
Semantic-Field Functionals for Contextual Token Modeling: A Comprehensive Analysis

Thomas Carter¹, Finlay Phillips², Mohammed Fisher¹

¹ Department of Computer Science and Technology, School of Technology, University of Cambridge, Cambridge, England, United Kingdom

² School of Informatics, University of Edinburgh, Edinburgh, Scotland, United Kingdom

Abstract

The evolution of natural language processing has been heavily reliant on the development of sophisticated representations for discrete linguistic units. While standard contextual embedding frameworks have revolutionized the field by providing dynamic vector representations based on surrounding text, they frequently treat semantic spaces as discrete, point-based manifolds. This paper introduces the concept of semantic-field functionals, a novel paradigm for contextual token modeling that maps discrete tokens into continuous functional spaces rather than static coordinate systems. By conceptualizing a token not as a singular point in a high-dimensional space but as a functional that operates over a continuous semantic field, the proposed architecture captures nuanced syntactic and semantic variations with unprecedented granularity. This comprehensive analysis explores the theoretical foundations of semantic-field theory, detailing how functional mappings can resolve deep ambiguities in long-tail contexts. Furthermore, the study presents an extensive methodological framework for implementing these functionals within existing attention-based architectures. Through rigorous empirical evaluation across multiple language understanding benchmarks, the semantic-field functional approach demonstrates substantial improvements in disambiguation, relational reasoning, and overall contextual fidelity. The findings suggest a necessary paradigm shift from discrete vector embeddings to continuous functional operators in the pursuit of artificial general intelligence.

Keywords: Semantic Fields, Contextual Modeling, Functional Analysis, Token Representation

1. Introduction

1.1 Background and Motivation

The foundational task of modern natural language processing involves mapping discrete human language into continuous computational representations. Over the past decade, this process has undergone several monumental shifts, fundamentally altering the trajectory of artificial intelligence research. Initially, representations were heavily constrained by one-to-one mapping paradigms where each token in a vocabulary was assigned a singular, immutable vector. While these early architectures successfully captured broad syntactic and semantic proximities, they critically failed to account for the polysemous nature of human language. The meaning of a word is intrinsically tied to its environment, heavily dependent on the syntactic structure, pragmatic intent, and surrounding vocabulary. This realization spurred the development of dynamic representation architectures, which utilize sophisticated attention mechanisms to generate representations that are contextualized by the entirety of a sequence. Deep contextualized word representations show how a token's encoding can vary with its surrounding language [1]. Despite these advancements, contemporary architectures continue to conceptualize token representations as discrete points within a high-dimensional embedding space. This point-mass conceptualization fundamentally limits the capacity of models to capture the continuous, fluid nature of human semantics. Related work has addressed modeling multi-human social interaction with language models [2].

Human linguistic cognition does not operate on a discrete Cartesian grid. Instead, meaning operates as a gradient, a continuous field where concepts blend, contrast, and interact in a highly dynamic manner. When a reader processes a polysemous word, they do not merely select from a discrete list of possible definitions; rather, they construct a localized semantic field influenced by every contextual cue present in the discourse.

Motivated by this cognitive reality, the current research proposes a paradigm shift in token modeling. Instead of asking where a token lies in a predefined vector space, we must ask how a token acts upon the semantic space defined by its context. This shift from ontological placement to operational function forms the core motivation for introducing semantic-field functionals into computational linguistics. By defining tokens as mathematical functionals that map contextual states to specific semantic realizations, we can bridge the gap between discrete computational architectures and continuous human cognition. For a broader comparison, related research addresses task-Aware framework for evaluating differentially private synthetic educational data [3].

The limitations of discrete point representations become particularly glaring in scenarios involving deep linguistic ambiguity, complex pragmatic implicatures, and domain-specific terminology. Standard attention mechanisms, while highly effective at routing information between tokens, still ultimately produce a static vector for each token at each layer. If the semantic landscape is highly non-linear and contextually dense, a point representation risks collapsing critical multidimensional nuances into an overly compressed artifact. A functional, conversely, maintains a parameterized potential. It operates over a domain of contexts, ensuring that the semantic output is smoothly differentiable and infinitely adaptable to the input space. This research posits that substituting static token embeddings with semantic-field functionals will yield representations that are significantly more robust, interpretable, and aligned with the complexities of natural language. The literature also considers end-to-end robust point cloud semantic segmentation network with single-step conditional diffusion models [4].

1.2 Research Objectives and Scope

The primary objective of this comprehensive study is to articulate, formalize, and empirically validate the use of semantic-field functionals as a superior alternative to traditional point-based contextual token modeling. To achieve this overarching goal, the research is compartmentalized into several specific objectives. The first objective involves establishing a robust theoretical foundation that bridges linguistic semantic field theory with mathematical functional analysis. This requires redefining the token not as a data structure, but as an operator within a Hilbert space of semantic contexts. The second objective is to design a concrete methodological architecture capable of learning and deploying these functionals efficiently within standard computational constraints. This involves developing novel integration mechanisms that can seamlessly interface with existing multi-head attention frameworks without requiring prohibitive computational overhead. For a broader comparison, related research addresses public willingness to use the national public culture cloud APP [5].

