

Received 24 November 2025, accepted 24 December 2025, date of publication 1 January 2026,
date of current version 6 January 2026.

Digital Object Identifier 10.1109/ACCESS.2025.3650352



Interactive Narrative Analytics: Bridging Computational Narrative Extraction and Human Sensemaking

BRIAN KEITH^{ID}

Universidad Católica del Norte, Antofagasta 1270709, Chile

e-mail: brian.keith@ucn.cl

This research is funded by the Agencia Nacional de Investigación y Desarrollo (ANID) FONDECYT de Iniciación 11250039 Project. The author is also supported by Project 202311010033-VRIDT-UCN.

ABSTRACT Information overload and misinformation create significant challenges in extracting meaningful narratives from large news collections. This paper defines the nascent field of **Interactive Narrative Analytics** (INA), which combines computational narrative extraction with interactive visual analytics to support sensemaking. INA approaches enable the interactive exploration of narrative structures through computational methods and visual interfaces that facilitate human interpretation. The field faces challenges in *scalability*, *interactivity*, *knowledge integration*, and *evaluation standardization*, yet offers promising opportunities across news analysis, intelligence, scientific literature exploration, and social media analysis. Through the combination of computational and human insight, INA addresses complex challenges in narrative sensemaking.

INDEX TERMS Human-AI collaboration, information extraction, interactive visual analytics, knowledge integration, narrative extraction, narrative sensemaking, semantic interaction, visual analytics.

I. INTRODUCTION

The digital landscape presents an unprecedented volume of information, with news and social media content growing exponentially [1], [2]. This information overload is compounded by the proliferation of misinformation and disinformation [3], [4], [5], [6], creating a complex ecosystem that challenges our ability to make sense of important events and their connections. Traditional information processing approaches have become inadequate [7], [8], highlighting the need for more sophisticated computational and interactive tools to extract meaningful narratives from large text collections [9], [10].

This paper introduces and defines a new field of research that we call **Interactive Narrative Analytics** (INA), a multidisciplinary approach combining computational narrative extraction [9], [11], [12] with interactive visual analytics [13], [14], [15] to support sensemaking [16], [17]. By coining

The associate editor coordinating the review of this manuscript and approving it for publication was Derek Abbott^{ID}.

this term, we seek to establish a distinct research domain focused on addressing the unique challenges of narrative understanding in the digital age. Unlike traditional text analytics, which often focuses on statistical patterns or isolated entities, INA emphasizes temporal, causal, and relational aspects of information, capturing how events unfold and connect over time to form coherent stories [9].

The need to establish INA as a field has become increasingly evident as analysts, journalists, researchers, and the public struggle with information complexity. The sheer volume of event data makes manual analysis virtually impossible [8], [18], [19], [20], necessitating scalable computational approaches. The dynamic nature of news events requires methods that can adapt to evolving narratives in real-time. The complexity of narratives demands sophisticated representation models capable of capturing intricate relationships. Furthermore, the subjective nature of narrative interpretation requires interactive tools that can incorporate human feedback and domain knowledge [21], [22], [23], [24].

Current narrative extraction approaches employ techniques from natural language processing, machine learning, and graph theory [25], [26], [27], but fall short in critical areas. Most existing methods lack scalability, struggling to process the massive volumes of data generated in today's information ecosystem. They typically operate as black boxes with limited transparency and interactivity, failing to leverage human intuition and domain expertise. Moreover, they rarely incorporate external knowledge, leading to incomplete or inaccurate narrative representations. Most importantly, they often lack effective evaluation metrics, making it difficult to compare different approaches [9], [28].

This proposed field would bring together computational optimization, large language models (LLMs), and knowledge representation [29], [30], [31] with interactive visualization techniques [32], [33] to create an integrated approach to narrative understanding. While individual elements of this field exist in isolation across various research domains, we argue that establishing INA as a distinct discipline would foster the cross-disciplinary collaboration needed to address the complex challenges of narrative sensemaking.

In particular, this paper seeks to **define the scope of this nascent field**, examining its theoretical foundations, potential components, anticipated research challenges, promising applications, and future research directions. As a position paper proposing a new interdisciplinary field, we synthesize key concepts from multiple disciplines rather than providing a systematic literature review. By defining this new area of research, we aim to highlight the importance of integrating computational extraction with human sensemaking capabilities, providing a conceptual framework for future research in INA. For detailed surveys of the underlying technical areas, we refer interested readers to Keith et al. [9] on narrative extraction, Santana et al. [12] on narrative processing, and Keim et al. [34] on visual analytics foundations.

II. DEFINING INTERACTIVE NARRATIVE ANALYTICS

INA represents an emerging interdisciplinary field at the intersection of computational narrative extraction, visual analytics, and human-centered sensemaking. We define INA as the systematic approach to *extracting, representing, and exploring event-based narrative structures* from large textual collections through a combination of computational methods that identify temporal, causal, and semantic relationships between events, interactive visual interfaces that support human exploration and refinement of these structures, and knowledge integration mechanisms that enhance narrative understanding with external domain knowledge.

In particular, INA focuses primarily on event-based narratives, particularly those found in news articles, social media, and similar domains where understanding temporal and causal relationships between events is crucial [9]. It emphasizes the interactive aspect of narrative analysis, acknowledging that narrative understanding requires human interpretation [23], and incorporates knowledge integration to enhance extracted narratives with external domain knowledge

[35]. Thus, unlike traditional text analytics approaches that focus on statistical patterns or isolated entities, INA is concerned with the holistic structure of narratives, including events, entities, temporal relationships, and causal connections. It extends beyond traditional narrative extraction by incorporating interactive visualization and human feedback into the extraction process itself, rather than treating these as separate downstream tasks [36].

A. DISTINGUISHING INA FROM RELATED FIELDS

To clarify INA's unique position in the research landscape, we must distinguish it from several related but distinct areas. Visual Analytics (VA) is defined as “the science of analytical reasoning facilitated by interactive visual interfaces” [37]. VA emerged from intelligence needs following the 9/11 attacks [34], [37] and provides the interactive visualization foundation upon which INA builds. However, while VA broadly addresses all data types and analytical tasks, INA specifically focuses on *narrative structures* within textual data, emphasizing temporal event sequences and causal relationships rather than general data patterns. This distinction is crucial: VA might visualize any patterns in document collections, but INA specifically seeks to extract and represent the stories those documents tell.

Intelligence and security analysis applications, which historically drove VA development [34], employ visual analytics primarily for threat detection, anomaly identification, and real-time monitoring. While INA can support intelligence applications, it extends far beyond security contexts to support narrative understanding across diverse domains including scientific literature, news media, and legal documents. The key difference lies in the analytical focus: intelligence analysis prioritizes identifying threats and anomalies, while INA emphasizes understanding narrative coherence, evolution, and the relationships between different storylines.

Event-based text mining represents another related area that shares INA's interest in events but differs fundamentally in scope and purpose. Traditional event extraction methods identify isolated events and their attributes—who, what, when, where—from text [38]. INA goes beyond this extraction to *connect events into coherent narratives*, emphasizing not just what happened but how events relate to each other, how they form storylines, and how these storylines evolve and interact over time. Where event mining might produce a database of discrete events, INA produces an interconnected narrative structure that reveals the flow and development of stories.

The distinction between INA and data storytelling or narrative visualization [22], [39] is particularly important given the overlapping terminology. Data storytelling and narrative visualization focus on *communicating* predetermined insights through carefully crafted narrative techniques—they help authors tell stories with data. In contrast, INA supports *discovering* narratives within data through interactive

exploration. This fundamental difference in purpose shapes everything else: data storytelling optimizes for clarity and engagement in presenting known insights, while INA optimizes for exploration and discovery of unknown narrative structures. Similarly, story visualization techniques visualize existing, authored stories such as novels or films where the narrative structure is explicit. INA instead must *extract* implicit narratives from unstructured document collections where no explicit story structure exists, making it a fundamentally different challenge.

B. CORE PRINCIPLES AND OBJECTIVES

INA is guided by several core principles and objectives that shape its approach to narrative understanding. At its foundation, INA integrates computational algorithms with human intuition and domain expertise, addressing what Wenskovitch and North [23] call the “two black-box” problem where both artificial intelligence (AI) systems and human users may lack understanding of each other’s reasoning.

INA also strives to balance computational scalability with interactive responsiveness, enabling the processing of large text volumes while supporting real-time user interaction and feedback. This balance is essential for analyzing rapidly evolving news narratives and other high-volume text sources, as noted by Liu et al. [40] in their work on narrative extraction from breaking news.

Knowledge enhancement represents another crucial principle, as INA incorporates external knowledge resources to improve the accuracy and relevance of narrative extraction. By leveraging domain-specific ontologies, frames, and knowledge graphs, INA systems can produce more contextually appropriate and complete narrative representations, as demonstrated in recent work by Yan and Tang [35] and Blin [41].

Supporting narrative sensemaking is the primary objective of INA. Thus, this field seeks to provide tools that help users discover, explore, and understand complex narratives. These tools should enable manipulation and refinement based on user understanding and objectives, creating a more intuitive analytical environment as shown in the work of Keith et al. on semantic interaction for narrative visualization [36].

Finally, INA emphasizes evaluation and quality assessment through methods for measuring narrative quality, coherence, and utility. These evaluation approaches include metrics for assessing narrative coherence, coverage, and relevance, addressing a significant gap identified in current narrative extraction research [9]. Figure 1 presents the five core components of INA.

This comparative framework draws on established characterizations from foundational works and recent surveys across multiple domains. Visual analytics characterizations are based on seminal frameworks by Thomas and Cook [37] and Keim et al. [34], with interactive capabilities informed by Endert et al.’s semantic interaction work [14], [15] and recent

advances in large-scale visual analytics [42]. Event mining approaches reflect methodologies established by Chieu and Lee [38] and evolutionary timeline work by Yan et al. [43], while our broader narrative extraction characterization synthesizes comprehensive surveys [9], [11], [12]. Data storytelling and narrative visualization draw from Segel and Heer’s foundational taxonomy [22] and recent surveys of automation in narrative visualization [39], with visualization scalability considerations informed by information cartography approaches [44]. Intelligence analysis applications reflect the motivations and requirements articulated in Thomas and Cook’s work [37] and practical visual analytics deployments [45]. The INA characterization synthesizes our proposed framework with current capabilities in narrative extraction [9], [26], [40], interactive visualization [14], [15], [36], and knowledge integration [35], [41]. Scalability assessments for each field reflect typical system capabilities and architectural constraints discussed in domain literature, including recent work on parameter-efficient adaptation for large-scale text processing [46], [47], progressive visualization techniques [48], and multi-scale processing systems [49]. Evaluation metrics represent those commonly reported in each field’s literature, with narrative quality assessment drawing from recent work on subjective task evaluation [50], [51] and established visual analytics evaluation frameworks [34].

C. INTEGRATION OF EXISTING APPROACHES

The defining attributes of INA work in tandem to enable narrative understanding, building upon and extending established concepts from visual analytics and narrative extraction. The *human-in-the-loop* approach, while present in earlier visual analytics work [14], [15], takes on new significance in INA by incorporating user feedback directly into the narrative extraction process rather than merely the visualization layer. This allows algorithms to adapt to user understanding of narrative coherence and importance, which varies by domain and analytical goal.

