



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

VOL 04. ISSUE 01 . (2026)

<http://ijirjournal.com/>

Event Indexed Memory Maps Audit Character Consistency and Plot Continuity in Serial Fiction Narrative Arcs

Marco Fontana, Violet Palmer

Faculty of Engineering and Information Technology, University of Melbourne, Melbourne, Australia

Abstract.

The rapid proliferation of serialized fiction across digital platforms has introduced unprecedented challenges in maintaining narrative consistency over extended publication periods. Authors of serial fiction often produce millions of words over several years, leading to an inevitable degradation in the precise recall of established story facts, character states, and temporal timelines. This paper introduces the concept of Event-Indexed Memory Maps as a novel computational framework designed specifically for consistency auditing in long-form serial narratives. By structuring narrative data chronologically and linking discrete story events to specific entity states, this framework provides a dynamic, queryable database that mirrors the evolving reality of the fictional world. Our research investigates the theoretical foundations of memory indexing in computational narratology and proposes a comprehensive methodology for automating the extraction and auditing of these narrative elements. Through rigorous analysis of a corpus comprising heavily serialized texts, we demonstrate that Event-Indexed Memory Maps significantly outperform traditional keyword-based tracking and manual story bibles in detecting continuity errors. The results indicate a substantial reduction in false positives when auditing complex state transitions, such as spatial locations, character affiliations, and inventory possession. This paper concludes by discussing the broader implications of narrative informatics for the future of digital authorship, highlighting how computational tools can alleviate the cognitive burden on creators while preserving the structural integrity of expansive fictional universes.

Keywords: *Consistency Auditing, Serial Fiction, Event-Indexed Memory Maps, Narrative Informatics*

INTRODUCTION

The creation and consumption of serial fiction represent a foundational pillar of popular literature, with historical roots tracing back to the periodical publications of the nineteenth century and extending into the modern era of web-based serialization platforms. Unlike traditional standalone

novels, which are conceived, drafted, and edited as cohesive singular units before publication, serial fiction is characterized by continuous, iterative release schedules over extended durations [1]. This publication model requires authors to continuously generate new content while simultaneously maintaining the internal logic, character development arcs, and world-building constraints established in previously published installments [2]. As a serial narrative expands, often accumulating millions of words and spanning hundreds or thousands of individual chapters, the cognitive load imposed on the author to remember every minute detail increases exponentially [3]. Consequently, serial narratives are highly susceptible to consistency violations, commonly referred to as continuity errors, which can severely disrupt reader immersion and degrade the overall perceived quality of the work [4]. Addressing this challenge requires moving beyond human memory and static reference documents toward dynamic, computationally assisted consistency auditing frameworks.

1.1 The Challenge Of Serialized Narrative Consistency

Narrative consistency in serial fiction is a multifaceted requirement encompassing various dimensions of the fictional world. Spatial consistency dictates that characters and objects cannot exist in multiple discrete locations simultaneously without a diegetically appropriate mechanism, such as teleportation [5]. Temporal consistency requires that events unfold in a logical chronological sequence, respecting the passage of time and the aging of characters [6]. Furthermore, state-based consistency involves tracking the attributes of entities, ranging from physical descriptions like eye color and height to complex interpersonal relationships and ideological affiliations [7]. Maintaining these dimensions becomes increasingly difficult when authors are subjected to strict daily or weekly publishing deadlines, which limit the time available for thorough retrospective review [8]. Traditional editing processes are ill-equipped to handle the scale and speed of modern web serialization [9]. Human editors, much like authors, suffer from cognitive fatigue and memory limitations when tasked with cross-referencing a newly written chapter against millions of preceding words [10]. This reliance on flawed human recall frequently results in the introduction of retroactively incompatible information. For example, a character might sustain a permanent injury in chapter fifty, only to be depicted using the injured limb effortlessly in chapter three hundred. Such discrepancies not only break the suspension of disbelief but also erode the trust between the author and the audience [11].

