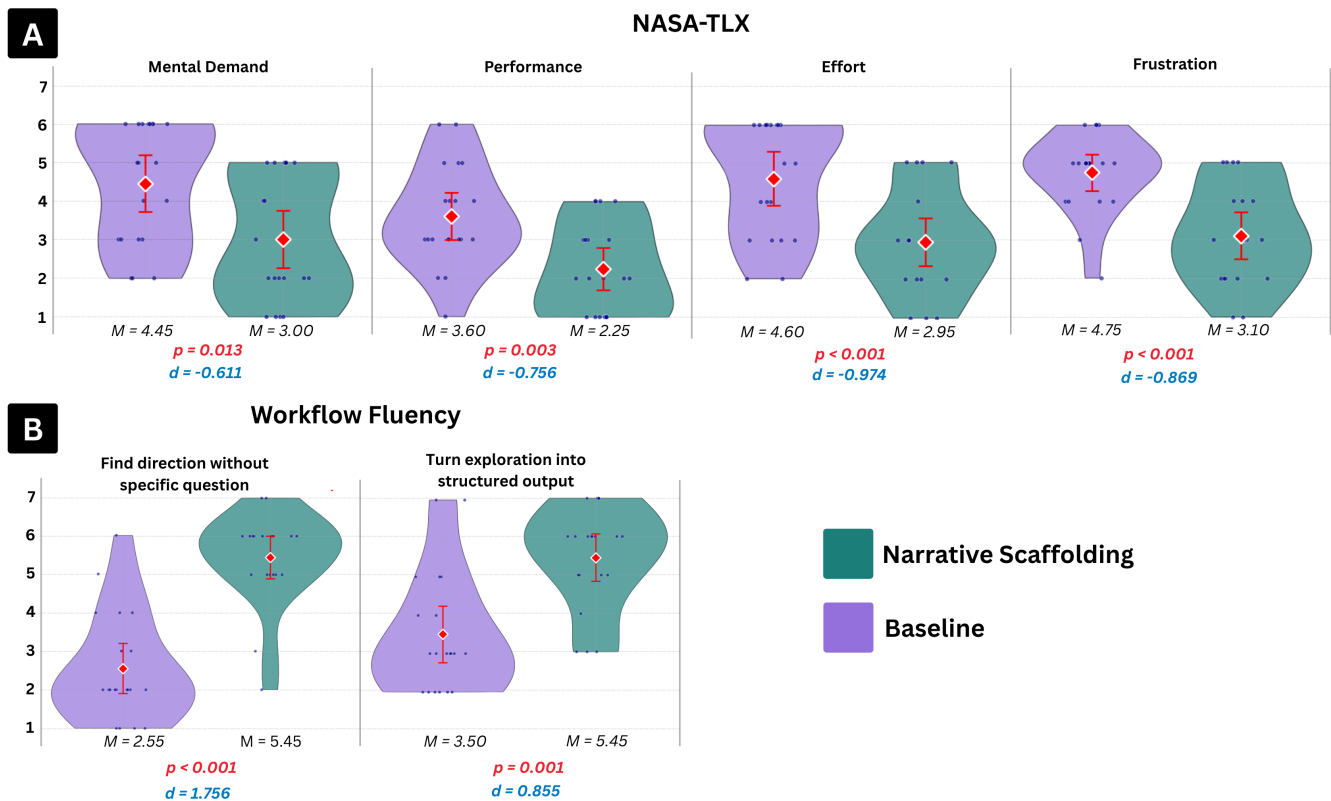
NS participants described using reflection for ongoing reasoning validation rather than just problem-fixing.  $P_{13}$  noted relying on "the timeline to check when and why earlier reasoning might need to be revisited," while  $P_{10}$  found "the inquiry board helps me track which questions still need attention." These findings indicate many reflection events in NS serve to confirm conclusions or validate reasoning. Despite lower reflection-edit rates, NS participants reported significantly higher perceived benefits for reasoning quality, suggesting that NS successfully externalized inquiry intent and maintained reflective continuity (DG2).

### 7.3 [RQ3]: Exploration Divergence

**7.3.1 Participants Explored More Alternatives and Factors.** With NS, participants considered significantly more alternatives during their analysis ( $M = 8.75, SD = 3.03, p < 0.001$ ) compared to Baseline participants ( $M = 3.25, SD = 1.58$ ). They also explored a greater breadth of available factors, examining on average 7.0 out of 8 possible factors ( $SD = 0.9$ ), compared to Baseline participants who explored 5.3 factors ( $SD = 1.4$ ). As  $P_8$  explained, "With NS I felt like I could easily try different takes on the data, instead of sticking to one angle."

Self-reported measures confirmed that participants felt NS better supported exploratory breadth (Figure 12, D). Participants rated NS significantly higher for encouraging different alternatives for analyzing the same data (NS:  $M = 6.25, SD = 0.91$ ; Baseline:  $M = 3.95, SD = 1.47, p < .001$ ) and for making it easier to consider and contrast alternative perspectives (NS:  $M = 5.80, SD = 1.20$ ; Baseline:  $M = 3.10, SD = 1.29, p < .001$ ). Participants ( $P_4, P_7, P_{13}, P_{16}, P_{17}$ ) themselves described this sense of breadth, noting that "[Baseline] made me tunnel on one or two things, but [with NS] I actually considered a range of options before narrowing down." –  $P_4$

$P_{15}$  emphasized how NS encouraged breadth across factors: "The system had nudged me to look at other dimensions when I reflected, like community well-being alongside cost." Together, these findings suggest that NS not only increased the number of alternatives participants explored but also helped them consider broader coverage of



**Figure 11: Workflow Fluency and Cognitive Load Comparison of (A) NASA-TLX workload dimensions and (B) workflow fluency measures between NS and Baseline conditions. NS demonstrated significantly lower cognitive demands ( $p < .05$ ) and improved support for exploration initiation and output structuring. Violin plots show response distributions with means marked by red diamonds.**

available factors, creating more opportunities to surface contrasting perspectives.

This broader and more varied exploration pattern is also visible in the activity timelines (Figure 10): NS participants interleaved investigation with reflection and branching throughout their sessions, while Baseline activity remained concentrated on linear investigation. This illustrates how NS not only encouraged participants to try more alternatives but also supported a workflow where reflection and exploration reinforced one another.

### 7.3.2 Participants Integrated More Evidence into Decision-Making.

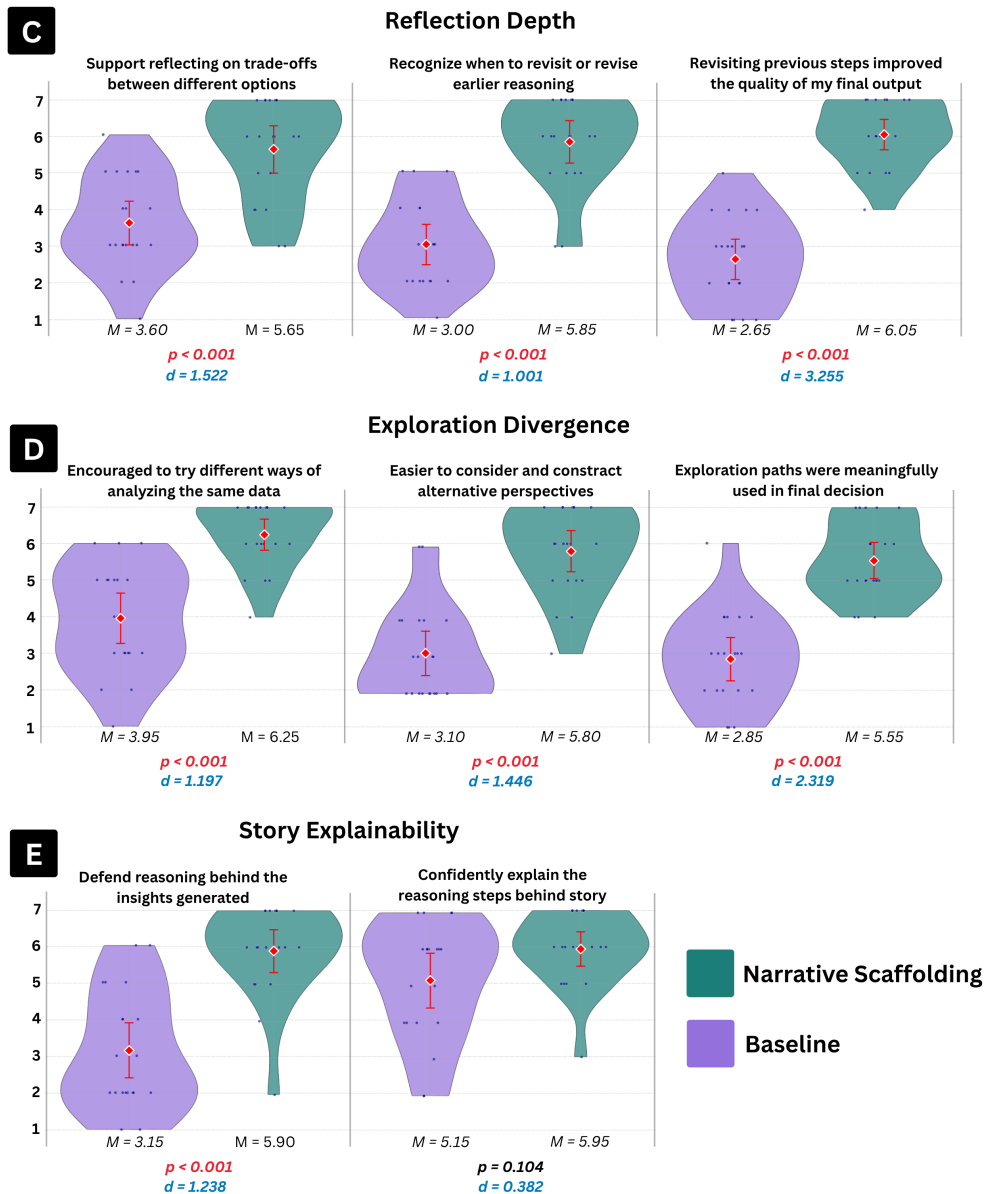
Beyond exploring more alternatives, NS participants integrated more of their exploration into final decisions. They incorporated significantly more factors into their decision-making process ( $M = 3.8, SD = 0.7$ ) compared to Baseline participants ( $M = 2.5, SD = 0.7$ ). This suggests that NS participants not only explored more broadly but also utilized their exploration more meaningfully. Participants emphasized that the more exhaustive exploration in NS translated into having more actionable options to weigh. As  $P_{11}$  explained, "Because I had looked at more angles, I actually had choices, I could go in a few different directions instead of just one." Similarly,  $P_6$  elaborated, "Having more alternatives in [NS] meant I didn't have to stick with my first idea [...] I could line up different options and