Semantic interaction techniques in INA extend the foundational work of Endert et al. [14], [15] by enabling users to manipulate narrative visualizations in ways that reflect their understanding of story structure. When an analyst moves events closer together or groups related storylines, these interactions propagate to the computational model, refining its understanding of narrative relationships. This bidirectional flow between user and system, recently advanced through deep learning approaches [32], [33], creates a more intuitive environment specifically tailored for narrative analysis rather than general data exploration.

Multi-scale representation in narrative contexts presents unique challenges beyond traditional overview-plus-detail visualizations [52]. INA must maintain narrative coherence across scales—from individual events to complete storylines to entire narrative landscapes. Building on information cartography concepts [44], INA systems support fluid transitions between reading individual documents, following specific

Interactive Narrative Analytics: Conceptual Framework

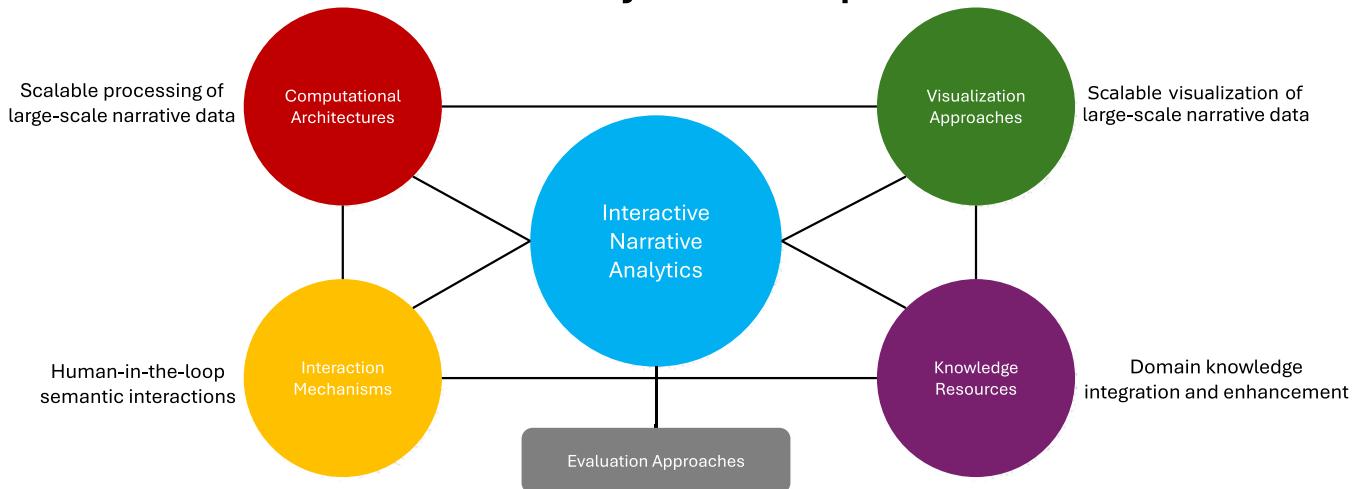


FIGURE 1. The five core components of interactive narrative analytics and their interconnections. Each component addresses specific challenges while working together in an integrated system.

TABLE 1. Comparative Framework: INA and Related Fields.

Aspect	Interactive Narrative Analytics	Visual Analytics	Event Mining	Data Storytelling	Intelligence Analysis
Primary Goal	Extract & explore narrative structures	Analyze complex data patterns	Extract discrete events	Communicate insights	Detect threats
Input Output	Document collections Interactive narrative maps	Any data type Interactive visualizations	Text corpora Event databases	Analyzed data Story presentations	Multi-source intel Threat assessments
Temporal Focus	Central (event sequences)	Variable	Event timestamps	Narrative flow	Real-time monitoring
User Role	Explorer/Analyst	Analyst	Consumer	Audience	Intelligence officer
Interactivity	Semantic interaction	Data manipulation	Query-based	Limited/guided	Alert-driven
Knowledge Integration	Domain ontologies	Statistical models	NLP models	Author knowledge	Classified databases
Target Scalability	Thousands of documents with real-time interaction	Varies by data type and visualization	Document-level batch processing	Single dataset/analysis	Real-time streaming with alerts
Evaluation Metrics	Narrative coherence, coverage, utility, user task performance	Visual scalability, interaction latency, insight quality	Precision/recall for event extraction	Engagement, comprehension, retention	Detection accuracy, false alarm rate, response time

narrative threads, and understanding the overall narrative structure. This requires careful attention to preserving temporal and causal relationships across scales, ensuring that zooming in or out does not break the narrative flow as the user progressively interacts with the system [53].

Knowledge integration in INA goes beyond simple entity linking to incorporate rich domain knowledge about narrative structures, event types, and causal relationships. Recent work on narrative-enhanced knowledge graphs [35], [41] demonstrates how external knowledge can fill gaps in extracted narratives, provide historical context, and validate causal connections. This integration must be carefully balanced—too little leaves narratives incomplete, while too much can overwhelm the extraction process with irrelevant information.

The adaptive extraction algorithms in INA learn not just from user feedback but from the evolving understanding of

what constitutes a coherent narrative in specific domains. This adaptation, advocated by Wenskovitch et al. [24], enables systems to improve their narrative extraction over time, learning domain-specific patterns of narrative development, common storyline structures, and important narrative elements that might not be explicitly stated in the text.

III. THEORETICAL FOUNDATIONS

INA draws from four key theoretical traditions that have evolved over decades of research in their respective fields. In particular, INA draws from four key theoretical traditions: interactive visual analytics [23], computational narratives [9], [11], [12], sensemaking [17], [54], [55], [56], and knowledge representation [57]. This section provides a focused overview of how these theoretical foundations converge to establish the conceptual framework of INA.

A. VISUAL ANALYTICS FOUNDATIONS

The theoretical foundations of visual analytics, essential to INA, were established through seminal works that shaped the field. Thomas and Cook's "Illuminating the Path" [37] provided the foundational definition of visual analytics as "the science of analytical reasoning facilitated by interactive visual interfaces," explicitly connecting it to intelligence analysis needs. This work, emerging from post-9/11 security concerns, established core principles of combining automated analysis with human insight that remain central to INA.

Keim et al.'s framework [34] formalized visual analytics as combining "automated analysis techniques with interactive visualizations for an effective understanding, reasoning and decision making on the basis of very large and complex data sets." This definition explicitly acknowledges the "two black box" problem that Wenskovitch and North [23] later address—neither humans nor machines fully understand each other's reasoning processes. Ben Shneiderman's "eyes have it" mantra [58]—overview first, zoom and filter, then details on demand—provides an interaction paradigm that INA extends to narrative exploration.

The intelligence analysis roots of visual analytics directly inform INA's development. Pike et al.'s work on the science of interaction [59] established principles for how analysts make sense of complex information through interactive visualization. The VAST Challenge series, beginning in 2006, has driven innovation in visual analytics for narrative understanding through scenarios involving terrorism, disease outbreaks, and social unrest. These challenges revealed the need for tools that can extract and visualize narrative structures from massive text collections—precisely the gap INA addresses.

B. NARRATIVE THEORY AND COMPUTATIONAL APPROACHES

Narrative theory provides essential concepts of narrative structure—events, entities, temporal and causal relations—that computational narratology has translated into formal models [9], [11], [12]. Event-based approaches conceptualize narratives as sequences of connected events, while entity-based approaches focus on characters and relationships. The concept of narrative coherence [25], [26] is central to computational modeling, measured through semantic similarity, temporal consistency, and causal connectedness.

C. SENSEMAKING

Sensemaking, characterized by Pirolli and Card as an iterative process involving information foraging and schema formation, is inherently subjective and context-dependent [16], [17]. Visual analytics [42] supports this process through interactive visualizations that enable exploration at multiple scales, from overview to detail [44]. Semantic interaction—enabling users to express analytical reasoning through natural interactions that propagate to underlying computational

models—is particularly significant for INA [14], [15], [32]. Figure 2 shows a mockup of a multi-scale INA system.

D. KNOWLEDGE REPRESENTATION AND INTEGRATION

Knowledge representation provides frameworks for incorporating external knowledge into narrative extraction through ontologies, frames, and knowledge graphs [35], [60]. Knowledge graphs align naturally with narrative structures [31], representing entities, events, and relationships in graph-based formats—though challenges remain in adapting general knowledge resources to specific narrative domains [41]. Figure 3 shows a representation of how knowledge could augment interactive narrative systems.

E. CONVERGENCE IN INTERACTIVE NARRATIVE ANALYTICS

INA represents the convergence of these theoretical traditions, synthesizing concepts from narrative theory, sensemaking, visual analytics, and knowledge representation. This convergence is guided by several key principles:

- **Complementarity of Computation and Cognition:** INA recognizes that computational models and human cognition offer complementary strengths in narrative understanding [23].
- **Iterative Narrative Construction:** INA conceptualizes narrative extraction as an iterative process involving continuous refinement based on user feedback and interaction [24].
- **Knowledge-Enhanced Understanding:** INA incorporates external knowledge resources to enhance narrative extraction and visualization, recognizing that narrative understanding often requires background knowledge not explicitly stated in text [35].
- **Interactive Exploration:** INA emphasizes interactive exploration in narrative understanding, providing tools that enable users to navigate and manipulate narrative representations according to their analytical goals [36].

This theoretical integration provides a foundation for addressing the challenges of narrative sensemaking. Drawing on established theories across multiple disciplines, INA offers a principled approach to combining computational methods with human insight for narrative understanding.

IV. CORE ELEMENTS: CURRENT APPROACHES AND RESEARCH CHALLENGES

The development of Interactive Narrative Analytics requires advancing five interconnected elements: computational architectures, visualization approaches, interaction mechanisms, knowledge resources, and evaluation frameworks. For each element, we examine current approaches, their limitations, and the research challenges that must be addressed to realize INA's potential.

A. THE NEED FOR INTERACTIVE NARRATIVE ANALYTICS

The limitations identified in current approaches highlight the need for an integrated field of INA that addresses these challenges holistically. By combining advances in

Multi-scale Narrative Visualization

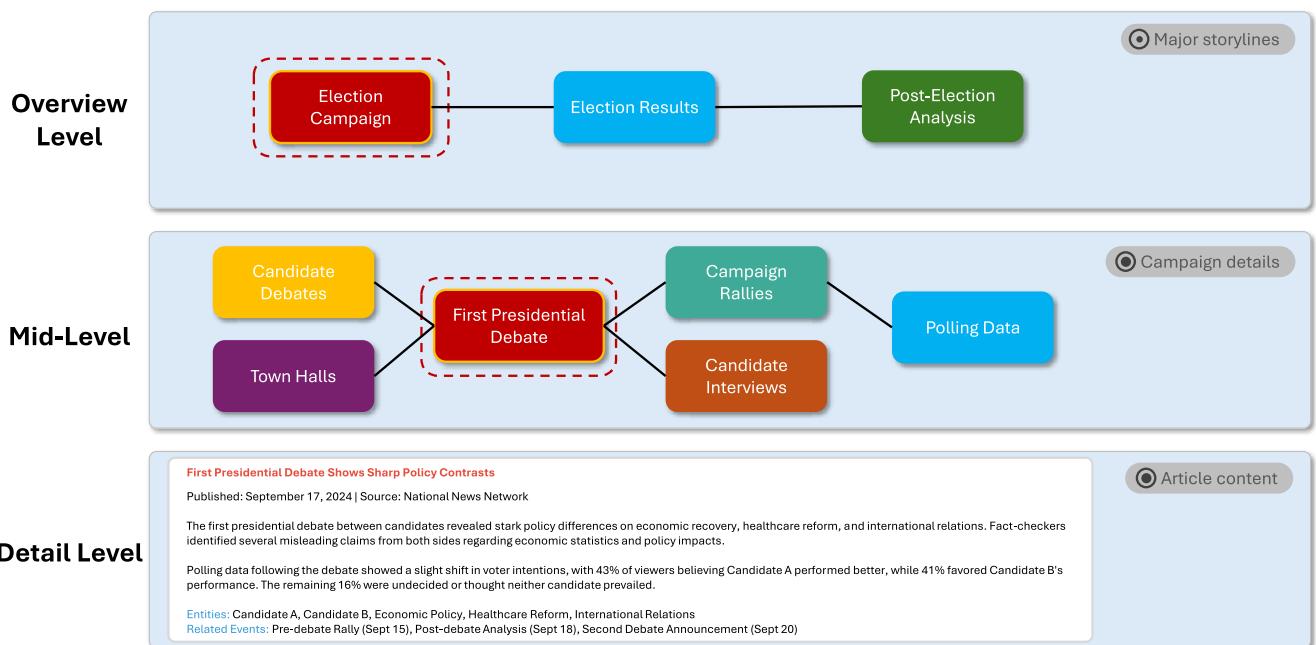


FIGURE 2. Multi-scale narrative visualization showing different levels of granularity while maintaining context across levels.