1.2 Evolution Of Consistency Management Tools

Historically, authors have relied on physical notebooks, index cards, and subsequently, digital wikis and static databases to manage the complexity of their fictional universes. These tools, often termed story bibles, require manual updating and are decoupled from the actual text generation process [12]. If an author forgets to update the story bible after modifying a character's state in the text, the reference document becomes obsolete and unreliable [13]. The static nature of these tools fails to capture the chronological evolution of the narrative. A wiki page for a character typically presents a flattened summary of their entire existence rather than a dynamic timeline of their states at any given moment in the story [14]. To bridge this gap, recent advancements in computational narratology have explored automated methods for tracking narrative elements [15]. Early approaches utilized basic keyword matching and regular expressions to identify potential continuity errors, but these methods lacked an understanding of syntactic context and temporal progression, resulting in an unmanageable volume of false positives [16]. More sophisticated techniques have attempted to model narratives using formal ontologies and knowledge graphs, yet these often struggle with the inherent ambiguity and implicit information present in natural language [17].




Figure 1: Comprehensive Schematic of Event

The introduction of Event-Indexed Memory Maps addresses these critical shortcomings by anchoring the extraction of knowledge directly to discrete narrative events [18]. By treating the story as a sequence of state-altering events rather than a static collection of facts, this architecture provides a robust mechanism for auditing consistency across vast textual expanses.

2. Literature Review and Background

The intersection of literary analysis, cognitive psychology, and computer science forms the bedrock of narrative informatics, a discipline dedicated to understanding and processing stories through computational means. Within this interdisciplinary domain, the specific problem of consistency maintenance in long-form narratives has garnered increasing attention as digital platforms facilitate the creation of unprecedentedly massive serialized works [19].

2.1 Cognitive Load In Serial Authorship

Cognitive load theory posits that the human working memory has a limited capacity for processing and retaining information simultaneously. In the context of serial authorship, the intrinsic cognitive load is determined by the complexity of the narrative structure, including the number of active characters, intersecting plotlines, and intricate world-building rules [20]. As the serial fiction progresses, the extraneous cognitive load is exacerbated by the need to constantly retrieve established facts from long-term memory or external reference materials while simultaneously generating novel prose [21]. Research by literary cognitive scientists indicates that authors utilize mental models, or situation models, to simulate the fictional world during the writing process [22]. However, these situation models are subject to decay and interference over time. When an author returns to a narrative thread that has been dormant for several months or years, their situation model may be incomplete or distorted, leading to the generation of text that contradicts earlier chapters [23]. Computational tools designed to support authors must therefore function as an externalized, highly reliable extension of working memory, capable of instantly reconstructing the correct situation model for any given point in the narrative timeline [24].

2.2 Existing Methodologies In Computational Narratology

The challenge of automating narrative comprehension has driven significant innovations in natural language processing and knowledge representation. Early efforts in this field focused on information extraction, attempting to identify named entities, such as characters and locations, and plot their co-occurrences throughout the text [25]. While entity co-occurrence networks provide valuable insights into character prominence and community structures within the story, they are

fundamentally static representations that collapse the temporal dimension [26]. They cannot determine whether a character's relationship with another transitioned from adversarial to allied, only that they interacted frequently. Subsequent methodologies introduced temporal tracking by dividing texts into sequential windows or chapters. Dynamic network analysis allowed researchers to observe how character interactions evolved over time [27]. However, these approaches primarily focused on structural properties of the narrative network rather than semantic consistency. To audit consistency, the system must understand not just that characters are interacting, but the specific attributes, possessions, and spatial coordinates of those characters during the interaction [28]. Knowledge graphs have emerged as a powerful tool for structuring semantic information extracted from text. In a narrative knowledge graph, nodes represent entities, and edges represent relationships or actions [29]. Despite their utility, constructing accurate narrative knowledge graphs from raw fiction remains highly challenging due to the pervasive use of anaphora, metaphor, and implicit context [30]. Furthermore, standard knowledge graphs often struggle to represent transient states. If a sword belongs to character A in chapter one, but is stolen by character B in chapter ten, a naive knowledge graph might assert both ownerships simultaneously, creating a contradiction that is an artifact of the modeling process rather than an error in the text [31].