actually remember how they compared, which made it easier to choose from."

Self-reports reinforced this pattern (Figure 12, E), participants rated NS significantly higher for ensuring exploration paths were meaningfully used in their final story (NS:  $M = 5.55, SD = 1.05$ ; Baseline:  $M = 2.85, SD = 1.27, p < .001$ ). These results indicate that NS enabled not only wider exploration but also more systematic integration of diverse evidence into coherent analytical conclusions (DG3), providing participants with more usable options when making decisions.

## 7.4 [RQ4]: Reasoning Traceability

We assessed whether NS improved users' ability to trace and justify their analytical reasoning process. Participants using NS reported significantly greater confidence in defending the reasoning behind their insights ( $M = 5.90, SD = 1.25$ ) compared to Baseline ( $M = 3.15, SD = 1.60, p < .001$ ). Participants described how timelines and provenance features supported justification.  $P_2$  explained, "The timelines show the whole process which guarantees transparency and provides relevant evidence."  $P_{13}$  noted, "Since the insights were linked to evidence and captured in the timeline, it is pretty straightforward to justify my reasoning." Similarly,  $P_{17}$  emphasized, "Since every




**Figure 12: Reflection, Divergence, and Explainability Self-reported measures of (C) reflection depth, (D) exploration divergence, and (E) story explainability across conditions. NS participants reported significantly stronger support ( $p < .05$ ) for revisiting reasoning, exploring alternatives, and defending conclusions. Violin plots show response distributions with means marked by red diamonds.**

*captured insight is automatically linked back to its supporting data and visualizations, I would always have clear evidence to justify my claims.* P<sub>20</sub> summarized this benefit as, "I felt more confident that I could defend my choices because I had pulled in evidence from different places, not just one metric."

When asked about presenting conclusions to others, participants also reported higher average confidence with NS ( $M = 5.95$ ,  $SD = 1.00$ ) than with Baseline ( $M = 5.15$ ,  $SD = 1.60$ ). However, this difference was not statistically significant ( $p = .112$ ). Participants

( $N = 13$ ) rated themselves as confident presenting with NS, citing features such as organized evidence, visible branches, and insight timelines, which serve as valuable sources of support. For instance, P<sub>1</sub> stated, "With all evidence organized, branches visible, and a polished story generated, I'd feel confident explaining my analysis to anyone." Similarly, P<sub>10</sub> noted, "Every insight is backed by reasoning, and the story mode is also helpful; therefore, it will aid in justifying my recommendations."



**Figure 13: Temporal distribution of reflection events across three types throughout the analysis session. Each dot represents a reflection event by a participant, positioned along normalized session time (0.0 = start, 1.0 = end). Narrative Scaffolding participants engaged in reflection more frequently and distributed more evenly across the session, particularly at the sentence and insight levels. Baseline participants showed concentrated reflection activity toward the end of sessions, with notably fewer inquiry-level reflections overall.**

Among these same participants, ( $N = 9$ ) also rated themselves as confident presenting with Baseline. In these cases, confidence stemmed from their own skills rather than the tool itself:  $P_1$  remarked, "I feel really confident because once I make a presentation, I will be able to address the weak points in my insights and conclusions." Similarly,  $P_{12}$  explained, "I'd be comfortable with the recommendations themselves, but less confident explaining the full path I took to get there since some of that context lives outside the tool."

This contrast suggests that while participants often considered themselves capable of presenting regardless of the tool, the source of their confidence differed. In NS, confidence was rooted in system-supported provenance and narrative structure (DG4), whereas in Baseline, confidence was self-reliant, grounded in personal presentation skills rather than the analytic workflow.

## 7.5 [RQ5]: Output Reliability

We analyzed expert ratings ( $E_1 - E_6$ ) along two dimensions: (i) alignment with the intent of vague prompts and (ii) support for further exploration. The first captures whether outputs were judged as plausible and relevant responses to the expressed inquiry, recognizing that vague prompts can admit multiple reasonable interpretations. The second captures whether outputs were seen as useful entry points for continuing analysis, since vague prompts are rarely end goals but instead serve to open inquiry. For each dimension, we report quantitative results followed by qualitative feedback from experts on cases they felt did not work well.

Expert ratings indicated that the system outputs were generally robust across vague prompts. For intent alignment, 77% received "agree" or "strongly agree" ratings, 20% fell between "slightly agree" and "neutral," and only 3% were rated "slightly disagree" or lower. For supporting exploration, 79% were rated "agree" or "strongly agree," 18% between "slightly agree" and "neutral," and 3% "slightly disagree" or below. These distributions suggest that the large majority of outputs were considered both relevant to the prompt and helpful in sustaining inquiry.

Experts explained that lower ratings often reflected a failure to capture implied intent. While generated dashboards were usually relevant to the surface wording of the prompt, they sometimes missed the deeper analytical angle that vague phrasing implied. For example,  $E_3$  noted a prompt about "some neighbourhoods are

unexpectedly safe" suggested an outlier-focused analysis, but the system produced generic averages and correlations. Similarly,  $E_5$  explained that when asked about "travel growth differences between countries" the system produced growth rate visualizations that "It showed similar percentages across countries, but missed that 10% growth could mean millions in one place and only thousands in another."  $E_2$  described such cases as "technically correct but superficial". This implies that while the majority of outputs were judged as good matches to the prompts, a smaller subset still fell short by failing to respond to the implied meaning behind vague inquiries.

In general, experts appreciated that the system consistently produced plausible and relevant outputs, though they also noted occasional reliance on generic templates or literal matches that did not surface distinctive analytical directions. These suggest that the system is sufficient for producing semantically aligned dashboards from vague prompts at scale. Targeted refinements (diversifying templates, capturing implied intent) could further enhance its utility.

## 8 Discussion

### 8.1 The Narrative-First Paradigm: Supporting Natural Reasoning Through Writing

Our findings reveal how narrative construction can function as a primary reasoning interface during data analysis, aligning with narrative psychology's view that storytelling structures how people make sense of complex information [8, 58]. Rather than treating writing as post-analysis documentation [47, 105], participants used NS not merely to document completed insights, but as a thinking medium, a space where vague hunches evolved into precise, evidence-backed claims through the act of writing.

A key aspect of this narrative-as-thinking-medium approach was adapted to how participants actually think. This paradigm accommodated diverse analytical entry points and supported continuous interpretive refinement throughout exploration. Unlike visualization-first tools that require users to translate thoughts into structured queries [20, 79], NS accommodated diverse entry points—whether specific hypotheses or exploratory questions. This addresses the "gulf of envisioning" [85]: the difficulty of translating vague intentions into system-actionable inputs. Participants could write underspecified hunches, generate visualizations, then iteratively refine their narrative to pursue new directions—each refinement triggering updated views without breaking flow. This fluidity

reduced friction not just during early-stage exploration [39, 102], but throughout the entire sensemaking process, as participants' understanding evolved and required new framings.