Knowledge-Enhanced Narrative Example

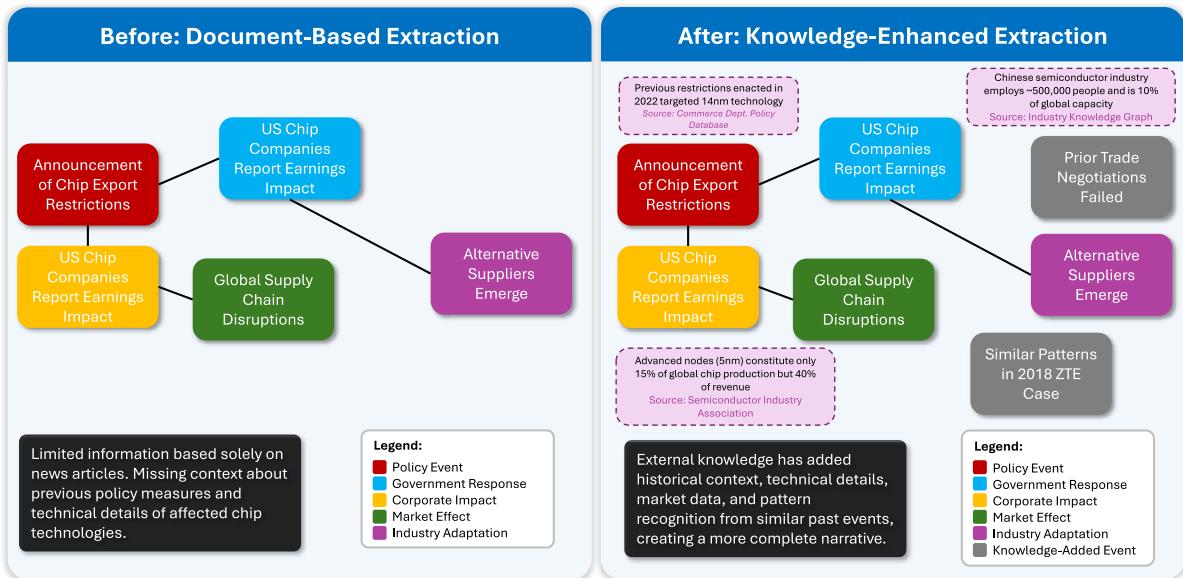


FIGURE 3. Before-and-after comparison showing how external knowledge enhances a narrative about semiconductor export restrictions with contextual information, historical patterns, and domain expertise.

computational narrative extraction, visual analytics, and interactive systems, INA aims to provide more effective tools for understanding complex narrative landscapes.

Specifically, INA seeks to address scalability challenges through optimized extraction algorithms and efficient visualization techniques, enabling analysis of large-scale

narrative collections. It aims to enhance interactivity through semantic interaction models specifically designed for narrative exploration and manipulation, allowing analysts to incorporate their expertise into the extraction process. It strives to integrate external knowledge resources to improve narrative completeness and accuracy, providing context

Traditional vs. INA Workflow Comparison

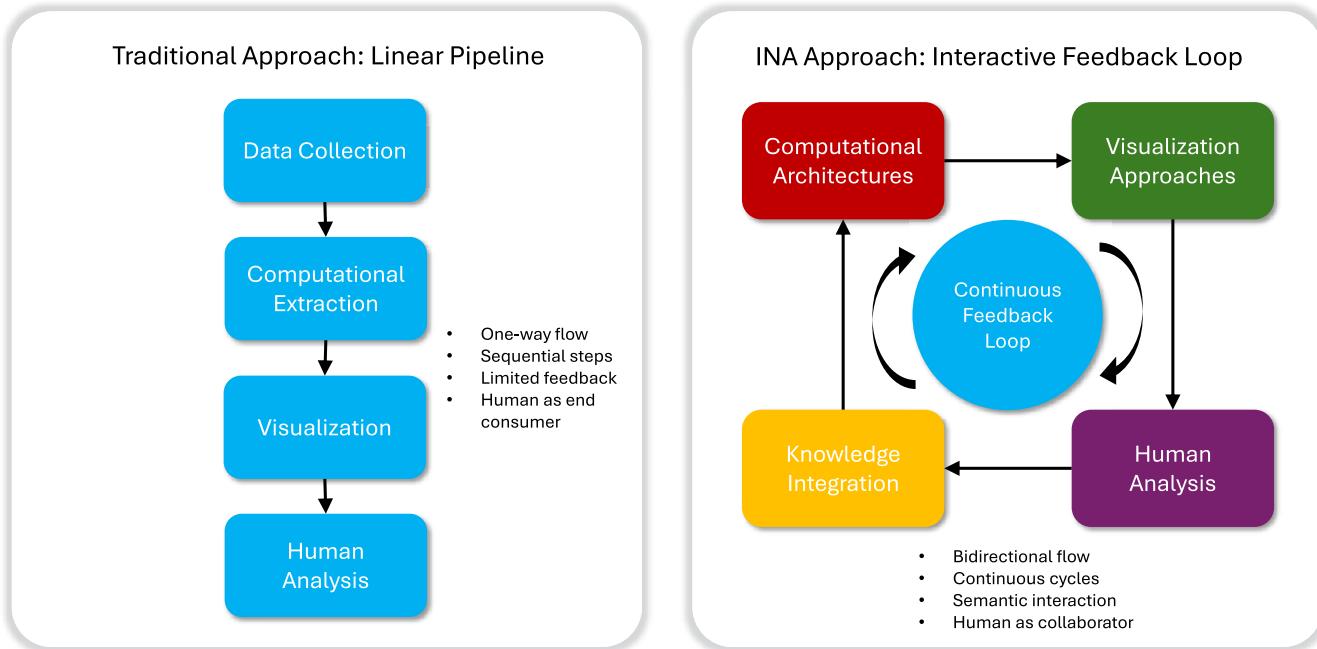


FIGURE 4. Comparison between the traditional linear pipeline approach and the Interactive Narrative Analytics integrated approach with continuous feedback between computational and human processes.

and background information that may be implicit in the text.

In addition, INA strives to develop integrated evaluation frameworks that assess both computational performance and user experience, accounting for the subjective nature of narrative interpretation while providing meaningful quality metrics. It also aims to address misinformation challenges through credibility assessment mechanisms and narrative verification techniques, helping analysts identify and understand misleading content in narrative contexts.

By addressing these challenges, INA offers a promising approach to narrative understanding in complex information environments, supporting analysts in making sense of the ever-growing volume of textual data in domains ranging from news analysis to intelligence gathering, from scientific literature to social media discourse. Figure 4 shows the comparison between INA approaches and traditional pipelines, demonstrating the paradigm shift from sequential processing to integrated, iterative analysis.

B. COMPUTATIONAL ARCHITECTURES FOR NARRATIVE PROCESSING

Current narrative extraction methods face fundamental scalability challenges that limit their practical application. Timeline-based approaches like those of Chieu and Lee [38] and Yan et al. [43] typically exhibit polynomial time complexity in the number of documents, making them impractical for large corpora. More sophisticated timeline extraction by Shahaf and Guestrin [25] identifies coherent chains

of news articles, yet still faces computational bottlenecks with large collections. Graph-based approaches such as Narrative Maps [26] and Story Forest [40] face even steeper computational challenges, with complexity often quadratic or higher relative to document count. These methods involve computationally intensive operations including pairwise document similarity calculations, global optimization over entire document collections, and complex graph algorithms that become prohibitive as collections grow [9].

The computational demands become particularly acute when narratives evolve in real-time, requiring systems to process streaming data and incrementally update narrative structures without complete recalculation. Current approaches struggle with this temporal constraint, creating tension between computational thoroughness and responsiveness. Shahaf et al. [44] address this through approximation techniques and incremental updates, but these solutions often introduce trade-offs between narrative quality and computational performance that remain unresolved. Most existing systems operate in batch mode, processing entire document collections offline before presenting results, which fails to support the iterative exploration central to INA.

The integration of Large Language Models presents both opportunities and challenges for computational narrative extraction. While LLMs offer powerful semantic processing capabilities that could enhance narrative understanding, they introduce substantial computational costs, potential biases, and transparency issues [29], [30], [61]. Their black-box nature often obscures the reasoning behind

extracted narrative structures, making it difficult for analysts to understand how particular narratives were derived. Current research explores using LLMs for specific subtasks like event extraction or coherence scoring, but fully integrating these models into interactive narrative systems while maintaining computational efficiency and interpretability represents a critical research challenge.

Addressing these computational challenges requires developing architectures that balance multiple competing demands. Systems must process document collections ranging from thousands to millions of documents while maintaining sub-second response times for interactive operations. They must support both batch processing for initial narrative extraction and stream processing for real-time updates. They must integrate multiple computational approaches—from traditional NLP to deep learning to graph algorithms—within a unified framework. Recent work on distributed computing frameworks and GPU acceleration shows promise [42], but adapting these technologies specifically for narrative processing while maintaining the semantic richness required for meaningful narrative extraction remains an open challenge.

Recent advances in parameter-efficient fine-tuning provide practical pathways for adapting large language models to narrative extraction tasks while maintaining computational feasibility for interactive systems. For example, Low-Rank Adaptation (LoRA) [46] decomposes weight updates into low-rank matrices, reducing trainable parameters by orders of magnitude while matching full fine-tuning performance. For resource-constrained environments, QLoRA [47] combines 4-bit quantization with LoRA to enable fine-tuning of very large models on single GPUs. Prompt tuning approaches [62], [63] offer even greater efficiency by optimizing only continuous task-specific vectors while freezing model parameters entirely. For INA systems, modular integration architectures could allow the selective use of LLMs for specific subtasks (e.g., event extraction, coherence scoring, and entity linking) while maintaining traditional approaches for computationally intensive operations. However, the choice among these techniques involves trade-offs between model size, computational efficiency, adaptation quality, and interactive responsiveness.