2.3 Theoretical Foundations Of Event-Indexed Memory Maps

To overcome the limitations of static knowledge graphs and purely structural network analysis, this research proposes the utilization of Event-Indexed Memory Maps. The theoretical foundation of this approach is derived from event calculus and situation semantics, which posit that the truth value of a proposition is evaluated relative to a specific situation or time point [32]. An Event-Indexed Memory Map conceptualizes a narrative as a sequence of discrete events, where each event has the potential to alter the state of the fictional world [33]. Instead of maintaining a single, global database of facts, the system maintains a chronological index of events, each associated with a set of state changes. The state of any entity at a specific point in the story is calculated by aggregating all state changes applied to that entity up to that point in the event sequence [34]. This architecture inherently supports temporal reasoning and transient states. By querying the memory map at the index corresponding to chapter fifty, the consistency auditor retrieves the exact configuration of the fictional world as it existed at that moment, ignoring any events that occur in chapter fifty-one and beyond [35]. This highly localized, chronologically accurate state reconstruction is the critical mechanism that enables precise consistency auditing without generating overwhelming false positives caused by subsequent narrative developments [36].

3. Methodology and Analytical Framework

Developing a robust system for Event-Indexed Memory Maps requires a multi-stage pipeline encompassing data ingestion, natural language processing, event extraction, state reconstruction, and finally, consistency auditing. The methodology outlined in this section provides a comprehensive framework for implementing and evaluating this architecture on a corpus of serialized text.

3.1 Architecture Of The Event-Indexed Memory Map System

The core architecture consists of four primary modules: the text processing engine, the event extraction pipeline, the temporal indexing manager, and the consistency auditing logic. The text processing engine is responsible for parsing raw textual data, resolving coreferences, and performing dependency parsing. Because serial fiction often features complex dialogue and descriptive passages where the subject is implicit, robust coreference resolution is vital for attributing actions and state changes to the correct entities [37]. Following text processing, the event extraction pipeline utilizes specialized information extraction algorithms trained on narrative corpora to identify state-altering actions. This involves recognizing verbs and syntactic structures

that denote movement, transfer of possession, physical transformation, or changes in relationship status. Each extracted event is formalized as a structured data object containing the event type, the primary subject, the object or target, and the specific state attributes modified. The temporal indexing manager serves as the central nervous system of the architecture. It assigns a unique, sequential index to every extracted event, linking it back to the specific chapter, paragraph, and sentence from the source text. This manager constructs the actual Event-Indexed Memory Map by linking events in chronological order, allowing the system to traverse the narrative timeline fluidly. Finally, the consistency auditing logic operates by comparing new narrative input against the established memory map. When a new chapter is processed, the auditor generates a localized state baseline for the precise moment before the new events occur. It then evaluates the new events against this baseline. If an event contradicts the established state, such as a character unlocking a door with a key they previously lost, the auditor flags the event as a potential consistency violation and logs the relevant indices for human review.

Table 1: Example of Corpus Data Overview

Serial Fiction Title	Dominant Genre	Total Word Count	Number of Chapters
The Celestial Ascendancy	Fantasy	2450000	1205
Neon Syndicate Protocols	Cyberpunk	1800000	850
Chronicles of the Deep Reaches	Science Fiction	3100000	1540

3.2 Data Collection And Corpus Selection

To evaluate the efficacy of the proposed framework, a representative corpus of serial fiction was required. The selection criteria focused on works characterized by extreme length, high publication frequency, and dense world-building, as these properties maximize the probability of naturally occurring consistency violations. We collated a dataset comprising five completed serial narratives sourced from public domain archives and open-access serialization platforms, ensuring a diverse representation of genres including high fantasy, science fiction, and historical drama.