Beyond these cognitive benefits for users, narrative-first design provided unique system-level benefits by making reasoning evolution visible and trackable. Externalizing thoughts in narrative form allowed participants to monitor how their assumptions shifted and revisit earlier claims. This highlights how narrative construction serves dual purposes: supporting human reasoning while providing systems with insight-level data that traditional interaction logs cannot capture [65, 73].

Our results demonstrate that when narrative construction becomes the primary reasoning interface, systems can scaffold the process of thinking rather than just its outputs. This opens possibilities for more sophisticated provenance, adaptive scaffolding, and reasoning-aware recommendations. However, some participants noted that highly quantitative comparisons — such as spatial relationships (e.g., visually scanning which of 50 cities cluster together geographically) or multi-dimensional trade-offs (e.g., simultaneously meeting cost and safety thresholds) — were difficult to express through text alone, as these tasks benefit from parallel visual processing. These limitations suggest that while narrative-first design offers powerful scaffolds for reasoning, future mixed-initiative systems should also integrate complementary modalities that better support quantitative and spatial thinking.

## 8.2 Redirecting Cognitive Effort and Supporting Productive Friction

Most interactive systems aim to reduce user effort under the assumption that easier interactions lead to better outcomes. However, cognitive load theory distinguishes between productive cognitive effort that supports learning and extraneous overhead that impedes it [38, 92]. Rather than simply minimizing effort, our findings suggest the importance of redirecting cognitive resources toward valuable reasoning activities. Our findings highlight that participants using NS reported significantly lower mental demand, effort, and frustration while engaging in more complex reasoning. For example, they generated 2.3 times more insights and conducting 2.7 times more reflection events. Also, they engaged in reflection throughout the whole process instead of only towards the end. This pattern suggests that NS reshaped how participants allocated cognitive resources during analysis.

These results suggest a nuanced view of the relationship between analysis depth and cognitive burden. Our findings reveal two distinct types of cognitive load in analytical work. NS reduced overhead load by eliminating cognitive friction around procedural tasks. As  $P_{12}$  noted: "I didn't have to remember where I wrote things down or try to piece together my thinking." Simultaneously, NS redirected effort toward reasoning activities (reflection, comparison, and synthesis) rather than procedural management.  $P_{10}$  described this shift: "Instead of spending time on how to organize my thoughts, I now focus on what the data actually telling me." This redistribution of cognitive effort explains why participants could engage in more sophisticated reasoning while experiencing less subjective difficulty.

The system's workflow effectively managed the demands of tracking and reconnecting insights, allowing participants to focus on *interpretive reasoning*.

Overall, our study highlights the importance of optimizing the type of effort users expend rather than simply minimizing user effort, suggesting that design evaluation should focus on whether systems enhance reasoning-focused effort, rather than prioritizing workload reduction solely.

## 8.3 Design Implications for Mixed-Initiative Analytical Systems

A key challenge in designing mixed-initiative analytical systems is determining when and how AI should participate in users' reasoning processes. Our findings reveal that effective analysis requires treating exploration, reflection, and articulation as an integrated process. We identified three interdependent stages: (1) divergence, which creates raw material by generating multiple alternatives and factors; (2) reflection, which evaluates and refines these alternatives; and (3) integration, which incorporates them into defensible conclusions.

Our findings suggest these stages are interdependent. Divergence without reflection produces scattered exploration; reflection without alternatives limits reasoning depth; integration depends on both. This interdependence suggests that individual features, while valuable, may provide incomplete support for deep reasoning. Our findings indicate that sophisticated analysis benefits from systemic design where exploration, reflection, and integration are scaffolded as interconnected processes rather than isolated capabilities.

*Supporting Divergence Through Alternative Engagement.* Effective divergence depends on helping users engage with multiple alternatives, rather than merely producing superficial options [72, 86]. Our bidirectional narrative–evidence coupling illustrates this. Features like *Show View* and *Capture* let users test a narrative hunch with visualizations, then save what they discover back into their reasoning. This cycle helps alternatives develop enough depth to support later reflection. While not every system needs a narrative-first interface, the principle applies broadly: mixed-initiative systems can support divergence by enabling exploration of multiple paths that help develop competing explanations from initial hunches into well-reasoned, evidence-backed alternatives.

*Provenance Needs to Capture Why, Not Just What.* Productive reflection requires access to the underlying reasoning, not just a record of interactions. Our insight timeline and branching features illustrate one approach: when users create branches or mark shifts, the system captures both the change and its surrounding narrative context, making earlier reasoning states revisitable and comparable. Whether using narratives or other forms, the aim is to make the evolution of reasoning both visible and revisitable. By capturing the reasons behind users' changes in direction, rather than merely tracking their clicks, these systems foster genuine reflection instead of simple interaction replay.

## 8.4 Supporting Reasoning Evolution Across Domains

The principles underlying NS address a challenge that extends beyond data analysis: how to support reasoning that develops over time. A clear example is intelligence analysis, where analysts must synthesize fragmented information into coherent assessments while remaining open to contradictory evidence. NS suggests potential to track reasoning from initial hypotheses to evidence integration and final judgments, helping analysts maintain clear records of their reasoning and remain open to alternative explanations on ambiguous problems [30, 69]. Similar needs arise in policy development, where decision-makers must manage unresolved questions, weigh competing factors, and show how recommendations evolve from evidence [26]. In research contexts, tools could likewise benefit from insight timelines that track the progression of hypotheses across literature reviews, experiments, and analysis.

More broadly, these examples highlight a shared design gap: most platforms capture either final outputs (documents, decisions) or low-level interactions (clicks, searches) but miss the interpretive work that connects them. By treating reasoning threads as first-class objects, systems could help users revisit earlier assumptions, compare alternative explanations, and refine conclusions over time. Whether in intelligence, policy, or research, our framework highlights opportunities for system architectures that make reasoning evolution explicit, which may improve transparency, accountability, and coherence across domains.

## 8.5 Ethical Implications for Analytical Integrity

AI-assisted analysis systems raise concerns about analytical integrity when algorithmic decisions become invisible to users [25]. We discuss three areas where NS's design could unintentionally affect reasoning quality: algorithmic bias in view generation, narrative framing effects, and system steering through workflow design.

*Algorithmic Bias in Semantic Alignment.* While utilizing LLMs for semantic alignment enables efficient translation from narrative to visualizations, the model generation could contain inherent biases that negatively impact the quality of analysis. Studies have shown that LLMs could exhibit systematic errors in query generation, inconsistent interpretation and visualization selection. When a user writes "Asia is cheaper but more crowded," the system makes multiple interpretive choices: which columns represent "cheap," what aggregation level defines regions, and which chart types show comparisons. These encoded assumptions can privilege certain interpretations while obscuring others. To mitigate these risks, NS maintains human control over AI-generated outputs. Views are generated on explicit user request rather than automatically, and users decide whether to include them in their narrative. The bidirectional capture mechanism provides additional context to prompt context-aware generation. This human-in-the-loop approach ensures that AI assists rather than dictates analytical reasoning. However, significant risks remain. Users may not recognize when generated queries misrepresent their intent or when visualizations appear plausible but encode problematic assumptions. Future interventions can further mitigate these risks by exposing generated SQL queries for user verification, providing confidence scores for query correctness,

and explaining the reasoning behind visualization choices to help users recognize when outputs misrepresent their analytical intent.

*Narrative Framing Effects.* Confirmation bias is a risk that analysts actively work to avoid in rigorous analysis. NS explicitly surfaces this risk by enabling users to start exploration with narrative hypotheses. When users write claims and the system generates aligned views, they may selectively interpret patterns as confirming their narrative while overlooking contradictions. Following design guidelines that emphasize exploring alternatives to reduce confirmation bias, NS supports narrative branching to explore alternative explanations with proactive reflection for revisit. The inquiry board maintains visibility of unresolved questions, encouraging users to examine competing interpretations. However, these mechanisms rely on users recognizing when their reasoning has narrowed and taking action to branch or reflect. Users deeply committed to their initial narrative may not recognize the need for alternatives. Future systems could explore proactive interventions that detect when narrative commitment narrows exploration and suggest contradictory evidence or alternative framings to maintain analytical rigor.