Another challenge in this context is addressing the “two black-box” problem, which requires integrating explainable AI techniques specifically adapted for narrative extraction. SHAP-based approaches [64] or LIME variants [65] could be used to quantify the contribution of individual narrative elements (events, entities, temporal markers) to extraction decisions, though adaptation from tabular to textual narrative data requires careful consideration [66]. Counterfactual generation methods [67], [68] could enable analysts to understand how narrative structures would change under alternative event configurations—for instance, how removing specific events would affect inferred causal chains. Attention visualization techniques that go beyond raw attention weights [69] could be useful to reveal which textual spans most influence narrative connection decisions across transformer layers. For

narrative-specific explanations, causal inference approaches [70], [71] could help surface why particular events are grouped into storylines or why certain temporal/causal relationships are inferred. In this context, developing *interactive explanation interfaces* could allow users to query why specific narrative elements were connected, drill down into model reasoning, and provide corrective feedback that propagates through the extraction model—creating a transparent and collaborative human-in-the-loop system where analysts understand and can guide computational narrative construction.

C. VISUALIZATION APPROACHES FOR NARRATIVE STRUCTURES

Creating effective visual representations of narrative structures presents unique challenges beyond traditional information visualization. Narratives inherently involve complex structures with multiple storylines, branching paths, and interconnected elements that resist straightforward visual encoding. Current approaches broadly fall into timeline-based, graph-based, and hybrid visualizations, each with distinct capabilities and limitations.

Timeline-based visualizations excel at representing temporal evolution but struggle with complex non-linear narratives. Systems like CloudLines [72] address scalability through compact visualization techniques using logarithmic time scaling and distortion lenses, yet they primarily support linear narrative flows. Shahaf et al.’s Information Cartography [44] extends timeline concepts with zoomable maps supporting multi-scale exploration, but maintaining narrative coherence across scales remains challenging. The fundamental limitation of timeline approaches lies in their difficulty representing parallel storylines, narrative branches, and convergent paths that characterize complex real-world narratives [9].

Graph-based visualizations offer greater expressiveness for complex narrative structures but introduce severe scalability and interpretability challenges. As narrative graphs grow, they quickly become cluttered and difficult to navigate [42]. Keith and Mitra’s Narrative Maps [26] attempt to address this through hierarchical clustering and progressive disclosure, but users still struggle to maintain awareness of overall narrative structure while examining specific details. Liu et al. [40] organize events into tree structures that capture evolutionary relationships, yet these still face visualization challenges as trees grow. The cognitive load of interpreting complex graph structures can overwhelm analysts, particularly when narratives involve hundreds of events and multiple intersecting storylines. Furthermore, standard graph layout algorithms often fail to preserve the temporal and causal relationships crucial to narrative understanding.

The design of these visualizations must consider human perceptual and cognitive constraints [73], [74]. Limited working memory, attentional bottlenecks, and interpretation biases all affect how effectively analysts can make sense of visual narrative representations [75], [76], [77]. Visualization approaches must work within these constraints while still

Human-AI Narrative Sensemaking Interface

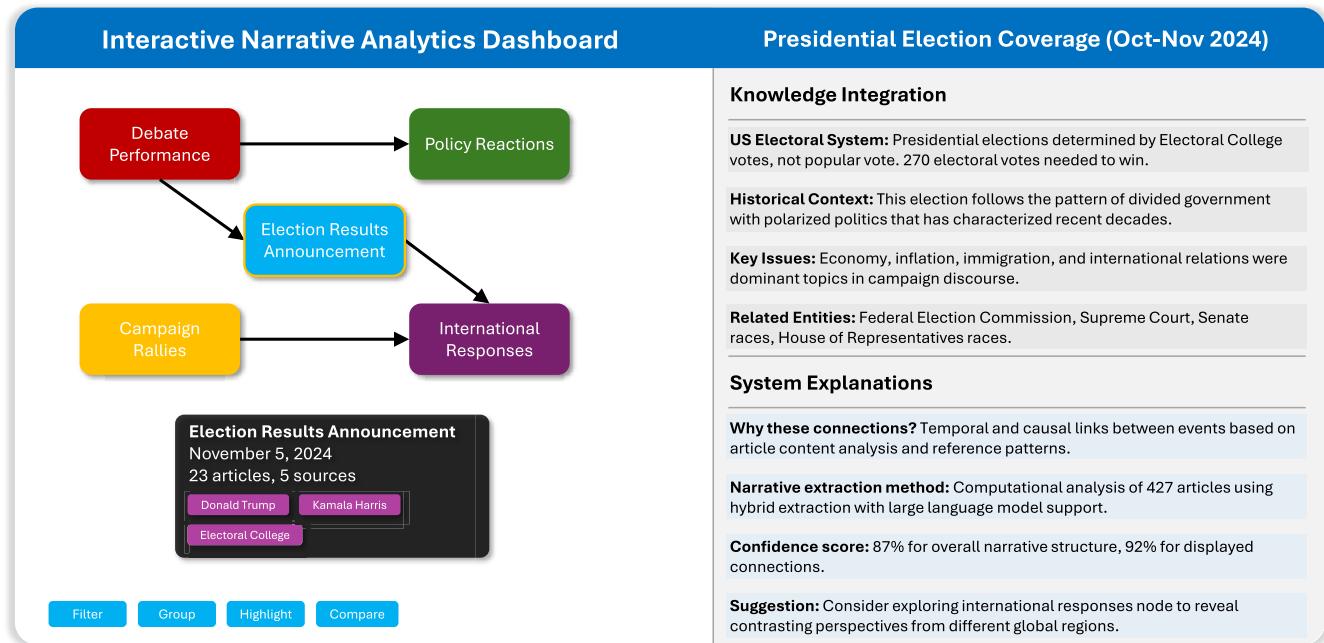


FIGURE 5. Mockup of an Interactive Narrative Analytics interface showing a narrative map visualization with semantic interaction capabilities and knowledge integration features.

conveying the richness and complexity of narrative structures. Figure 5 shows a mockup of a potential INA system that addresses these challenges through multi-level representation, semantic interaction capabilities, and integrated knowledge features.

Hybrid approaches attempt to combine the strengths of timeline and graph representations. Angus et al.'s Discourse Lines [27] blend timeline representations with narrative maps to show policy and media storylines, using spatial positioning to encode both temporal progression and thematic relationships. These hybrid visualizations show promise but introduce new challenges in maintaining visual coherence, managing screen real estate, and helping users understand multiple visual encodings simultaneously. The design space for hybrid narrative visualizations remains largely unexplored, with no established guidelines for when to use which visual encoding or how to transition between them [36].

Topic and trend visualization techniques, while not specifically designed for narratives, provide relevant insights for INA design. Topic models visualize thematic patterns in text collections, though they typically operate at document level rather than capturing fine-grained event structures [42]. Trend visualization shows temporal patterns but often lacks the causal and explanatory connections between events that define narratives [11]. These approaches highlight the gap between general text visualization and narrative-specific visual representations.

In this context, several multi-scale techniques could address some of the visualization challenges inherent in

complex narrative structures. Edge bundling methods [78], [79] reduce visual clutter in narrative graphs by routing related connections along shared paths while maintaining readability, with recent optimizations achieving significant speedup for large-scale graphs. Timeline compression approaches enable semantic zoom where temporal granularity adjusts based on user focus—for instance, showing individual events during detailed examination but aggregating to weekly or monthly summaries in overview mode [49]. Hierarchical clustering techniques organize narratives into multi-level structures allowing top-down exploration from major storylines to constituent events [80], with progressive visualization methods [48] enabling incremental refinement as users navigate across scales. The key challenge in using these approaches lies in maintaining narrative coherence across scale transitions—ensuring that temporal ordering, causal relationships, and thematic connections remain interpretable whether viewing individual documents, specific narrative threads, or the entire narrative landscape.

D. INTERACTION MECHANISMS FOR NARRATIVE EXPLORATION

The interactive dimension fundamentally distinguishes INA from automated narrative extraction, yet current systems provide limited support for meaningful interaction with narrative structures. Most existing narrative extraction methods operate in batch mode with minimal user involvement [36], while those that do support interaction typically limit it to basic operations like filtering, zooming, and highlighting. Search

and filtering interfaces rely heavily on users formulating appropriate queries and synthesizing results into coherent narratives themselves [42]. This gap between automated extraction and interactive exploration represents a critical limitation in current approaches.

Semantic interaction, pioneered by Endert et al. [14], [15] and extended through deep learning by Bian and North [33], enables analysts to express analytical reasoning through natural interactions with visualizations. When users spatially arrange documents or adjust similarities, these interactions propagate to underlying computational models. Recent advances by Han et al. [32] demonstrate how explainable projections can make these interactions more interpretable. However, current semantic interaction systems focus primarily on document-level representations rather than narrative structures [9]. Extending semantic interaction to narrative elements—events, entities, relationships, and storylines—requires new interaction paradigms that can capture the complex reasoning involved in narrative sensemaking.

Keith et al. [36] introduce mixed multi-model semantic interaction specifically for narrative visualizations, allowing users to manipulate graph-based narrative representations with changes propagating through multiple computational models. Their approach demonstrates how user interactions with visual narrative elements can update both the extraction model’s understanding of narrative coherence and the relevance weights for different narrative components. Yet this work also reveals fundamental challenges: inferring user intent from potentially ambiguous interactions, translating those intentions into meaningful model updates, and managing conflicts between user preferences and computational recommendations all remain difficult problems.

The challenge of interaction design for narrative analytics extends beyond technical implementation to understanding how analysts naturally think about and work with narratives. Cognitive processes involved in narrative sensemaking [81], [82] include identifying key events, establishing causal connections, recognizing patterns across storylines, and synthesizing multiple perspectives into coherent understanding. Current interaction mechanisms often fail to support these natural reasoning processes, forcing analysts to translate their narrative understanding into system-specific operations. Developing interaction vocabularies that align with analysts’ mental models of narratives while remaining computationally tractable represents an ongoing challenge [14], [15].

System responsiveness creates additional technical challenges when interactions trigger computationally intensive operations. Re-extracting narratives based on user feedback, updating visualizations to reflect new parameters, or computing alternative narrative structures all potentially require significant computation. Current systems often sacrifice either interactivity (by processing updates offline) or quality (by using simplified models for real-time response). Progressive computation approaches that provide initial results quickly while refining them over time show promise [83], but implementing these approaches specifically for narrative

analytics while maintaining narrative coherence requires careful algorithm design and avoiding the issue of introducing potentially spurious patterns [84].

Multimodal interaction modalities could enhance narrative exploration through natural, accessible interfaces. For example, voice command interfaces could enable hands-free navigation of narrative structures, particularly valuable for large display environments where analysts may be distant from traditional input devices [85]. Natural language query systems allow analysts to ask questions like “show me events related to trade negotiations in March” rather than constructing complex filter sequences [86]. Gesture control for spatial narrative manipulation enables direct manipulation of graph-based narrative visualizations through touch or mid-air gestures, creating more intuitive workflows for exploring connections between events. Eye-tracking systems could implicitly detect analyst interest, automatically providing additional detail for fixated elements or adapting visualizations based on attention patterns. However, such multimodal interaction capabilities remain unexplored in INA and could present interesting opportunities for future work in the field.

E. KNOWLEDGE RESOURCES FOR NARRATIVE ENHANCEMENT

External knowledge integration fundamentally enhances narrative extraction by providing context, background information, and domain expertise that may not be explicit in text. However, incorporating knowledge resources into narrative analytics systems presents significant challenges in representation, alignment, and reasoning that current approaches only partially address.

Knowledge representation for narratives must balance expressive power with computational feasibility. Ontologies and frames can capture rich narrative concepts—event types, participant roles, causal relationships—but often become unwieldy for large-scale processing. Knowledge graphs offer more scalable representations and align naturally with graph-based narrative structures [35], [41], yet they struggle to represent the temporal dynamics and uncertainty inherent in evolving narratives. Porzel et al. [57] explore narrativizing knowledge graphs to better capture story-like structures, but finding representations that can capture narrative complexity while supporting efficient computation remains challenging.