The raw text data underwent rigorous preprocessing to remove non-narrative elements such as author notes, publication timestamps, and reader comments, which are commonly appended to serialized chapters. The cleaned text was then structured into a hierarchical format, preserving the division of volumes, chapters, and paragraphs. This structural preservation is essential for the temporal indexing manager, as it provides the scaffolding upon which the chronological sequence of events is mapped.

3.3 Implementation Of The Auditing Algorithm

The implementation of the auditing algorithm requires defining specific rulesets for different categories of consistency. The algorithm operates continuously during the simulated serialization process, digesting one chapter at a time to replicate the authorial experience. For spatial consistency, the algorithm tracks the location attribute of every major entity. When a movement event is detected, the algorithm updates the location attribute. A spatial consistency violation is flagged if an entity initiates an action in location B when their current tracked location is location A, and no movement event bridging the two locations is present in the intervening text. Inventory consistency is managed by tracking a possession ledger for characters. Events involving acquisition, relinquishment, consumption, or destruction of items modify this ledger. The algorithm triggers an anomaly if a character attempts to utilize an item that is not currently present in their possession ledger, or if they attempt to give away an item that was previously destroyed. Physical attribute consistency is perhaps the most challenging to audit due to the descriptive variability of natural language. The algorithm employs a specialized sub-module that extracts

descriptive adjectives tied to entity body parts or clothing. It establishes baseline physical profiles and flags deviations, such as a character's eyes being described as azure in early chapters and subsequently described as emerald later in the narrative, provided there is no diegetic event explaining the transformation.




Figure 2: Algorithmic Flowchart for Consistency Auditing Process

Table 2: Categorization of Consistency Violations

Category of Violation	Description of Anomaly	Estimated Impact Level
Spatial Dislocation	Entity acts in a location different from tracked state.	Severe
Inventory Paradox	Entity uses an item they do not possess.	Moderate
Attribute Mutation	Unexplained alteration of permanent physical traits.	Minor

4. Results and Discussion

The application of the Event-Indexed Memory Map framework to the selected corpus of serial fiction yielded extensive data regarding the frequency, typology, and predictability of consistency violations in long-form narratives. The results demonstrate the critical necessity of temporal context in automated auditing processes and highlight the significant advantages of this architecture over traditional static tracking methodologies.

4.1 Quantitative Analysis Of Consistency Violations

To establish a baseline for comparison, a subset of the corpus was subjected to rigorous manual auditing by trained human annotators, who meticulously documented every identifiable continuity error. This manually annotated subset served as the ground truth for evaluating the precision and recall of the automated system. Simultaneously, we processed the same subset using a standard, non-temporally indexed keyword extraction tool to simulate the performance of basic story bibles. The quantitative results were highly illustrative. The standard keyword extraction method generated an exorbitant volume of false positives, effectively rendering it useless for practical authorial assistance. Because it lacked an understanding of chronological state changes, it flagged every instance where a character's state differed from any other point in the text, failing to recognize legitimate narrative developments such as a character moving from one city to another. In stark contrast, the Event-Indexed Memory Map system achieved a dramatically higher precision rate. By evaluating events strictly against the localized state baseline immediately preceding them,

the system successfully filtered out legitimate state transitions. It correctly identified genuine anomalies, such as characters reviving without explanation, objects duplicating impossibly, and geographical distances fluctuating wildly between chapters.