*System Steering Through Workflow Design.* While NS reduces cognitive overhead, the workflow design creates structural pressures. Prior work identifies steering challenges in AI-assisted analysis and proposes task decomposition to address them [40]. Narrative-driven exploration faces similar challenges at the workflow level and introduces friction through explicit actions for generating views, capturing insights, and creating branches. However, this friction occurs at execution points rather than validation points. Users must independently recognize when alternatives are needed or when reasoning should be questioned, and the system's editorial choices during synthesis remain implicit. While NS makes branching and reflection visible workflow components, these depend on users initiating them. Future systems could explore friction-induced interventions that prompt alternative exploration when analytical patterns suggest narrowing, or make the system's synthesis criteria more transparent.

## 9 Limitations and Future Work

### 9.1 Open-Ended Exploration Without Ground-Truth Validation

Our evaluation focused on interpretive sensemaking in open-ended exploratory contexts. We selected tasks with multiple valid interpretations where correctness is context-dependent, and assessed process quality rather than outcome correctness. The datasets have no ground-truth answers, as different stakeholders might legitimately prioritize different factors.

This design choice means our findings may not generalize to high-stakes analytical contexts where correctness and bias have serious consequences [50, 96]. In domains such as medical diagnosis [76] or policy evaluation [44], reducing confirmation bias and ensuring analytical rigor are critical. While our study suggests NS supports broader exploration and deeper reflection, we cannot claim these process improvements translate into reduced bias or increased accuracy without targeted evaluation in such contexts.

The limitation extends beyond measurement to the appropriateness of design. High-stakes domains may require scaffolding

mechanisms different from those we implemented. For instance, medical diagnosis benefits from structured checklists that enforce consideration of differential diagnoses, while intelligence analysis requires explicit documentation of assumptions and confidence levels. Our narrative-first approach prioritizes flexible exploration and interpretive freedom, which may be less appropriate when analytical protocols must be followed rigorously or when premature narrative commitment could reinforce dangerous biases.

While NS incorporates design principles aimed at mitigating cognitive bias, their effectiveness in reducing analytical errors remains unverified. Future research should extend this work to high-stakes domains with ground-truth solutions to empirically assess where and how analytical errors occur during narrative-driven exploration. This includes developing methods to detect when reasoning diverges from accuracy, designing interventions that interrupt problematic patterns (e.g., premature commitment to flawed narratives), and understanding optimal timing and modality for such interventions. By studying how narrative scaffolding performs when correctness can be verified, we can identify specific failure modes and develop safeguards for contexts where analytical errors have serious consequences.

## 9.2 Supporting Convergence in Divergent Exploration

NS creates scaffolds for divergence through branching and alternative exploration but relies heavily on user-driven reflection to synthesize and converge insights. Without active reflection, participants may end up with scattered and unproductive exploration. Our analysis supports convergence through reflection was necessary for productive outcomes. However, systems could better support convergence processes through automated synthesis suggestions or propagation mechanisms that maintain coherence across alternative reasoning paths as users refine and integrate scattered explorations [87]. This could include developing mechanisms to detect when users have sufficient alternatives for meaningful comparison and suggesting synthesis strategies.

## 10 Conclusion

Narrative-driven exploration represents a natural mode of human reasoning with data, yet it carries inherent risks of premature commitment and confirmation bias. Rather than avoiding this approach, we argue the path forward lies in developing AI systems that support and safeguard how people naturally think with data. In this work, we introduced the Narrative Scaffolding framework to address persistent challenges in data-driven sensemaking, particularly in supporting workflow fluency, reflection, divergence, and explainability. Building on the view that data exploration is a process of interpretive reasoning, Narrative Scaffolding treats narrative construction not as an after-the-fact report but as the primary interface for reasoning, coupling writing with semantically aligned view generation, provenance tracking, and branching mechanisms. Our formative studies revealed four key challenges in managing evolving reasoning, and our evaluations of the Narrative Scaffolding system that instantiates the framework showed that it enables more fluid workflows, deeper reflection, broader exploration of alternatives, and more defensible narratives. This work opens new ground for AI-integrated data analysis by treating writing as a reasoning

interface for narrative-driven sensemaking. We surface key challenges in balancing human initiative with system guidance and managing trade-offs between exploratory breadth and analytical precision, providing a foundation for future systems that enhance rather than suppress human reasoning with data.

## Usage of Generative AI Statement

We disclose all instances of GenAI use throughout this research:

- **Prototype development:** Claude-4 was used to assist with coding support during prototype implementation. All generated code was reviewed, modified, and tested by the authors to ensure correctness.
- **Data analysis:** Claude-4 was used to assist with Python notebooks and D3 charts for data analysis. The authors reviewed, validated, and interpreted all analysis results independently.
- **Manuscript writing:** GPT-5 and Claude-4 were used for grammar refinement, sentence restructuring, and improving the clarity and flow of the manuscript. Assistance was limited to language editing and organization.

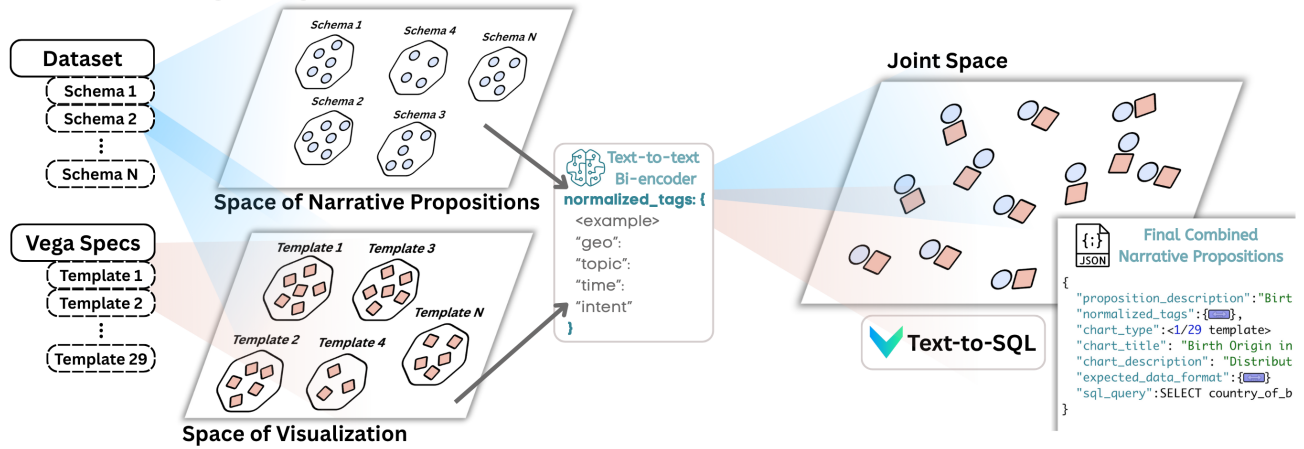
## References

- [1] Simon Attfield and John Dowell. 2003. Information seeking and use by newspaper journalists. *Journal of documentation* 59, 2 (2003), 187–204.
- [2] Aaron Bangor, Philip Kortum, and James Miller. 2009. Determining what individual SUS scores mean: Adding an adjective rating scale. *Journal of usability studies* 4, 3 (2009), 114–123.
- [3] Lyn Bartram, Michael Correll, and Melanie Tory. 2021. Untidy data: The unreasonable effectiveness of tables. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 686–696.
- [4] Andrea Batch and Niklas Elmqvist. 2018. The Interactive Visualization Gap in Initial Exploratory Data Analysis. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan 2018), 278–287.
- [5] Yoav Benjamini and Yoel Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* 57, 1 (1995), 289–300.
- [6] Charles Berret and Tamara Munzner. 2024. Iceberg Sensemaking: A Process Model for Critical Data Analysis. *IEEE Transactions on Visualization and Computer Graphics* (2024).
- [7] John Brooke et al. 1996. SUS-A quick and dirty usability scale. *Usability evaluation in industry* 189, 194 (1996), 4–7.
- [8] Jerome Bruner. 1991. The narrative construction of reality. *Critical inquiry* 18, 1 (1991), 1–21.
- [9] Steven P Callahan, Juliana Freire, Emanuele Santos, Carlos E Scheidegger, Cláudio T Silva, and Huy T Vo. 2006. VisTrails: visualization meets data management. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*. 745–747.
- [10] Akhilesh Camisetty, Chaitanya Chandurkar, Maoyuan Sun, and David Koop. 2018. Enhancing web-based analytics applications through provenance. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 131–141.
- [11] Jeff Conklin and Michael L Begeman. 1988. gIBIS: A hypertext tool for exploratory policy discussion. *ACM Transactions on Information Systems (TOIS)* 6, 4 (1988), 303–331.
- [12] Zach Cutler, Kiran Gadhave, and Alexander Lex. 2020. Ttrack: A library for provenance-tracking in web-based visualizations. In *2020 IEEE Visualization Conference (VIS)*. IEEE, 116–120.
- [13] Michael F Dahlstrom. 2014. Using narratives and storytelling to communicate science with nonexpert audiences. *Proceedings of the national academy of sciences* 111, supplement\_4 (2014), 13614–13620.
- [14] Brenda Dervin. 1998. Sense-making theory and practice: An overview of user interests in knowledge seeking and use. *Journal of knowledge management* 2, 2 (1998), 36–46.
- [15] Wenwen Dou, Dong Hyun Jeong, Felesia Stukes, William Ribarsky, Heather Richter Lipford, and Remco Chang. 2009. Recovering reasoning processes from user interactions. *IEEE computer graphics and applications* 29, 3 (2009), 52–61.
- [16] David W Eccles and Güler Arsal. 2017. The think aloud method: what is it and how do I use it? *Qualitative Research in Sport, Exercise and Health* 9, 4 (2017), 514–531.