The alignment between textual mentions and knowledge resources introduces fundamental difficulties in narrative contexts. Entity linking in narratives must handle not just ambiguity and coreference but also temporal evolution—entities change over time, relationships shift, and the same entity may play different narrative roles in different contexts. Event detection and alignment face even greater challenges, as the same real-world event may be described differently across sources, at different granularities, and from different perspectives. Current approaches using neural methods show promise for entity and event alignment, but they often fail to capture the narrative-specific aspects of these problems,

such as maintaining character continuity or preserving causal chains [35].

Domain adaptation presents another critical challenge for knowledge-enhanced narrative analytics. Different narrative domains—news, scientific literature, legal documents—require different knowledge resources and reasoning patterns. A system analyzing news narratives needs knowledge about current events, political entities, and geographic relationships, while scientific narrative analysis requires understanding of research concepts, methodologies, and citation patterns. Developing domain-specific knowledge resources requires substantial expertise and effort, yet general-purpose resources often lack the specificity needed for meaningful narrative enhancement [41]. Transfer learning and few-shot adaptation techniques offer potential solutions, but their application to narrative-specific knowledge remains largely unexplored.

Handling noisy and conflicting knowledge graph data requires robust mechanisms at multiple levels [87], [88]. For instance, confidence scoring systems could weight knowledge sources based on provenance, recency, and historical reliability, allowing narrative extraction to prioritize high-quality information while still considering alternative perspectives. Provenance tracking [89] mechanisms must maintain detailed records of where knowledge originated, enabling analysts to trace claims back to sources and assess credibility in context. In this context, conflict resolution heuristics become necessary when multiple sources provide contradictory information about events or relationships: voting-based approaches can favor majority consensus, recency-based methods can prioritize newer information for evolving situations, and source reliability-based techniques can defer to authoritative sources.

The integration of multiple knowledge sources compounds these challenges. Real-world narrative analysis often requires combining general knowledge (from resources like WikiData or ConceptNet) with domain-specific knowledge (from specialized ontologies) and temporal knowledge (from event databases or news archives). These resources may use different representations, cover different time periods, and contain conflicting information. Reconciling these differences while maintaining narrative coherence requires sophisticated fusion techniques that can reason about source reliability, temporal validity, and contextual relevance [57].

F. EVALUATION APPROACHES FOR NARRATIVE QUALITY

The assessment of INA systems presents distinctive challenges related to the complex, subjective nature of narrative understanding. Developing metrics and methodologies that meaningfully capture narrative quality, system usability, and analytical utility remains an open research problem [9]. The lack of standardized evaluation frameworks significantly hinders comparative assessment and systematic progress in the field.

Narrative quality assessment involves multiple dimensions including coherence, completeness, accuracy, and relevance.

Formalizing these subjective qualities into measurable properties requires careful consideration of narrative theory, computational feasibility, and human judgment alignment. Different analytical contexts may prioritize different quality aspects, further complicating standardization efforts. Current metrics often capture isolated aspects of narrative quality but fail to provide comprehensive assessment aligned with human narrative perception.

User-centered evaluation introduces additional complexity, as it must assess how effectively systems support analytical tasks within realistic contexts. Methodologies must balance experimental control with ecological validity, requiring careful task design and appropriate measurement approaches. The diverse analytical needs across different application domains further complicate standardization efforts. Developing formal evaluation approaches that are also practical represents a significant challenge for advancing INA research. Future empirical studies comparing INA approaches should carefully consider experimental design choices—such as between-subjects versus within-subjects comparisons—and ensure sufficient statistical power to detect meaningful differences in user performance and narrative quality.

Comparative evaluation across different INA systems requires standardized datasets, tasks, and metrics that remain largely undeveloped. Creating representative benchmark materials that capture diverse narrative phenomena while supporting systematic comparison requires substantial community effort. The subjective nature of narrative interpretation further complicates ground truth establishment, as multiple valid interpretations may exist for the same narrative material. Addressing these evaluation challenges will require collaborative efforts across the research community and integration of perspectives from multiple disciplines.

In particular, addressing subjectivity in narrative interpretation requires evaluation frameworks that embrace rather than eliminate disagreement [51]. Thus, annotation protocols that preserve individual annotator judgments—rather than simply collapsing to a majority vote—are required to retain valuable disagreement signals, as perfect agreement may be impossible—and even undesirable—for subjective narrative tasks [50]. Multi-annotator modeling approaches [90] treat disagreement as informative signal, predicting distributions over possible interpretations rather than single labels, with evaluation metrics that reward capturing this diversity. Task-based evaluation that measures utility for downstream analytical goals (e.g., successful investigation completion, accurate threat assessment) can complement agreement-based metrics, assessing whether systems enable effective sensemaking regardless of “ground truth” [91]. Importantly, evaluation must recognize that disagreement sometimes reflects systematic differences across demographic or ideological groups [92] rather than random noise, requiring fairness metrics that ensure minority perspectives receive appropriate representation in both extraction and evaluation.

Ethical considerations in evaluation add another layer of complexity. Evaluation must consider not just technical performance but also potential biases, privacy implications, and societal impacts [30], [93]. Systems that excel at extracting narratives might also be effective at surveillance or spreading misinformation [3], [4], [94]. Evaluation frameworks must therefore incorporate ethical dimensions, assessing whether systems respect privacy, avoid amplifying biases, and prevent misuse. The challenge of detecting and mitigating misinformation within narrative structures requires specific evaluation criteria [95], [96], as systems must help analysts identify potentially misleading content while understanding how false narratives spread and evolve.

G. ETHICAL CONSIDERATIONS IN NARRATIVE ANALYTICS

The development and deployment of INA systems raise critical ethical considerations that must be addressed as the field matures. Fischer et al.'s work on ethical awareness in communication analysis [97] provides a framework particularly relevant to INA applications. Their MULTI-CASE framework [98] demonstrates how transformer-based systems can incorporate ethical considerations directly into the analytical process, implementing equal joint agency between human and AI for heterogeneous information analysis.

Privacy concerns in narrative extraction are paramount, as personal narratives extracted from social media, news, or other sources may contain highly sensitive information about individuals and communities. The aggregation of narrative fragments across multiple sources can reveal patterns and connections that individuals never intended to make public. INA systems must implement appropriate consent mechanisms, data minimization principles, and purpose limitation as required by regulations like GDPR. The challenge intensifies when narratives span multiple jurisdictions with different privacy regulations and cultural expectations about information sharing.

Algorithmic bias in narrative interpretation presents another critical challenge. Narrative understanding is inherently cultural, and systems trained on Western news sources may misinterpret narratives from other cultural contexts. The selection of which events to include, how to connect them, and what constitutes narrative coherence all embed cultural assumptions that may not transfer across contexts. Addressing this requires diverse training datasets, regular auditing of system outputs for bias, and inclusive design processes that involve stakeholders from multiple cultural backgrounds.

The potential for INA systems to be used for surveillance or population control cannot be ignored. The same capabilities that help journalists understand complex stories or researchers track scientific developments could be used to monitor dissent, track activist movements, or suppress minority narratives. Developers must consider implementing technical safeguards, use restrictions, and transparency

measures that prevent misuse while enabling legitimate applications. This includes careful consideration of who has access to these tools and for what purposes.

Thus, considering the previous concerns, operationalizing ethical INA systems requires concrete technical safeguards and governance mechanisms. Bias auditing pipelines should systematically evaluate narrative extraction across demographic groups, geographic regions, and ideological perspectives [99], testing whether certain communities' narratives are systematically underrepresented or mischaracterized. Fairness-aware modeling techniques must ensure balanced narrative extraction across diverse sources, avoiding over-reliance on dominant perspectives while properly weighting minority viewpoints. Privacy-preserving computation becomes critical when narratives involve sensitive information about individuals or communities; differential privacy approaches [100] and federated learning can enable aggregate narrative analysis while protecting individual data points. Transparency measures should include model cards [101] documenting intended use cases, known limitations, and performance across demographic groups, alongside datasheets [102] specifying data collection processes and potential biases. Governance frameworks must establish oversight mechanisms for INA deployment, particularly in high-stakes contexts like intelligence analysis or misinformation detection, ensuring that systems remain accountable and contestable rather than becoming opaque arbiters of narrative truth.

V. FUTURE RESEARCH DIRECTIONS

A. ADVANCED COMPUTATIONAL MODELS FOR NARRATIVE UNDERSTANDING

Future research should focus on developing more powerful and efficient computational models for narrative extraction and understanding. Narrative models based on LLMs offer significant potential but require adaptation for narrative-specific tasks [29], [30]. Research priorities include developing specialized architectures or training objectives that capture narrative structures, creating more efficient models suitable for interactive applications, and enhancing model interpretability to support human-AI collaboration [103], [104]. Incremental and adaptive extraction algorithms represent another critical research direction, enabling systems to efficiently update narrative structures as new information becomes available rather than reprocessing entire document collections [40]. These approaches must address challenges in maintaining narrative coherence across updates while ensuring computational efficiency suitable for interactive exploration.

B. HUMAN-AI COLLABORATION FOR NARRATIVE SENSEMAKING

Effective integration of human intelligence and computational capabilities represents a fundamental research direction for INA [23], [24]. This research should explore different

models for human-AI collaboration in narrative understanding, examining how to allocate responsibilities between analysts and algorithms to leverage their complementary strengths. Key research questions include developing appropriate division of cognitive labor, designing interaction mechanisms that support smooth transitions of control, and creating collaborative frameworks that enhance rather than diminish human analytical capabilities. These questions require interdisciplinary approaches that combine insights from human-computer interaction, cognitive science, and AI.

C. KNOWLEDGE-ENHANCED NARRATIVE ANALYTICS

The integration of external knowledge into narrative extraction presents significant opportunities for enhancing the completeness, accuracy, and context of extracted narratives [35], [41], [57]. Future research should develop approaches for effectively incorporating domain-specific knowledge and commonsense reasoning into narrative analytics systems. Research priorities include developing domain-specific knowledge resources for narrative domains, creating efficient alignment techniques for mapping textual mentions to knowledge bases, and designing inference mechanisms that apply external knowledge to enhance narrative understanding. This integration must balance knowledge depth with computational efficiency to maintain interactive performance.

D. EVALUATION FRAMEWORKS AND ETHICAL CONSIDERATIONS

The development of standardized evaluation frameworks represents a critical requirement for advancing INA research. These frameworks should include narrative quality metrics that formalize subjective aspects of narrative coherence and utility, benchmark datasets that enable systematic comparison between approaches, and user-centered evaluation methodologies that assess how effectively systems support analytical tasks [26]. In addition to technical development, research must address ethical considerations such as detection and mitigation of bias, protection of privacy in narrative analysis, and transparency improvements that make extraction processes more understandable for users [30], [93].

E. ANTICIPATED LIMITATIONS FOR INA SYSTEMS

In summary, while INA offers promising approaches to narrative understanding, several fundamental limitations must be acknowledged. Computational complexity presents a persistent challenge: extracting coherent narratives from massive document collections while maintaining interactive responsiveness requires careful algorithmic design and potentially costly infrastructure. The subjective nature of narrative interpretation means that “correct” narrative extraction may not exist—different analysts with different goals may legitimately construct different narrative representations from the same documents. Domain transfer remains difficult, as narrative extraction systems trained on news data may perform poorly on scientific literature or legal documents without

substantial adaptation. Evaluation standardization proves challenging given the lack of objective ground truth for many narrative sensemaking tasks, requiring development of task-based metrics that measure analytical utility rather than algorithmic accuracy. Finally, misinformation and adversarial narratives pose unique challenges, as malicious actors may deliberately construct misleading narrative structures that exploit system assumptions about coherence and plausibility.