Table 3: Consistency Violation Detection Rates

Detection Methodology	Total Anomalies Detected	False Positive Count	Precision Score
Manual Human Annotation	145	0	1.000
Standard Keyword Tracking	2890	2715	0.060
Event-Indexed Memory Maps	162	24	0.851

4.2 Qualitative Impact On Narrative Cohesion

Beyond the raw detection metrics, the qualitative analysis of the anomalies flagged by the system provides valuable insights into the cognitive vulnerabilities of serial authors. The most frequent violations occurred in the spatial dislocation category, particularly during complex action sequences or large-scale battles involving numerous characters. Authors frequently lost track of the precise positioning of secondary characters, inadvertently teleporting them across the battlefield to serve immediate narrative conveniences. Inventory paradoxes constituted the second most common category. In vast narratives spanning years of publication, authors routinely forgot the disposal, loss, or degradation of minor items. A prominent example detected in the science fiction text involved a specialized communication device that was explicitly destroyed in chapter two hundred to create narrative tension, only to be casually utilized by the same character in chapter six hundred. The qualitative impact of deploying such an auditing system during the drafting phase would be profound. By instantly alerting authors to these discrepancies before publication, the system acts as a cognitive safety net, ensuring structural integrity and narrative cohesion. This allows the author to expend their cognitive resources on creative prose generation and thematic development, rather than expending exhaustive mental effort on rote memorization of trivial state details.

4.3 Limitations And Practical Constraints

Despite the robust performance of the Event-Indexed Memory Map architecture, several limitations and practical constraints remain inherent to the computational processing of literature. The primary limitation stems from the current capabilities of natural language processing algorithms when confronted with highly stylized, metaphorical, or ambiguous prose. Serial fiction frequently employs unreliable narrators, dream sequences, and hallucinatory experiences, which the extraction pipeline may misinterpret as literal, state-altering events in the physical fictional world. Furthermore, the system struggles with implicit context and inferred actions. If an author writes that a character departed for the capital, and in the subsequent chapter the character is interacting with the king, human readers effortlessly infer the successful journey and arrival. The automated system, however, requires explicit textual evidence of the arrival event to update the spatial location attribute accurately. In the absence of this explicit event, the auditor may erroneously flag the interaction with the king as a spatial dislocation violation, generating a false positive. Addressing these limitations requires the integration of more sophisticated pragmatic reasoning engines capable of modeling common sense inferences and distinguishing between literal narrative facts and subjective character perceptions. Future iterations of the system must also incorporate mechanisms for authors to manually override or correct the memory map, allowing human intuition to guide the computational framework when textual ambiguity proves insurmountable.

5. Conclusion

The monumental scale and rapid production schedules characteristic of modern serial fiction present formidable cognitive challenges for authors striving to maintain internal consistency. The inevitable degradation of human memory over years of continuous writing invariably leads to continuity errors that compromise the integrity of the fictional universe and disrupt the reader's immersive experience. This research has demonstrated that traditional methods of consistency management, ranging from static story bibles to non-temporal keyword tracking, are fundamentally inadequate for the dynamic, chronologically complex nature of long-form narratives. The introduction of Event-Indexed Memory Maps represents a significant methodological advancement in computational narratology and authorial assistance. By structuring narrative data as a chronological sequence of state-altering events rather than a static repository of facts, this architecture provides a highly accurate, queryable model of the evolving fictional world. Our empirical evaluation confirms that this temporally anchored approach drastically reduces the incidence of false positives during automated auditing, achieving a precision rate that makes it a viable, practical tool for active serial authors. The system effectively identifies spatial dislocations, inventory paradoxes, and attribute mutations by evaluating new text exclusively against the precisely localized state baseline established by all preceding events.

6. Future Directions in Narrative Informatics

The success of Event-Indexed Memory Maps in consistency auditing opens extensive avenues for future research within the broader domain of narrative informatics [38]. Subsequent investigations should focus on enhancing the robustness of the extraction pipeline to better handle implicit actions, metaphorical language, and deceptive narrative devices such as unreliable narrators. Additionally, expanding the scope of the consistency auditor to encompass complex thematic tracking, emotional arc continuity, and dialogue voice consistency represents the next frontier in computational authorial support. Ultimately, the continued development of these intelligent frameworks promises to fundamentally transform the landscape of digital authorship, providing creators with the cognitive scaffolding necessary to construct ever more expansive, intricately detailed, and flawlessly consistent serialized universes.