- [17] Mi Feng, Cheng Deng, Evan M. Peck, and Lane Harrison. 2017. HindSight: Encouraging Exploration through Direct Encoding of Personal Interaction History. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan. 2017), 351–360. <https://doi-org.myaccess.library.utoronto.ca/10.1109/TVCG.2016.2599058>
- [18] Yu Fu, Dennis Bromley, and Vidya Setlur. 2025. DataWeaver: Authoring Data-Driven Narratives through the Integrated Composition of Visualization and Text. In *Computer Graphics Forum*. e70098.
- [19] Yu Fu and John Stasko. 2023. More than data stories: Broadening the role of visualization in contemporary journalism. *IEEE Transactions on Visualization and Computer Graphics* (2023).
- [20] Tong Gao, Mira Dontcheva, Eytan Adar, Zhicheng Liu, and Karrie G Karahalios. 2015. Datatone: Managing ambiguity in natural language interfaces for data visualization. In *Proceedings of the 28th annual acm symposium on user interface software & technology*. 489–500.
- [21] David Gotz and Michelle X Zhou. 2009. Characterizing users' visual analytic activity for insight provenance. *Information Visualization* 8, 1 (2009), 42–55.
- [22] Samuel Gratzl, Alexander Lex, Nils Gehlenborg, Nicola Cosgrove, and Marc Streitz. 2016. From visual exploration to storytelling and back again. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 491–500.
- [23] Garrett Grolemond and Hadley Wickham. 2014. A cognitive interpretation of data analysis. *International Statistical Review* 82, 2 (2014), 184–204.
- [24] Yi Guo, Danqing Shi, Mingjuan Guo, Yanqiu Wu, Nan Cao, and Qing Chen. 2024. Talk2Data: A natural language interface for exploratory visual analysis via question decomposition. *ACM Transactions on Interactive Intelligent Systems* 14, 2 (2024), 1–24.
- [25] Sunwoo Ha, Shayan Monadjemi, and Alvitia Ottley. 2024. Guided By AI: Navigating Trust, Bias, and Data Exploration in AI-Guided Visual Analytics. In *Computer Graphics Forum*, Vol. 43. Wiley Online Library, e15108.
- [26] Hanne Foss Hansen and Olaf Rieper. 2010. The politics of evidence-based policy-making: The case of Denmark. *German Policy Studies* 6, 2 (2010), 87.
- [27] Idit Ed Harel and Seymour Ed Papert. 1991. *Constructionism*. Ablex Publishing.
- [28] Sandra G Hart. 2006. NASA-task load index (NASA-TLX); 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting*, Vol. 50. Sage publications Sage CA: Los Angeles, CA, 904–908.
- [29] Jeffrey Heer, Jock Mackinlay, Chris Stolte, and Maneesh Agrawala. 2008. Graphical histories for visualization: Supporting analysis, communication, and evaluation. *IEEE transactions on visualization and computer graphics* 14, 6 (2008), 1189–1196.
- [30] Richards J Heuer. 1999. *Psychology of intelligence analysis*. Center for the Study of Intelligence.
- [31] Oliver Huang and Carolina Nobre. 2025. Toward Supporting Narrative-Driven Data Exploration: Barriers and Design Opportunities. *arXiv preprint arXiv:2508.04920* (2025).
- [32] Oliver Huang and Carolina Nobre. 2025. ViStruct: Simulating Expert-Like Reasoning Through Task Decomposition and Visual Attention Cues. *arXiv preprint arXiv:2506.21762* (2025).
- [33] Jessica Hullman. 2022. Designing for judgment in data analysis. *Commun. ACM* 65, 1 (2022), 62–71.
- [34] Mohammed Saidul Islam, Md Tahmid Rahman Laskar, Md Rizwan Parvez, Enamul Hoque, and Shafiq Joty. 2024. DataNarrative: Automated data-driven storytelling with visualizations and texts. *arXiv preprint arXiv:2408.05346* (2024).
- [35] Peiling Jiang, Jude Rayan, Steven P Dow, and Haijun Xia. 2023. Graphologue: Exploring large language model responses with interactive diagrams. In *Proceedings of the 36th annual ACM symposium on user interface software and technology*. 1–20.
- [36] Nazanin Kadivar, Victor Chen, Dustin Dunsmuir, Eric Lee, Cheryl Qian, John Dill, Christopher Shaw, and Robert Woodbury. 2009. Capturing and supporting the analysis process. In *2009 IEEE Symposium on Visual Analytics Science and Technology*. Ieee, 131–138.
- [37] Alex Kale, Matthew Kay, and Jessica Hullman. 2022. DisCERN: Mitigating Bias in Data Exploration via Analysis History Visualizations. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. ACM, 1–16.
- [38] Slava Kalyuga. 2011. Cognitive load theory: How many types of load does it really need? *Educational psychology review* 23, 1 (2011), 1–19.
- [39] Sean Kandel, Andreas Paepcke, Joseph M Hellerstein, and Jeffrey Heer. 2012. Enterprise data analysis and visualization: An interview study. *IEEE transactions on visualization and computer graphics* 18, 12 (2012), 2917–2926.
- [40] Majeed Kazemitabaar, Jack Williams, Ian Drosos, Tovi Grossman, Austin Zachary Henley, Carina Negreanu, and Advait Sarkar. 2024. Improving steering and verification in AI-assisted data analysis with interactive task decomposition. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–19.
- [41] Brian Felipe Keith Norambuena, Tanushree Mitra, and Chris North. 2022. Design guidelines for narrative maps in sensemaking tasks. *Information Visualization* 21, 3 (2022), 220–245.
- [42] Gary Klein, Jennifer K Phillips, Erica L Rall, and Deborah A Peluso. 2007. A data–frame theory of sensemaking. In *Expertise out of context*. Psychology Press, 118–160.
- [43] Gary Klein, Jennifer K. Phillips, Emily L. Rall, and David A. Peluso. 2007. A data-frame theory of sensemaking. *Expertise out of context* (2007), 113–155.
- [44] Johan E Korteling, Geerte L Paradies, and Josephine P Sassen-van Meer. 2023. Cognitive bias and how to improve sustainable decision making. *Frontiers in Psychology* 14 (2023), 1129835.
- [45] Werner Kunz and Horst WJ Rittel. 1970. Issues as elements of information systems. (1970).
- [46] Shahid Latif, Zheng Zhou, Yoon Kim, Fabian Beck, and Nam Wook Kim. 2021. Kori: Interactive synthesis of text and charts in data documents. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 184–194.
- [47] Bongshin Lee, Nathalie Henry Riche, Petra Isenberg, and Sheelagh Carpendale. 2015. More than telling a story: Transforming data into visually shared stories. *IEEE computer graphics and applications* 35, 5 (2015), 84–90.
- [48] Haotian Li, Yun Wang, and Huamin Qu. 2024. Where are we so far? understanding data storytelling tools from the perspective of human-ai collaboration. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [49] Haotian Li, Lu Ying, Haidong Zhang, Yingcai Wu, Huamin Qu, and Yun Wang. 2023. Notable: On-the-fly Assistant for Data Storytelling in Computational Notebooks. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 173, 16 pages. doi:10.1145/3544548.3580965
- [50] Shiyao Li, Thomas James Davidson, Cindy Xiong Bearfield, and Emily Wall. 2025. Confirmation Bias: The Double-Edged Sword of Data Facts in Visual Data Communication. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 1175, 16 pages. doi:10.1145/3706598.3713831
- [51] Heather Richter Lipford, Felesia Stukes, Wenwen Dou, Matthew E Hawkins, and Remco Chang. 2010. Helping users recall their reasoning process. In *2010 IEEE Symposium on Visual Analytics Science and Technology*. IEEE, 187–194.
- [52] Maxim Lisnic, Vidya Setlur, and Nicole Sultanum. 2025. Plume: Scaffolding Text Composition in Dashboards. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [53] Zhicheng Liu and John Stasko. 2010. Mental models, visual reasoning and interaction in information visualization: A top-down perspective. *IEEE transactions on visualization and computer graphics* 16, 6 (2010), 999–1008.
- [54] Karthic Madanagopal, Eric D Ragan, and Perakath Benjamin. 2019. Analytic provenance in practice: The role of provenance in real-world visualization and data analysis environments. *IEEE Computer Graphics and Applications* 39, 6 (2019), 30–45.
- [55] Andreas Mathisen, Tom Horak, Clemens Nylandstedt Klokmose, Kaj Grønbaek, and Niklas Elmqvist. 2019. InsideInsights: Integrating data-driven reporting in collaborative visual analytics. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 649–661.
- [56] Microsoft. 2024. Create a narrative Visual with Copilot for Power BI. <https://learn.microsoft.com/en-us/power-bi/create-reports/copilot-create-narrative?tabs=powerbi-service>. Accessed: 2025-08-30.
- [57] Microsoft. 2024. Introduction to Copilot in Power BI. <https://learn.microsoft.com/en-us/power-bi/create-reports/copilot-introduction>. Accessed: 2025-08-30.
- [58] Michael Murray. 2003. Narrative psychology and narrative analysis. *Palliative Medicine in Practice* 2, 3 (2003), 111–120.
- [59] Michael Murray et al. 2015. Narrative psychology. *Qualitative psychology: A practical guide to research methods* (2015), 85–107.
- [60] Arpit Narechania, Shunan Guo, Eunye Koh, Alex Endert, and Jane Hoffswell. 2025. Utilizing Provenance as an Attribute for Visual Data Analysis: A Design Probe with ProvenanceLens. *IEEE Transactions on Visualization and Computer Graphics* (2025), 1–14. doi:10.1109/TVCG.2025.3571708
- [61] Arpit Narechania, Kaustubh Odak, Mennatallah El-Assady, and Alex Endert. 2025. ProvenanceWidgets: A Library of UI Control Elements to Track and Dynamically Overlay Analytic Provenance. *IEEE Transactions on Visualization and Computer Graphics* 31, 1 (2025), 1235–1245. doi:10.1109/TVCG.2024.3456144
- [62] Arpit Narechania, Arjun Srinivasan, and John Stasko. 2020. NL4DV: A toolkit for generating analytic specifications for data visualization from natural language queries. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 369–379.
- [63] Raymond S Nickerson. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology* 2, 2 (1998), 175–220.
- [64] Chris North. 2006. Toward measuring visualization insight. *IEEE computer graphics and applications* 26, 3 (2006), 6–9.
- [65] Chris North, Remco Chang, Alex Endert, Wenwen Dou, Richard May, Bill Pike, and Glenn Fink. 2011. Analytic provenance: process+ interaction+ insight. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*. 33–36.
- [66] Lorelli S Nowell, Jill M Norris, Deborah E White, and Nancy J Moules. 2017. Thematic analysis: Striving to meet the trustworthiness criteria. *International journal of qualitative methods* 16, 1 (2017), 1609406917733847.