VI. APPLICATIONS AND OPPORTUNITIES

A. NEWS ANALYSIS AND MISINFORMATION DETECTION

News analysis represents a primary application domain for INA, addressing challenges in tracking evolving stories and identifying misinformation. INA systems can extract and visualize narrative structures from diverse sources, revealing how stories develop over time and how different publications characterize events [40]. These capabilities support comparative analyses, identification of underreported aspects, and contextual verification of claims. In misinformation detection, INA can provide contextual verification mechanisms that place specific claims within broader narrative structures [95], [96]. In particular, social media platforms generate vast amounts of content with complex diffusion patterns and embedded misinformation. INA can help track how narratives spread and evolve across platforms and communities, revealing adaptation patterns and transformations [105], [106].

In particular, INA systems could integrate multiple complementary components into a unified workflow. Recent work demonstrates the importance of narrative structure analysis for identifying coordinated disinformation campaigns [107], [108]. An INA-based misinformation detection system would combine: (1) stance detection using retrieval-augmented approaches to identify how claims propagate across narratives [109], (2) fact-checking knowledge base integration drawing from datasets like AVeriTeC [110] and FEVER [111] for evidence-based verification, (3) source credibility assessment [112] tracking provenance and historical reliability patterns within narrative contexts, (4) inconsistency detection [113] identifying contradictions in evolving narratives over time, and (5) narrative-based claim verification that examines not just isolated statements but their embedding within broader story structures [108]. This integrated approach addresses the limitation that misinformation often manifests through narrative framing and psychological mechanisms [114] rather than simple factual errors, requiring systems that can surface how stories evolve and how different sources characterize events.

B. INTELLIGENCE ANALYSIS AND SECURITY APPLICATIONS

Intelligence analysis involves making sense of complex, incomplete, and sometimes contradictory information. INA can assist by extracting narrative elements from intelligence reports and visualizing connections between them [45]. These

capabilities help analysts discover non-obvious relationships between events, identify gaps in intelligence coverage, and construct a more thorough understanding of complex situations. INA systems can also support the generation and testing of hypotheses by enabling the exploration of alternative narrative explanations, comparing these hypotheses with available evidence [16], [17].

In particular, building on the visual analytics frameworks that emerged from post-9/11 intelligence needs [37], an INA-based intelligence system would integrate: (1) multi-source narrative fusion combining structured intelligence reports, intercepted communications, open-source news, and social media into unified narrative representations that reveal connections across disparate information streams, (2) temporal pattern detection identifying recurring narrative structures that may signal coordinated activities or evolving threats [49], (3) anomaly identification highlighting events or relationships that deviate from expected narrative patterns within specific threat contexts, (4) hypothesis generation and testing enabling analysts to construct alternative narrative explanations and evaluate them against available evidence [16], [17], and (5) multi-scale exploration allowing seamless navigation from strategic overviews of entire situations to tactical details of specific events or actors [48]. Knowledge integration with threat databases, historical case studies, and domain expertise could enhance narrative extraction while privacy-preserving techniques [100] could enable collaborative analysis across organizational boundaries without compromising sensitive sources.

C. SCIENTIFIC LITERATURE EXPLORATION

The exponential growth of scientific literature creates challenges for researchers attempting to understand research landscapes and identify relevant prior work. INA can help researchers map these landscapes by extracting and visualizing narrative structures from scientific publications [44]. INA systems can also support the identification of literature gaps by revealing underexplored connections or questions within research narratives, and scientific claim verification by tracking how findings evolve through subsequent studies [31].

In particular, a scientific literature INA system would combine: (1) research trajectory mapping that traces how specific research questions, methodologies, or findings develop across publications and citation networks, revealing the narrative arc of scientific inquiry [44], (2) literature gap identification by analyzing narrative structures to reveal under-explored connections, contradictory findings that lack resolution, or logical next steps in research progressions, (3) claim verification across studies tracking how specific scientific claims are supported, refuted, or refined through subsequent research [110], enabling researchers to understand evidential strength and ongoing controversies, (4) cross-disciplinary connection discovery identifying narrative threads that span traditional disciplinary boundaries, revealing potential synergies or common patterns across fields, and (5) methodological evolution tracking showing

how experimental approaches, analytical techniques, or theoretical frameworks develop and spread through research communities, while acknowledging that scientific narrative interpretation may vary across theoretical traditions or methodological schools [50].

VII. CONCLUSION

We have defined **Interactive Narrative Analytics** as a nascent field addressing the challenges of understanding complex narrative structures in our information-rich landscape. Rather than treating narrative extraction and interactive visualization separately, we establish INA as a distinct interdisciplinary field that integrates these approaches to solve challenges neither computational nor manual methods can address alone. Furthermore, current narrative analysis suffers from fundamental limitations: computational methods operate as black boxes, visualization systems lack narrative-specific capabilities, and both fail to incorporate external knowledge effectively. These limitations create a significant gap between theoretical potential and practical utility.

INA represents a paradigm shift by placing human-machine collaboration at the center of narrative sensemaking, challenging the notion that extraction should merely precede human analysis. Our five core components—scalable computational architectures, interactive visualization techniques, semantic interaction models, knowledge integration approaches, and evaluation metrics—define the scope of this field.

Finally, as information volumes grow and narrative manipulation becomes more sophisticated, INA offers a fundamental reconceptualization of how research communities approach complex narrative landscapes. We invite researchers across disciplines to contribute to developing this field's foundations, methodologies, and applications.

REFERENCES

- [1] A. Gandomi and M. Haider, "Beyond the hype: Big data concepts, methods, and analytics," *Int. J. Inf. Manage.*, vol. 35, no. 2, pp. 137–144, Apr. 2015.
- [2] P. G. Roetzel, "Information overload in the information age: A review of the literature from business administration, business psychology, and related disciplines with a bibliometric approach and framework development," *Bus. Res.*, vol. 12, no. 2, pp. 479–522, Dec. 2019.
- [3] D. Lazer, M. Baum, Y. R. Benkler, A. J. Berinsky, K. M. Greenhill, F. Menczer, M. J. Metzger, B. Nyhan, G. Pennycook, D. Rothschild, M. Schudson, S. A. Sloman, C. R. Sunstein, E. Thorson, D. J. Watts, and J. Zittrain, "The science of fake news," *Science*, vol. 359, no. 6380, pp. 1094–1096, 2018.
- [4] S. Vosoughi, D. Roy, and S. Aral, "The spread of true and false news online," *Science*, vol. 359, no. 6380, pp. 1146–1151, Mar. 2018.
- [5] X. Zhou and R. Zafarani, "A survey of fake news: Fundamental theories, detection methods, and opportunities," *ACM Comput. Surveys*, vol. 53, no. 5, pp. 1–40, Sep. 2021.
- [6] O. D. Apuke, B. Omar, E. A. Tunca, and C. V. Gever, "Information overload and misinformation sharing behaviour of social media users: Testing the moderating role of cognitive ability," *J. Inf. Sci.*, vol. 50, no. 6, pp. 1371–1381, Dec. 2024.
- [7] M. J. Eppler and J. Mengis, "The concept of information overload: A review of literature from organization science, accounting, marketing, MIS, and related disciplines," *Inf. Soc.*, vol. 20, no. 5, pp. 325–344, 2004.

[8] S. L. Cohen, J. T. Hamilton, and F. Turner, "Computational journalism," *Commun. ACM*, vol. 54, no. 10, pp. 66–71, 2011.

[9] B. Keith, T. Mitra, and C. North, "A survey on event-based news narrative extraction," *ACM Comput. Surv.*, vol. 55, no. 14s, pp. 1–39, Jul. 2023, doi: [10.1145/3584741](https://doi.org/10.1145/3584741).

[10] K. Dalkir and R. Katz, *Navigating Fake News, Alternative Facts, and Misinformation in a Post-Truth World*. Hershey, PA, USA: IGI Global, 2020.

[11] P. Ranade, S. Dey, A. Joshi, and T. Finin, "Computational understanding of narratives: A survey," *IEEE Access*, vol. 10, pp. 101575–101594, 2022.

[12] B. Santana, R. Campos, E. Amorim, A. Jorge, P. Silvano, and S. Nunes, "A survey on narrative extraction from textual data," *Artif. Intell. Rev.*, vol. 56, no. 8, pp. 8393–8435, Aug. 2023.

[13] C. Andrews, A. Endert, and C. North, "Space to think: Large high-resolution displays for sensemaking," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, New York, NY, USA, Apr. 2010, pp. 55–64.

[14] A. Endert, C. Han, D. Maiti, L. House, S. Leman, and C. North, "Observation-level interaction with statistical models for visual analytics," in *Proc. IEEE Conf. Vis. Analytics Sci. Technol. (VAST)*, Providence, RI, USA, Oct. 2011, pp. 121–130.

[15] A. Endert, P. Fiaux, and C. North, "Semantic interaction for visual text analytics," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, May 2012, pp. 473–482.

[16] P. Pirolli and S. Card, "Information foraging in information access environments," in *Proc. SIGCHI Conf. Human factors Comput. Syst.*, New York, NY, USA, 1995, pp. 51–58.

[17] P. Pirolli and S. Card, "The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis," in *Proc. Int. Conf. Intell. Anal.*, vol. 5, McLean, VA, USA, 2005, pp. 2–4.

[18] S. C. Lewis, R. Zamith, and A. Hermida, "Content analysis in an era of big data: A hybrid approach to computational and manual methods," *J. Broadcast. Electron. Media*, vol. 57, no. 1, pp. 34–52, Jan. 2013.

[19] C. S. Park, "Does too much news on social media discourage news seeking? Mediating role of news efficacy between perceived news overload and news avoidance on social media," *Social Media Soc.*, vol. 5, no. 3, Apr. 2019, Art. no. 2056305119872956.

[20] S. H. Soroya, A. Farooq, K. Mahmood, J. Isoaho, and S.-E. Zara, "From information seeking to information avoidance: Understanding the health information behavior during a global health crisis," *Inf. Process. Manage.*, vol. 58, no. 2, Mar. 2021, Art. no. 102440.

[21] J. Heer and B. Shneiderman, "Interactive dynamics for visual analysis: A taxonomy of tools that support the fluent and flexible use of visualizations," *Queue*, vol. 10, no. 2, pp. 30–55, Feb. 2012.

[22] E. Segel and J. Heer, "Narrative visualization: Telling stories with data," *IEEE Trans. Vis. Comput. Graphics*, vol. 16, no. 6, pp. 1139–1148, Nov. 2010.

[23] J. Wenskovitch and C. North, "Interactive artificial intelligence: Designing for the," *Computer*, vol. 53, no. 8, pp. 29–39, 2020.

[24] J. Wenskovitch, C. Fallon, K. Miller, and A. Dasgupta, "Beyond visual analytics: human-machine teaming for AI-driven data sensemaking," in *Proc. IEEE Workshop Trust Expertise Vis. Analytics (TREX)*, New Orleans, LA, USA, Oct. 2021, pp. 40–44.