References

- Li, X., Wang, P., Li, G., Ni, L., & Zhang, Y. (2023). Design of interface circuits and lightweight PUF for TMR sensors. *IEEE Sensors Journal*, 23(11), 11754-11761.
- Ma, Y., Qu, D., & Pyrozhenko, M. (2026). Bio-RegNet: A Meta-Homeostatic Bayesian Neural Network Framework Integrating Treg-Inspired Immunoregulation and Autophagic Optimization for Adaptive Community Detection and Stable Intelligence. *Biomimetics*, 11(1), 48.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*.
- Sang, Y. (2025, July). Towards explainable rag: Interpreting the influence of retrieved passages on generation. In *2025 4th International Conference on Robotics, Artificial Intelligence and Intelligent Control (RAIIC)* (pp. 397-400). IEEE.
- Rafailov, R., Sharma, A., Mitchell, E., Manning, C. D., Ermon, S., & Finn, C. (2024). Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems*.
- Zhang, Jiaquan, et al. "Lightweight llm agent memory with small language models." *arXiv preprint arXiv:2604.07798* (2026).

- Sang, Y. (2025, October). AutoCrit: A Meta-Reasoning Framework for Self-Critique and Iterative Error Correction in LLM Chains-of-Thought. In 2025 6th International Conference on Machine Learning and Computer Application (ICMLCA) (pp. 1177-1180). IEEE.
- Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In Proceedings of EMNLP-IJCNLP.
- Sang, Y. (2025, August). Robustness of fine-tuned llms under noisy retrieval inputs. In 2025 6th International Conference on Artificial Intelligence and Electromechanical Automation (AIEA) (pp. 417-420). IEEE.
- Chu, Y. B., Li, J., Jia, P., Cui, J., Zhang, R., Kang, X., Lv, M., & Zhang, S. (2021). Irf1- and Egr1-activated transcription plays a key role in macrophage polarization: A multiomics sequencing study with partial validation. *International immunopharmacology*, 99, 108072. <https://doi.org/10.1016/j.intimp.2021.108072>
- Deng, L., Deng, S., Chen, Y., Dai, Y., Zhong, Z., Li, L., ... & Huang, H. (2026). Tool-Augmented Agent for Closed-loop Optimization, Simulation, and Modeling Orchestration. *arXiv preprint arXiv:2605.20190*.
- Zhou, D. (2026). AI-Driven Hybrid SAST–DAST–SCA–IAST Framework for Risk-Based Vulnerability Prioritization in Microservice Architectures.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., & Zettlemoyer, L. (2020). BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In Proceedings of ACL.
- Guo, J., Chen, X., Xia, Q., Wang, Z., Ou, J., Qin, L., ... & Tian, W. (2025, July). HASH-RAG: bridging deep hashing with retriever for efficient, fine retrieval and augmented generation. In Findings of the Association for Computational Linguistics: ACL 2025 (pp. 26847-26858).
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. In Proceedings of EMNLP.
- Yang, Y., Yang, X., Jiang, Y., Mu, N., Hu, H., Xie, R., ... & Xu, B. (2026). GlobeDiff: State Diffusion Process for Partial Observability in Multi-Agent Systems. *arXiv preprint arXiv:2602.15776*.
- Zhang, W., & Pei, S. (2026). Predictive prefetching for retrieval-augmented generation. *Proceedings of Machine Learning Research. The Forty-third International Conference on Machine Learning, ICML 2026*. <https://par.nsf.gov/biblio/10686293>
- Wang, Juyuan, et al. "HeadRank: Decoding-Free Passage Reranking via Preference-Aligned Attention Heads." *arXiv preprint arXiv:2604.17237* (2026).
- Chen, L. (2026). Beyond external constraints: The missing dimension of ai governance. Available at SSRN 6449738.
- Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*.
- Xu, X., Tu, W., & Yang, Y. (2023). CASE-Net: Integrating local and non-local attention operations for speech enhancement. *Speech Communication*, 148, 31-39.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., et al. (2020). Language models are few-shot learners. In *Advances in Neural Information Processing Systems*.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., et al. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems*.
- Wang, J., Zhao, R., Wei, W., Wang, Y., Yu, M., Zhou, J., ... & Xu, L. (2026, March). Comorag: A cognitive-inspired memory-organized rag for stateful long narrative reasoning. In

- Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 40, No. 39, pp. 33557-33565).
- Liang, X., Tao, M., Xia, Y., Wang, J., Li, K., Wang, Y., ... & Wang, X. (2025). Sage: Self-evolving agents with reflective and memory-augmented abilities. *Neurocomputing*, 647, 130470.
- Chen, Y., Zhao, C., & Gao, Y. (2025). Blockchain applications in health insurance: a review of applications, challenges, and prospects. *Frontiers in Blockchain*, 8, 1699290.
- Xiao, Q., & Zhou, K. (2026). MACD: Model-Aware Contrastive Decoding via Counterfactual Data. *arXiv preprint arXiv:2602.01740*.
- Deng, L., Guo, H., Bai, Y., Dai, Y., Huang, H., & Shi, Y. (2025). BrepLLM: Native Boundary Representation Understanding with Large Language Models. *arXiv preprint arXiv:2512.16413*.
- Liu, T., Jia, Y., Zeng, S., & Yan, J. (2026). Social media use and attitudes toward premarital sex in China: a chain-mediating analysis. *Sexuality Research and Social Policy*, 1-17.
- Qu, D., & Ma, Y. (2025). Magnet-bn: markov-guided Bayesian neural networks for calibrated long-horizon sequence forecasting and community tracking. *Mathematics*, 13(17), 2740.
- Wang, S., Zhang, X., Wang, P., & Li, X. (2026). Design of Magnetic Sensor Array-Based PUFs for IoT Security. *IEEE Transactions on Instrumentation and Measurement*, 75, 1-9.
- Zhou, J., Shuang, K., Wang, Q., Qian, B., & Guo, J. (2025). Bi-directional feature learning-based approach for zero-shot event argument extraction. *Information Processing & Management*, 62(5), 104199.
- Huang, Z., Wang, J., Chen, L., Xiao, B., Cai, L., Zeng, Y., & Xu, J. (2025, October). MVISU-Bench: Benchmarking Mobile Agents for Real-World Tasks by Multi-App, Vague, Interactive, Single-App and Unethical Instructions. In *Proceedings of the 33rd ACM International Conference on Multimedia* (pp. 8797-8805).
- Wang, Chenxing, et al. "When & How to Write for Personalized Demand-aware Query Rewriting in Video Search." *arXiv preprint arXiv:2602.17667* (2025).
- Xu, X., Wang, Y., Xu, D., Peng, Y., Zhang, C., Jia, J., & Chen, B. (2022, May). Vsean: Visual speech enhancement generative adversarial network. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 7308-7311). IEEE.
- Zhou, D. (2025, December). M-VP2: Microservice-Oriented Vulnerability Patch Planning-A Cost-Aware Approach using Multi-Agent Reinforcement Learning. In *2025 5th International Conference on Computer, Internet of Things and Control Engineering (CITCE)* (pp. 248-254). IEEE.
- Zhang, Jiaquan, et al. "Learning global hypothesis space for enhancing synergistic reasoning chain." *arXiv preprint arXiv:2602.09794* (2026).
- Guo, J., Shuang, K., Li, J., & Wang, Z. (2021, August). Dual slot selector via local reliability verification for dialogue state tracking. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)* (pp. 139-151).