- [67] Jason Obeid and Enamul Hoque. 2020. Chart-to-text: Generating natural language descriptions for charts by adapting the transformer model. *arXiv preprint arXiv:2010.09142* (2020).
- [68] Kevin Padian. 2018. Narrative and "anti-narrative" in science: how scientists tell stories, and don't. *Integrative and Comparative Biology* 58, 6 (2018), 1224–1234.
- [69] Randolph H Pherson and Richards J Heuer Jr. 2019. *Structured analytic techniques for intelligence analysis*. Cq Press.
- [70] Peter Pirolli and Stuart Card. 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis*, Vol. 5. McLean, VA, USA, 2–4.
- [71] Prolific. 2024. Prolific. <https://www.prolific.com>. Version used: July 2025.
- [72] Kevin Pu, KJ Kevin Feng, Tovi Grossman, Tom Hope, Bhavana Dalvi Mishra, Matt Latzke, Jonathan Bragg, Joseph Chee Chang, and Pao Siangliulue. 2025. Ideasynt: Iterative research idea development through evolving and composing idea facets with literature-grounded feedback. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–31.
- [73] Eric D Ragan, Alex Endert, Jibonananda Sanyal, and Jian Chen. 2015. Characterizing provenance in visualization and data analysis: an organizational framework of provenance types and purposes. *IEEE transactions on visualization and computer graphics* 22, 1 (2015), 31–40.
- [74] Adam Rule, Aurélien Tabard, and James D Hollan. 2018. Exploration and explanation in computational notebooks. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–12.
- [75] Morgan Saletta, Ariel Kruger, Tamar Primoratz, Ashley Barnett, Tim van Gelder, and Robert E Horn. 2020. The role of narrative in collaborative reasoning and intelligence analysis: A case study. *Plos one* 15, 1 (2020), e0226981.
- [76] Gustavo Saposnik, Donald Redelmeier, Christian C Ruff, and Philippe N Tobler. 2016. Cognitive biases associated with medical decisions: a systematic review. *BMC medical informatics and decision making* 16, 1 (2016), 138.
- [77] Mike Scaife and Yvonne Rogers. 1996. External cognition: how do graphical representations work? *International journal of human-computer studies* 45, 2 (1996), 185–213.
- [78] Vidya Setlur. 2025. Supporting Human-Centric Data Exploration Through Semantics and Natural Language Interaction. In *Companion of the 2025 International Conference on Management of Data (Berlin, Germany) (SIGMOD/PODS '25)*. Association for Computing Machinery, New York, NY, USA, 851–854. doi:10.1145/3722212.3725628
- [79] Vidya Setlur, Sarah E Battersby, Melanie Tory, Rich Gossweiler, and Angel X Chang. 2016. Eviza: A natural language interface for visual analysis. In *Proceedings of the 29th annual symposium on user interface software and technology*. 365–377.
- [80] Vidya Setlur, Michael Correll, Arvind Satyanarayan, and Melanie Tory. 2024. Heuristics for Supporting Cooperative Dashboard Design. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (Jan. 2024), 370–380. doi:10.1109/TVCG.2023.3327158
- [81] Danqing Shi, Xinyue Xu, Fuling Sun, Yang Shi, and Nan Cao. 2020. Calliope: Automatic visual data story generation from a spreadsheet. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 453–463.
- [82] Arjun Srinivasan, Steven M Drucker, Alex Endert, and John Stasko. 2018. Augmenting visualizations with interactive data facts to facilitate interpretation and communication. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 672–681.
- [83] Arjun Srinivasan, Vidya Setlur, and Arvind Satyanarayan. 2025. Pluto: Authoring Semantically Aligned Text and Charts for Data-Driven Communication. In *Proceedings of the 30th International Conference on Intelligent User Interfaces (IUI '25)*. Association for Computing Machinery, New York, NY, USA, 1123–1140. doi:10.1145/3708359.3712122
- [84] Charles D. Stolper, Adam Perer, and David Gotz. 2014. Progressive Visual Analytics: User-Driven Visual Exploration of In-Progress Analytics. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 1653–1662. doi:10.1109/TVCG.2014.2346574
- [85] Hari Subramonyam, Roy Pea, Christopher Pondoc, Maneesh Agrawala, and Colleen Seifert. 2024. Bridging the gulf of envisioning: Cognitive challenges in prompt based interactions with llms. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [86] Sangho Suh, Meng Chen, Bryan Min, Toby Jia-Jun Li, and Haijun Xia. 2024. Luminare: Structured generation and exploration of design space with large language models for human-ai co-creation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–26.
- [87] Sangho Suh, Michael Lai, Kevin Pu, Steven P Dow, and Tovi Grossman. 2025. StoryEnsemble: Enabling Dynamic Exploration & Iteration in the Design Process with AI and Forward-Backward Propagation. *arXiv preprint arXiv:2508.03182* (2025).
- [88] Nicole Sultanum, Fanny Chevalier, Zoya Bylinskii, and Zhicheng Liu. 2021. 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*. 1–17.
- [89] Nicole Sultanum and Vidya Setlur. 2024. From instruction to insight: Exploring the functional and semantic roles of text in interactive dashboards. *IEEE Transactions on Visualization and Computer Graphics* (2024).
- [90] Nicole Sultanum and Arjun Srinivasan. 2023. Datatales: Investigating the use of large language models for authoring data-driven articles. In *2023 IEEE Visualization and Visual Analytics (VIS)*. IEEE, 231–235.
- [91] Mengdi Sun, Ligan Cai, Weiwei Cui, Yanqiu Wu, Yang Shi, and Nan Cao. 2022. Erato: Cooperative data story editing via fact interpolation. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2022), 983–993.
- [92] John Sweller. 1994. Cognitive load theory, learning difficulty, and instructional design. *Learning and instruction* 4, 4 (1994), 295–312.
- [93] Melanie Tory, Lyn Bartram, Brittany Fiore-Gartland, and Anamaria Crisan. 2021. Finding their data voice: Practices and challenges of dashboard users. *IEEE Computer Graphics and Applications* 43, 1 (2021), 22–36.
- [94] Melanie Tory and Torsten Moller. 2004. Human factors in visualization research. *IEEE transactions on visualization and computer graphics* 10, 1 (2004), 72–84.
- [95] Vanna.AI Team. 2024. *Vanna.AI: A retrieval-augmented text-to-SQL framework*. <https://vanna.ai/>. Accessed: 2025-08-29.
- [96] Emily Wall, Leslie M Blaha, Lyndsey Franklin, and Alex Endert. 2017. Warning, bias may occur: A proposed approach to detecting cognitive bias in interactive visual analytics. In *2017 IEEE conference on visual analytics science and technology (vast)*. IEEE, 104–115.
- [97] Huichen Will Wang, Larry Birnbaum, and Vidya Setlur. 2025. 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*. 1–24.
- [98] Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2023. Improving text embeddings with large language models. *arXiv preprint arXiv:2401.00368* (2023).
- [99] Yun Wang, Zhida Sun, Haidong Zhang, Weiwei Cui, Ke Xu, Xiaojuan Ma, and Dongmei Zhang. 2019. Datashot: Automatic generation of fact sheets from tabular data. *IEEE transactions on visualization and computer graphics* 26, 1 (2019), 895–905.
- [100] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* 35 (2022), 24824–24837.
- [101] Wesley Willett, Jeffrey Heer, Joseph Hellerstein, and Maneesh Agrawala. 2011. CommentSpace: structured support for collaborative visual analysis. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 3131–3140.
- [102] Kanit Wongsuphasawat, Yang Liu, and Jeffrey Heer. 2019. Goals, process, and challenges of exploratory data analysis: An interview study. *arXiv preprint arXiv:1911.00568* (2019).
- [103] Dylan Wootton, Amy Rae Fox, Evan Peck, and Arvind Satyanarayan. 2024. Charting EDA: Characterizing Interactive Visualization Use in Computational Notebooks with a Mixed-Methods Formalism. *IEEE Transactions on Visualization and Computer Graphics* (2024).
- [104] Ji Soo Yi, Youn-ah Kang, John T. Stasko, and Julie A. Jacko. 2008. Understanding and characterizing insights: how do people gain insights using information visualization?. In *Proceedings of the 2008 Workshop on Beyond Time and Errors: Novel Evaluation Methods for Information Visualization (Florence, Italy) (BELIV '08)*. Association for Computing Machinery, New York, NY, USA, Article 4, 6 pages. doi:10.1145/1377966.1377971
- [105] Jian Zhao, Shenyu Xu, Senthil Chandrasegaran, Chris Bryan, Fan Du, Aditi Mishra, Xin Qian, Yiran Li, and Kwan-Liu Ma. 2021. Chartstory: Automated partitioning, layout, and captioning of charts into comic-style narratives. *IEEE transactions on visualization and computer graphics* 29, 2 (2021), 1384–1399.
- [106] Chen Zhu-Tian and Haijun Xia. 2022. CrossData: Leveraging text-data connections for authoring data documents. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–15.