[25] D. Shahaf and C. Guestrin, "Connecting the dots between news articles," in *Proc. 16th ACM SIGKDD Int. Conf. Knowl. discovery data mining*, New York, NY, USA, Jul. 2010, pp. 623–632.

[26] B. Keith and T. Mitra, "Narrative maps: An algorithmic approach to represent and extract information narratives," in *Proc. ACM Hum.-Comput. Interact.*, New York, NY, USA, Dec. 2020, pp. 1–33.

[27] S. Angus, S. Borgohain, S. Fan, S. Goodwin, A. Kristanto, L. O'Neill, H. C. Purchase, N. Van Nieuwenhove, Y. Yang, Y. Zhang, and T. Dwyer, "Discourse lines: Visualising current policy and media storylines of opportunity and disadvantage with narrative exploration maps," in *Proc. NLVIZ*, Melbourne, VIC, Australia, Jun. 2023, pp. 1–7.

[28] A. S. Vivacqua, R. Stelling, A. C. B. Garcia, and L. C. Gouveia, "Explanations and sensemaking with AI and HCI," in *Proc. 9th Latin Amer. Conf. Human Comput. Interact.*, New York, NY, USA, Sep. 2019, pp. 1–4.

[29] Y. Zhao, Y. Zhang, Y. Zhang, X. Zhao, J. Wang, Z. Shao, C. Turkay, and S. Chen, "LEVA: Using large language models to enhance visual analytics," *IEEE Trans. Vis. Comput. Graphics*, vol. 31, no. 3, pp. 1830–1847, Mar. 2025.

[30] Y. Chang, X. Wang, J. Wang, Y. Wu, L. Yang, K. Zhu, H. Chen, X. Yi, C. Wang, Y. Wang, W. Ye, Y. Zhang, Y. Chang, P. S. Yu, Q. Yang, and X. Xie, "A survey on evaluation of large language models," *ACM Trans. Intell. Syst. Technol.*, vol. 15, no. 3, pp. 1–45, 2024.

[31] L. Yang, H. Chen, Z. Li, X. Ding, and X. Wu, "Give us the facts: Enhancing large language models with knowledge graphs for fact-aware language modeling," *IEEE Trans. Knowl. Data Eng.*, vol. 36, no. 7, pp. 3091–3110, Jul. 2024.

[32] H. Han, R. Faust, B. F. Keith Norambuena, J. Lin, S. Li, and C. North, "Explainable interactive projections of images," *Mach. Vis. Appl.*, vol. 34, no. 6, p. 100, Nov. 2023.

[33] Y. Bian and C. North, *DeepSI: Interactive Deep Learning for Semantic Interaction*. New York, NY, USA: ACM, 2021, pp. 197–207.

[34] D. Keim, G. Andrienko, J.-D. Fekete, C. Görg, J. Kohlhammer, and G. Melançon, *Visual Analytics: Definition, Process, and Challenges*. Berlin, Germany: Springer, 2008, pp. 154–175.

[35] Z. Yan and X. Tang, "Narrative graph: Telling evolving stories based on event-centric temporal knowledge graph," *J. Syst. Sci. Syst. Eng.*, vol. 32, no. 2, pp. 206–221, Apr. 2023.

[36] B. F. Keith Norambuena, T. Mitra, and C. North, "Mixed multi-model semantic interaction for graph-based narrative visualizations," in *Proc. 28th Int. Conf. Intell. User Interfaces*, New York, NY, USA, Mar. 2023, pp. 866–888.

[37] J. J. Thomas and K. A. Cook, *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. Richland, WA, USA: National Visualization and Analytics Center, 2005.

[38] H. L. Chieu and Y. K. Lee, "Query based event extraction along a timeline," in *Proc. 27th Annu. Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, New York, NY, USA, Jul. 2004, pp. 425–432.

[39] Q. Chen, S. Cao, J. Wang, and N. Cao, "How does automation shape the process of narrative visualization: A survey of tools," *IEEE Trans. Vis. Comput. Graphics*, vol. 30, no. 8, pp. 4429–4448, Aug. 2024.

[40] B. Liu, F. X. Han, D. Niu, L. Kong, K. Lai, and Y. Xu, "Story forest: Extracting events and telling stories from breaking news," *ACM Trans. Knowl. Discovery from Data*, vol. 14, no. 3, pp. 1–28, Jun. 2020.

[41] I. Blin, "Building narrative structures from knowledge graphs," in *The Semantic Web: ESWC 2022 Satellite Events*, P. Groth, A. Rula, J. Schneider, I. Tiddi, E. Simperl, P. Alexopoulos, R. Hoekstra, M. Alam, A. Dimou, and M. Tamper, Eds., Cham, Switzerland: Springer, 2022, pp. 234–251.

[42] M. Dowling, N. Wycoff, B. Mayer, J. Wenskovitch, S. Leman, L. House, N. Polys, C. North, and P. Hauck, "Interactive visual analytics for sensemaking with big text," *Big Data Res.*, vol. 16, pp. 49–58, Jul. 2019.

[43] R. Yan, X. Wan, J. Otterbacher, L. Kong, X. Li, and Y. Zhang, "Evolutionary timeline summarization: A balanced optimization framework via iterative substitution," in *Proc. 34th Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, New York, NY, USA, Jul. 2011, pp. 745–754.

[44] D. Shahaf, J. Yang, C. Suen, J. Jacobs, H. Wang, and J. Leskovec, "Information cartography: Creating zoomable, large-scale maps of information," in *Proc. 19th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, New York, NY, USA, Aug. 2013, pp. 1097–1105.

[45] M. Shukla, R. Dos Santos, F. Chen, and C.-T. Lu, "DISCRN: A distributed storytelling framework for intelligence analysis," *Big Data*, vol. 5, no. 3, pp. 225–245, Sep. 2017.

[46] J. E. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in *Proc. Int. Conf. Learn. Represent.*, 2022.

[47] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer, "QLoRA: Efficient finetuning of quantized LLMs," in *Proc. Adv. Neural Inf. Process. Syst.*, 2023, pp. 10088–10115.

[48] A. Ulmer, M. Angelini, J.-D. Fekete, J. Kohlhammer, and T. May, "A survey on progressive visualization," *IEEE Trans. Vis. Comput. Graphics*, vol. 30, no. 9, pp. 6447–6467, Sep. 2024.

[49] E. Cakmak, U. Schlegel, D. Jäckle, D. Keim, and T. Schreck, "Multiscale snapshots: Visual analysis of temporal summaries in dynamic graphs," *IEEE Trans. Vis. Comput. Graphics*, vol. 27, no. 2, pp. 517–527, Feb. 2021.

[50] S. Frenda, G. Abercrombie, V. Basile, A. Pedrani, R. Panizzon, A. T. Cignarella, C. Marco, and D. Bernardi, "Perspectivist approaches to natural language processing: A survey," *Lang. Resour. Eval.*, vol. 59, no. 2, pp. 1719–1746, Jun. 2025.

[51] E. Leonardelli, G. Abercrombie, D. Almanea, V. Basile, T. Fornaciari, B. Plank, V. Rieser, A. Uma, and M. Poesio, "SemEval-2023 task 11: Learning with disagreements (LeWiDi)," in *Proc. The 17th Int. Workshop Semantic Eval.*, 2023, pp. 2304–2318.

[52] M. Angelini, G. Santucci, H. Schumann, and H. Schulz, "A review and characterization of progressive visual analytics," *Informatics*, vol. 5, no. 3, p. 31, 2018.

[53] J.-D. Fekete and R. Primet, "Progressive analytics: A computation paradigm for exploratory data analysis," 2016, *arXiv:1607.05162*.

[54] W. Kintsch, *Comprehension: A Paradigm for Cognition*. Cambridge, U.K.: Cambridge Univ. Press, 1998.

[55] R. A. Zwaan and G. A. Radvansky, "Situation models in language comprehension and memory," *Psychol. Bull.*, vol. 123, no. 2, pp. 162–185, 1998.

[56] J. S. Bruner, *Actual Minds, Possible Worlds*. Cambridge, MA, USA: Harvard Univ. Press, 2009.

[57] R. Porzel, M. Pomarlan, L. Spillner, J. Bateman, T. Mildner, and C. R. M. A. Santagiustina, "Narrativizing knowledge graphs," in *Proc. Int. Workshop Knowl. Graph Summarization*, Hangzhou, China, Oct. 2022, pp. 100–113.

[58] B. Shneiderman, "The eyes have it: A task by data type taxonomy for information visualizations," in *Proc. IEEE Symp. Vis. Lang.*, Boulder, CO, USA, Feb. 1996, pp. 336–343, doi: [10.1109/VL.1996.545307](https://doi.org/10.1109/VL.1996.545307).

[59] W. Pike, J. Stasko, R. Chang, and T. O'Connell, "The science of interaction," *Inf. Visualizat.*, vol. 8, no. 4, pp. 263–274, 2009.

[60] C. Trojahn, R. Vieira, D. Schmidt, A. Pease, and G. Guizzardi, "Foundational ontologies meet ontology matching: A survey," *Semantic Web*, vol. 13, no. 4, pp. 685–704, May 2022.

[61] S. Choudhary, N. Chatterjee, and S. K. Saha, "Interpretation of black box NLP models: A survey," 2022, *arXiv:2203.17081*.

[62] X. Li and P. Liang, "Prefix-tuning: Optimizing continuous prompts for generation," in *Proc. 59th Annu. Meeting Assoc. Comput. Linguistics 11th Int. Joint Conf. Natural Lang. Process.*, 2021, pp. 4582–4597.

[63] B. Lester, R. Al-Rfou, and N. Constant, "The power of scale for parameter-efficient prompt tuning," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2021, pp. 3045–3059.

[64] E. Mosca, F. Szigeti, S. Tragianni, D. Gallagher, and G. Groh, "Shape-based explanation methods: A review for nlp interpretability," in *Proc. 29th Int. Conf. Comput. Linguistics*, 2022, pp. 4593–4603.

[65] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you? Explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2016, pp. 1135–1144.

[66] K. Brian, G. Fausto, K. Eric, J. Sarah, and N. Chris, "Explainable AI components for narrative map extraction," in *Proc. Text2Story Workshop*, 2025, pp. 3–14.

[67] T. Wu, M. T. Ribeiro, J. Heer, and D. Weld, "Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models," in *Proc. 59th Annu. Meeting Assoc. Comput. Linguistics, 11th Int. Joint Conf. Natural Lang. Process.*, 2021, pp. 6707–6723.

[68] Z. Chen, Q. Gao, A. Bosselut, A. Sabharwal, and K. Richardson, "DISCO: Distilling counterfactuals with large language models," in *Proc. 61st Annu. Meeting Assoc. Comput. Linguistics*, vol. 1, 2023, pp. 5514–5528.

[69] H. Chefer, S. Gur, and L. Wolf, "Transformer interpretability beyond attention visualization," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 782–791.

[70] Z. Jin, J. von Kügelgen, J. Ni, T. Vaidhya, A. Kaushal, M. Sachan, and B. Schoelkopf, "Causal direction of data collection matters: Implications of causal and anticausal learning for NLP," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2021, pp. 9499–9513.

[71] Y. Zhou and Y. He, "Causal inference from text: Unveiling interactions between variables," in *Proc. Findings Assoc. Comput. Linguistics*, 2023, pp. 10559–10571.

[72] M. Krstajic, E. Bertini, and D. Keim, "CloudLines: Compact display of event episodes in multiple time-series," *IEEE Trans. Vis. Comput. Graphics*, vol. 17, no. 12, pp. 2432–2439, Dec. 2011.

[73] P. N. Johnson-Laird, *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Cambridge, MA, USA: Harvard Univ. Press, 1983.

[74] G. Repovš and A. Baddeley, "The multi-component model of working memory: Explorations in experimental cognitive psychology," *Neuroscience*, vol. 139, no. 1, pp. 5–21, Apr. 2006.

[75] L. Hasher and R. T. Zacks, "Working memory, comprehension, and aging: A review and a new view," *Psychol. Learn. Motivat.*, vol. 22, pp. 193–225, Oct. 1988.

[76] H. A. Slagter, A. Lutz, L. L. Greischar, A. D. Francis, S. Nieuwenhuis, J. M. Davis, and R. J. Davidson, "Mental training affects distribution of limited brain resources," *PLoS Biol.*, vol. 5, no. 6, p. e138, May 2007.

[77] D. A. Fair, N. U. F. Dosenbach, J. A. Church, A. L. Cohen, S. Brahmhadt, F. M. Miezin, D. M. Barch, M. E. Raichle, S. E. Petersen, and B. L. Schlaggar, "Development of distinct control networks through segregation and integration," *Proc. Nat. Acad. Sci. USA*, vol. 104, no. 33, pp. 13507–13512, Aug. 2007.

[78] M. Wallinger, D. Archambault, D. Auber, M. Nöllenburg, and J. Peltonen, "Edge-path bundling: A less ambiguous edge bundling approach," *IEEE Trans. Vis. Comput. Graphics*, vol. 28, no. 1, pp. 313–323, Jan. 2022.

[79] A. Lhuillier, C. Hurter, and A. Telea, "State of the art in edge and trail bundling techniques," in *Computer Graphics Forum*, vol. 36. Hoboken, NJ, USA: Wiley, 2017, pp. 619–645.

[80] W. Yang, X. Wang, J. Lu, W. Dou, and S. Liu, "Interactive steering of hierarchical clustering," *IEEE Trans. Vis. Comput. Graphics*, vol. 27, no. 10, pp. 3953–3967, Oct. 2021.

[81] Y. Lerner, C. J. Honey, L. J. Silbert, and U. Hasson, "Topographic mapping of a hierarchy of temporal receptive windows using a narrated story," *J. Neurosci.*, vol. 31, no. 8, pp. 2906–2915, Feb. 2011.

[82] L. M. Bietti, O. Tilstan, and A. Bangerter, "Storytelling as adaptive collective sensemaking," *Topics Cognit. Sci.*, vol. 11, no. 4, pp. 710–732, Oct. 2019.

[83] C. D. Stolper, A. Perer, and D. Gotz, "Progressive visual analytics: User-driven visual exploration of in-progress analytics," *IEEE Trans. Vis. Comput. Graphics*, vol. 20, no. 12, pp. 1653–1662, Dec. 2014.

[84] E. Zgraggen, Z. Zhao, R. Zeleznik, and T. Kraska, "Investigating the effect of the multiple comparisons problem in visual analysis," in *Proc. CHI Conf. Human Factors Comput. Syst.*, Apr. 2018, pp. 1–12.

[85] G. M. León, A. Bezerianos, O. Gladin, and P. Isenberg, "Talk to the wall: The role of speech interaction in collaborative visual analytics," *IEEE Trans. Vis. Comput. Graphics*, vol. 31, no. 1, pp. 941–951, Jan. 2025.

[86] L. Shen, E. Shen, Y. Luo, X. Yang, X. Hu, X. Zhang, Z. Tai, and J. Wang, "Towards natural language interfaces for data visualization: A survey," *IEEE Trans. Vis. Comput. Graphics*, vol. 29, no. 6, pp. 3121–3144, Jun. 2023.

[87] K. Liang, L. Meng, M. Liu, Y. Liu, W. Tu, S. Wang, S. Zhou, X. Liu, F. Sun, and K. He, "A survey of knowledge graph reasoning on graph types: Static, dynamic, and multi-modal," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 46, no. 12, pp. 9456–9478, Dec. 2024.

[88] Z. Deng, W. Wang, Z. Wang, X. Liu, and Y. Song, "Gold: A global and local-aware denoising framework for commonsense knowledge graph noise detection," in *Proc. Findings Assoc. Comput. Linguistics*, 2023, pp. 3591–3608.

[89] L. F. Sikos and D. Philp, "Provenance-aware knowledge representation: A survey of data models and contextualized knowledge graphs," *Data Sci. Eng.*, vol. 5, no. 3, pp. 293–316, Sep. 2020.

[90] T. Fornaciari, A. Uma, S. Paun, B. Plank, D. Hovy, and M. Poesio, "Beyond black & white: Leveraging annotator disagreement via soft-label multi-task learning," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol.*, 2021, pp. 2591–2597.

[91] P. Rottger, B. Vidgen, D. Hovy, and J. Pierrehumbert, "Two contrasting data annotation paradigms for subjective NLP tasks," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol.*, 2022, pp. 175–190.

[92] E. Fleisig, R. Abebe, and D. Klein, "When the majority is wrong: Modeling annotator disagreement for subjective tasks," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2023, pp. 6715–6726.

[93] B. Friedman, M. Harbers, D. G. Hendry, J. van den Hoven, C. Jonker, and N. Logler, "Eight grand challenges for value sensitive design from the 2016 Lorentz workshop," *Ethics Inf. Technol.*, vol. 23, no. 1, pp. 5–16, Mar. 2021.

[94] T. H. Davenport and J. C. Beck, "The attention economy," *Sci. Amer.*, vol. 2001, no. May, p. 1, 2001.

[95] E. Aïmeur, S. Amri, and G. Brassard, "Fake news, disinformation and misinformation in social media: A review," *Social Netw. Anal. Mining*, vol. 13, no. 1, pp. 30, Feb. 2023.

[96] K. Sharma, F. Qian, H. Jiang, N. Ruchansky, M. Zhang, and Y. Liu, “Combating fake news: A survey on identification and mitigation techniques,” *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 3, pp. 1–42, May 2019.

[97] M. T. Fischer, S. D. Hirsbrunner, W. Jentner, M. Miller, D. A. Keim, and P. Helm, “Promoting ethical awareness in communication analysis: Investigating potentials and limits of visual analytics for intelligence applications,” in *Proc. ACM Conf. Fairness Accountability Transparency*, New York, NY, USA, Jun. 2022, pp. 877–889.

[98] M. T. Fischer, Y. Metz, L. Joos, M. Miller, and D. A. Keim, “Multi-case: A transformer-based ethics-aware multimodal investigative intelligence framework,” 2024, *arXiv:2401.01955*.

[99] I. O. Gallegos, R. A. Rossi, J. Barrow, M. M. Tanjim, S. Kim, F. Dernoncourt, T. Yu, R. Zhang, and N. K. Ahmed, “Bias and fairness in large language models: A survey,” *Comput. Linguistics*, vol. 50, no. 3, pp. 1097–1179, Sep. 2024.

[100] J. Fu, Y. Hong, X. Ling, L. Wang, X. Ran, Z. Sun, W. H. Wang, Z. Chen, and Y. Cao, “Differentially private federated learning: A systematic review,” 2024, *arXiv:2405.08299*.

[101] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, and T. Gebru, “Model cards for model reporting,” in *Proc. Conf. Fairness, Accountability, Transparency*, Jan. 2019, pp. 220–229.

[102] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. D. Iii, and K. Crawford, “Datasheets for datasets,” *Commun. ACM*, vol. 64, no. 12, pp. 86–92, Dec. 2021.

[103] H. J. Wilson and P. R. Daugherty, “Collaborative intelligence: Humans and ai are joining forces,” *Harvard Bus. Rev.*, vol. 96, no. 4, pp. 114–123, 2018.

[104] M. Vaccaro, A. Almaatouq, and T. Malone, “When combinations of humans and AI are useful: A systematic review and meta-analysis,” *Nature Human Behaviour*, vol. 8, no. 12, pp. 2293–2303, Oct. 2024.

[105] J. Ansah, L. Liu, W. Kang, S. Kwashie, J. Li, and J. Li, “A graph is worth a thousand words: Telling event stories using timeline summarization graphs,” in *Proc. World Wide Web Conf.*, New York, NY, USA, May 2019, pp. 2565–2571.

[106] B. Liu, D. Niu, K. Lai, L. Kong, and Y. Xu, “Growing story forest online from massive breaking news,” in *Proc. ACM Conf. Inf. Knowl. Manage.*, New York, NY, USA, Nov. 2017, pp. 777–785.

[107] A. Nandi, M. Sundriyal, E. M. Khan, J. Sun, E. K. Vraga, J. Srivastava, and T. Chakraborty, “The psychology of falsehood: A human-centric survey of misinformation detection,” in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2025, pp. 8519–8536.

[108] W. Sosnowski, A. Modzelewski, K. Skorupska, and A. Wierzbicki, “DiNaM: Disinformation narrative mining with large language models,” in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2025, pp. 30212–30239.

[109] Z. Zhu, Z. Zhang, H. Zhang, and C. Li, “RATSD: Retrieval augmented truthfulness stance detection from social media posts toward factual claims,” in *Proc. Findings Assoc. Comput. Linguistics*, 2025, pp. 3366–3381.

[110] M. Schlichtkrull, Z. Guo, and A. Vlachos, “AVeriTeC: A dataset for real-world claim verification with evidence from the Web,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2023, pp. 65128–65167.

[111] J. Thorne, A. Vlachos, C. Christodoulopoulos, and A. Mittal, “FEVER: A large-scale dataset for fact extraction and verification,” in *Proc. Conf. North Amer. Chapter ACL*, 2018, pp. 809–819.

[112] A. Cosentino, C. De Maio, D. Furno, M. Gallo, and V. Loia, “Source credibility assessment in the realm of information disorder: A literature review,” *Int. J. Interact. Multimedia Artif. Intell.*, vol. 2025, pp. 1–15, Jan. 2025.

[113] L. Wu, Y. Long, C. Gao, Z. Wang, and Y. Zhang, “MFIR: Multimodal fusion and inconsistency reasoning for explainable fake news detection,” *Inf. Fusion*, vol. 100, Dec. 2023, Art. no. 101944.

[114] A. Ganti, E. A. H. Hussein, S. Wilson, Z. Ma, and X. Zhao, “Narrative style and the spread of health misinformation on Twitter,” in *Proc. Findings Assoc. Comput. Linguistics*, 2023, pp. 4266–4282.



BRIAN KEITH received the B.Sc. degree in engineering from the Universidad Católica del Norte (UCN), Antofagasta, Chile, in 2016, the B.Sc. degree in mathematics and the M.Sc. degree in informatics engineering from UCN, in 2017, and the Ph.D. degree in computer science and applications from Virginia Tech, Blacksburg, VA, USA, in 2023. He received the professional title in computing and informatics civil engineering from UCN. He is currently an Assistant Professor with the Department of Systems and Computing Engineering and a Secretary of research and technological development with the Faculty of Engineering and Geological Sciences, UCN. He is also the Director of the Artificial Intelligence Innovation Center for the Antofagasta Region (CIARA). He is the author of more than 60 research articles. His research interests include visual analytics, artificial intelligence, text analytics, computational narratives, and applied data analytics in mining, geochemistry, and education. He was a recipient of the Fulbright Faculty Development Scholarship (2019–2021) and the Becas Chile Doctoral Studies Scholarship (2019–2023). He was the Valedictorian of the 2016 and 2017 Graduating Classes at UCN. He is an Associate Editor of *Intelligent Data Analysis*.