## Preprocessing Stage



## Excution Stage

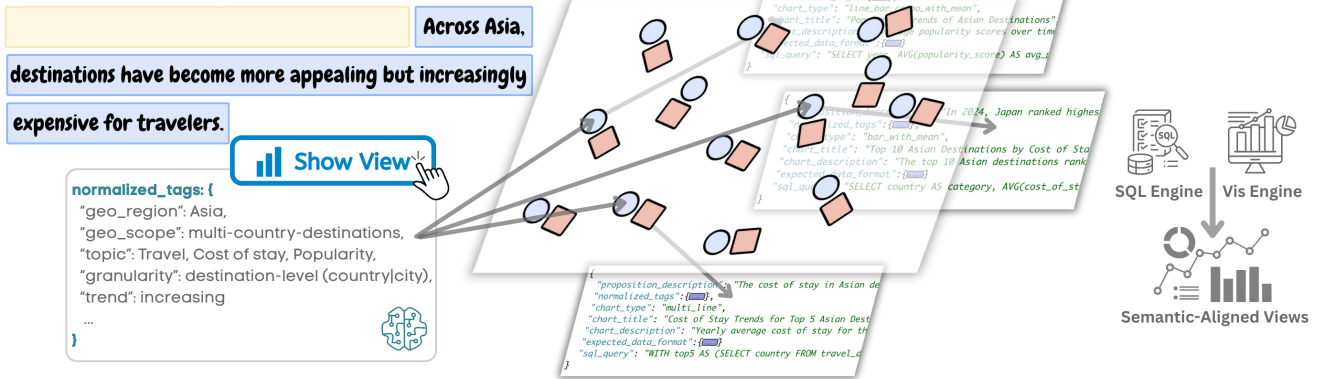


Figure 14: Semantic-aligned view generation pipeline. During preprocessing (top), the system generates schema-driven narrative proposition templates and pairs them with visualization specifications. These are normalized into shared semantic tags and embedded into a joint space using a bi-encoder architecture. During execution (bottom), user narratives are normalized into tags and matched against this joint space to retrieve relevant propositions. The system executes corresponding SQL queries, and renders aligned views.

## A Appendix

### A.1 Semantic-Aligned View Generation Pipeline

The narrative proposition space consists of structured templates grounded in the dataset schema, covering common analytical patterns such as ranking, change, composition, outlier detection, per-capita measures, and correlations. For example, from a crime dataset, the system may generate a template such as "In {year}, {borough} had the highest {metric} at {value}." Propositions are filtered with constraints (i.e. top-k limits, minimum timepoints) to ensure visibility. For statistical accuracy of the narrative propositions, templates are then instantiated with actual values through schema-driven SQL queries. Constraints, such as requiring multiple time points, further ensure that only meaningful propositions are retained.

The visualization proposition space is derived in parallel, adapting a library of Vega specifications to the dataset’s schema. Each chart template defines its expected data format (i.e. scatterplots with {x, y, group} or bar charts with {labels, values}) and is paired with an LLM-generated title and descriptions, ensuring interpretability across various views.

Narrative propositions and visualization descriptions are then normalized into shared tags (i.e. geo=, topic=, time=, and intent=) and embedded into a joint semantic space using a bi-encoder architecture [98]. The joint space consolidates all components needed for alignment. It contains the narrative proposition templates, LLM-generated visualization captions, normalized tags, the expected data formats from each visualization specification, and pre-stored SQL queries produced by a text-to-SQL model. By combining narrative and visualization in one embedding space, the system maintains semantic consistency across both directions of investigation.

When users write a narrative and request `Show view`, the system normalizes the text into tags (Figure 14, bottom). It matches against the list of narrative propositions. Based on the match, it determines whether a single chart or a group of charts provides the most appropriate alignment with the claim. The corresponding SQL queries are then executed, and the resulting data are rendered into either individual chart templates or grouped into dashboards when multiple views are needed to represent the claim.

### A.2 Prompt Used for Data Story

You are a professional data journalist. You will receive an exploration\_path as JSON (array of objects). Each object includes at least:

##### INPUT #####

sentence\_id  
sentence\_content  
drift\_type  
other optional fields may be present

##### TASK #####

Produce data\_story as a JSON array of 8--15 objects. Each object must contain:

"data\_story\_sentence": one concise, professional, data-driven sentence suitable for publication.  
"ref\_id": a single sentence\_id or an array of sentence\_ids from the exploration\_path that support this sentence.

##### GUIDELINES #####

Summarize evidence-driven insights; do not describe analysis steps or methods.  
Use objective, newsroom style (e.g., "The data shows," "Analysis reveals," "Destinations such as X and Y emerge").  
Cover the main branches of the exploration and converge on key conclusions.  
You may condense multiple related exploration sentences into one story point; if so, include all supporting sentence\_ids in ref\_id (use an array).  
Every ref\_id must exist in the exploration\_path.  
Avoid self-referential language (no "I," "we," or "the model").  
Do not invent facts; rely only on the content of exploration\_path.  
Keep each sentence concise and fact-focused.

##### ORDER AND SCOPE #####

Sequence the data\_story to follow a logical narrative from overview to key findings to conclusions. Ensure the story spans the breadth of the exploration\_path.

##### STYLE #####

Professional, objective, and data-driven tone.  
Present tense unless the data explicitly reference past periods.  
Avoid speculation; qualify claims only when supported by the data.

##### OUTPUT FORMAT #####

Output only a valid JSON array.  
Do not include any text before or after the JSON.  
Each element must include exactly "data\_story\_sentence" and "ref\_id" keys.

##### OPENING REQUIREMENT #####

The first data\_story\_sentence must begin with one of: "Across ", "From ", "Behind every ", "Focusing on ", or "In ".

### A.3 Prompt Used for Alternative Suggestions

You are an assistant helping a user reflect on their reasoning as they explore data and build a personal narrative.

Your reflection should help the user:

- Reconsider implicit assumptions
- Explore alternative explanations

```

- Think about branching to resolve contradictions

You must keep the reflection grounded within the scope of the provided datasets.

#####
RELEVANT DATA SOURCES:
The following dataset categories have been identified as relevant to the current insight. Stay within
  these domains when suggesting reflection prompts:

${JSON.stringify(firstStepData.related_source, null, 2)}

Only use reasoning related to these domains. Do not speculate outside them.

#####
CONTEXTUAL INSIGHTS:
Here are potentially related or contradictory insights previously written by the user. You may use them
  to explain reasoning tension or provide contrast.

${JSON.stringify(relatedSentencesContext, null, 2)}

Each item includes:
- node_id
- sentence_id
- sentence_content
- reason_for_relevance (e.g., "same topic", "opposite claim", "same borough", "time mismatch")

You may reference these insights in your reflection, but only if they meaningfully relate to the current
  sentence.

#####
INPUT:
{
  "node_id": ${firstStepData.node_id},
  "sentence_id": "${firstStepData.sentence_id}",
  "sentence_content": "${firstStepData.sentence_content}"
}

#####
OUTPUT FORMAT:
Return a single JSON object in this format:
{
  "node_id": ${firstStepData.node_id},
  "sentence_id": "${firstStepData.sentence_id}",
  "sentence_content": "${firstStepData.sentence_content}",
  "reflect": [
    {
      "prompt": "<detailed reflection question>",
      "reason": "<why this prompt is relevant>",
      "related_sentence": {
        "node_id": <number>,
        "sentence_content": "<string>"
      } | null
    }
  ]
}

If you cannot find anything to reflect on, return "reflect": []

```

#### A.4 Prompt Used for Interaction Capture

You are an assistant helping a user write a data-supported personal narrative. The user is exploring a dashboard to decide which borough in London to move to, based on factors like safety, income, housing, education, and quality of life.

They have just interacted with a specific view on the dashboard and clicked "Capture Insight." Your task is to write one concise, readable narrative sentence based on the data they were exploring.

**\*\*Your role\*\*:** write one concise, data-backed narrative sentence that would be a *natural follow-up* to the user's current sentence, using **ONLY** the evidence contained in the interaction log.

**\*\*Avoid\*\*:** Jargon, vague summaries, or speculative statements. Use the actual numbers and trends from the data. The sentence should sound human and fit into a longer narrative paragraph.

#####

INPUT JSON

#####

```
{
  "narrative_context": <string>,          // a full narration of the user's exploration
  "current_sentence": <string>,           // the sentence they just wrote or are editing
  "interaction_log": [                    // EXACTLY the last 5 dashboard events
    /* ARRAY IS IN CHRONOLOGICAL ORDER:
       index 0 = oldest, index 4 = most recent (most important) */
    {
      "elementId": <string>,              // unique view id
      "elementName": <string>,            // human-readable name
      "elementType": <"chart"|"map" ...>, // what kind of view
      "action": <string>,                 // e.g. "filter_change", "interactive_hover"
      "dashboardConfig": {                // static metadata about the view
        "title": <string>,
        "view_type": <string>,
        "variable_map": { ... }
      },
      "chartData": { ... } or [ ... ]     // numbers or categorical values revealed
    },
    ...
    /* 5 objects total */
  ]
}
```

**\*\*Important reading rules\*\***

1. The **\*\*4th and 5th objects (indexes 3 and 4)\*\*** represent the user's freshest focus; prefer those for your sentence.
2. Focus on whichever interaction(s) provide a clear, interesting, NUMERIC or CATEGORICAL takeaway.
  - \textbullet If multiple views are relevant, you may combine two, but keep the sentence short.
  - \textbullet Ignore hovers that do not reveal new numbers.
3. Do not invent numbers. Quote or paraphrase only what is present.
4. Write in **\*\*plain language\*\*** (journalistic tone, not code).
5. The sentence should help the user decide where to live (safety, income, diversity, housing, etc.).
6. If NO clear insight is possible, return `"narrative_suggestion": null``.

#####

YOUR SINGLE-LINE JSON OUTPUT

#####

Return a JSON object:

```
{
  "narrative_suggestion": "<one well-formed English sentence or null>",
  "source_elementId": "<elementId you relied on most>",
  "source_view_title": "<dashboardConfig.title>",
  "explanation": <explanation on why this is a reasonable followup, what insight would it provide in
    10-20 words>
}
```

## A.5 Prompt Used for Insight Timeline

You are an assistant helping a user track how their narrative insights evolve while exploring data.

```

    "related_columns": ["column1", "column2"]
  },
  "related_sentence": {
    "node_id": <number>,
    "reason": "<reason why they are related>"
  } | null
}
]

```

## A.6 Prompt Used for Inquiry Labels

You are a reasoning assistant enriching a user's inquiry graph for exploratory data analysis. The graph contains only ISSUE nodes (questions). For each issue, add (a) an optional position (likely answer), (b) an optional argument (justification), and (c) ISSUE-ISSUE links. Use ONLY what is present or clearly implied in the user's sentences. Do not speculate.

DEFINITIONS (IBIS-aligned, concise)

- position\_suggested\_by (optional): The likely or tentative ANSWER to the issue found in the user's text.
  - Include:
    - \textbullet text: short paraphrase or quote of the answer
    - \textbullet confidence: "high" (asserted/verified), "medium" (tentative), "low" (weak hint)
- argument\_suggested\_by (optional): The JUSTIFICATION supporting that position. Include:
  - \textbullet text: brief summary (data reference, mechanism, pattern, or comparison)
  - \textbullet basis: one of "data" | "mechanism" | "pattern" | "comparison" | "other"
- links (ISSUE-ISSUE only; use sparingly and only if justified):
  - \textbullet suggested\_by: Target issue was prompted as a follow-up/side question by this issue.
  - \textbullet generalized\_from: Current issue is a broader abstraction of a more specific earlier issue.
  - \textbullet specialized\_from: Current issue is a narrower/drilled-down version of a broader earlier issue.
  - \textbullet replaces: Current issue supersedes or rephrases an earlier issue (clearer scope/time/geo).

INPUT

Narrative sentences:

```

${JSON.stringify(sentenceList, null, 2)}

```

Extracted issues (from Layer 1):

```

${JSON.stringify(issueList, null, 2)}

```

```

// Each issue has: qid, title, sentenceRefs, status

```

OUTPUT

Return a JSON array where each item enriches one issue by qid. Do NOT repeat title/status/sentenceRefs. Use null for missing fields.

```

[
  {
    "qid": "<issue qid>",
    "position_suggested_by": {
      "text": "<answer from user text>",
      "confidence": "low" | "medium" | "high"
    } | null,
    "argument_suggested_by": {
      "text": "<justification from user text>",
      "basis": "data" | "mechanism" | "pattern" | "comparison" | "other"
    } | null,
    "links": [
      {
        "qid": "<target issue qid>",
        "type": "suggested_by" | "generalized_from" | "specialized_from" | "replaces",
        "explanation": "<brief, text-grounded reason>"
      }
    ]
  }
]

```

]

## A.7 Prompt Used for Inquiry Issues

You are an assistant that analyzes the user's narrative path during exploratory analysis and identifies which sentences raise meaningful questions (issues) worth tracking.

These issues may be:

- Explicit questions (e.g. "Did X happen in 2022?")
- Implicit inquiries (e.g. "I should check whether..." or "It seems like...")
- Investigative goals that signal something unresolved or still unfolding

Each issue should:

- Be framed as a **question** (``title``)
- Be assigned a **status**:
  - ``open`` if the question is raised but unanswered
  - ``resolved`` if the narrative includes clear follow-up or confirmation
  - ``stalled`` if the question was raised earlier but left unresolved and later abandoned or deferred (not in the final sentences)
- Be grounded in one or more **sentenceRefs** (e.g., ``["s3"]``)

#####

INPUT USER NARRATIVE:

```
`${JSON.stringify(sentenceList, null, 2)}`
```

Each sentence item contains:

- `sentence_id`
- `content`

#####

OUTPUT FORMAT:

Return a JSON array of identified issues:

```
[
  {
    "qid": "<unique string ID, prefixed with iss_>",
    "title": "<question title>",
    "status": "open" | "resolved" | "stalled",
    "sentenceRefs": ["<sentence_id>", ...]
  }
]
```

If no issues are identified, return an empty array: ``[]``