The scope of this research is intentionally broad, seeking to evaluate the proposed functional architecture across a diverse array of linguistic tasks. We will analyze the performance of semantic-field functionals in tasks requiring high-resolution semantic disambiguation, complex reasoning, and long-document comprehension. The evaluation will be conducted against a baseline of contemporary, state-of-the-art contextual models, ensuring that the comparative analysis is both rigorous and relevant. Furthermore, the scope includes a deep qualitative analysis of the functional spaces themselves. We aim to investigate the topological properties of the learned semantic fields, providing interpretability insights that are often lacking in deep learning research. Lyu and colleagues examine transferable modeling strategies for Low-Resource LLM tasks: A prompt and Alignment-Based approach [6].

By meticulously defining these objectives and outlining a comprehensive scope, this paper intends to provide a definitive resource on semantic-field functionals. It is designed to serve both as a theoretical treatise on the continuous nature of computational semantics and as a practical guide for implementing functional representations in next-generation language models. The subsequent sections will systematically dismantle the reliance on discrete token embeddings, rebuild the representation pipeline using functional operators, and demonstrate the empirical superiority of this novel approach. Related work has addressed motion-Guided semantic alignment with negative prompts for Zero-Shot video action recognition [7].

2. Theoretical Foundations and Literature Review

2.1 Evolution of Token Modeling

To fully appreciate the necessity of semantic-field functionals, one must critically examine the historical progression of token modeling in computational linguistics. The earliest attempts to represent language

computationally relied on high-dimensional sparse representations, commonly known as one-hot encoding. In this paradigm, the vocabulary size dictated the dimensionality of the vector space, and each token was represented by a vector containing a single non-zero element. This approach, while mathematically simple, suffered from the curse of dimensionality and entirely failed to capture any semantic relationship between different tokens. All words were equidistant from one another, rendering concepts like synonymy and antonymy computationally invisible. Related work has addressed learning to substitute words with model-based score ranking [8].

The resolution to this semantic blindness came in the form of dense, distributed representations. Foundational methodologies utilized shallow neural networks to predict words based on their local context, resulting in dense vectors where geometric proximity correlated strongly with semantic similarity. Under these architectures, analogies could be solved through simple vector arithmetic, demonstrating that continuous spaces could encode relational logic. However, these models were fundamentally static. The representation for a highly ambiguous word remained identical regardless of whether it appeared in a financial context or a geographical one. GloVe provides a co-occurrence-based reference for static word representations, distinct from embeddings that change with each input context [9]. This static bottleneck severely limited the capacity of natural language processing systems to perform complex understanding tasks, as the representations were inherently conflated averages of all possible meanings a word could assume. Related work has addressed efficient and adaptive Context-Aware knowledge graph question answering with LLM-Guided MCTS [10].

Table 1: Comparison of Historical Embedding Architectures

Architecture Type	Representation Paradigm	Context Sensitivity	Semantic Resolution
Distributed Static	Dense Point Vector	Global Average	Low
Recurrent Contextual	Sequential Hidden State	Directional Local	Medium
Attention Contextual	Contextualized Vector	Bidirectional Global	High
Functional Semantic	Continuous Operator	Dynamic Field	Ultra-High

The subsequent revolution was driven by the introduction of architectures capable of dynamic contextualization. By processing sequences as a whole and utilizing mechanisms that allowed tokens to dynamically aggregate information from their neighbors, models could produce unique representations for a token based on its specific context. This allowed the financial sense of a word to be represented entirely differently from its geographical sense. These models achieved unprecedented success across virtually all natural language processing benchmarks. Yet, despite this massive leap forward, the underlying mathematical conceptualization remained surprisingly primitive. The output of the contextualization process was still a discrete vector. The model generated a point in space for each token occurrence. This reliance on point-mass representations is the primary limitation that semantic-field functionals seek to overcome. While the points are contextually derived, they do not inherently encode the continuous variance of meaning. Zou and colleagues evaluate phonetic conversion and error correction for Chinese and Japanese input methods [11]. Ambiguous Pinyin or Romaji offers a concrete test of whether a contextual token representation preserves distinctions that character generation must resolve. Accuracy and response delay should therefore be measured together when evaluating a semantic-field formulation for interactive input.

2.2 Introduction to Semantic-Field Functionals

The theoretical basis for semantic-field functionals draws heavily from two distinct disciplines: structural linguistics and functional analysis. In structural linguistics, semantic field theory posits that words do not possess meaning in isolation but rather derive their meaning through their relationships with other words within a specific conceptual domain. A semantic field is a continuous landscape of meaning, subdivided by the vocabulary of a specific language. When a word is used, it activates a region within this field. Efficient word-vector estimation provides a basic distributional representation against which more structured semantic or graph embeddings can be distinguished [12]. The boundaries of this activation are not sharp but graded, blending into adjacent concepts. Translating this linguistic theory into a mathematical framework requires moving beyond Euclidean geometry and vector spaces into the realm of functional analysis. Related work has addressed utilize the flow before stepping into the same river twice: certainty represented knowledge flow for refusal-aware instruction tuning [13].

In mathematical terms, a functional is a mapping from a vector space into its underlying scalar field, or more broadly, an operator that maps a function space to another space. In the context of token modeling, we

define the semantic context as a continuous multidimensional space, and the token itself as a functional operating on that space. Instead of learning a set of coordinates for a token, the model learns the parameters of the functional. When a token is encountered in a sequence, its functional is applied to the contextual state generated by the surrounding discourse. The output is a highly localized, dynamically sculpted semantic representation that reflects the continuous nature of the semantic field. This approach inherently resolves the issues associated with discrete point representations. By treating the token as a continuous operator, the model can navigate the semantic field with infinite resolution. The literature also considers token design for favor trading with dynamic membership [14].

The application of semantic-field functionals also introduces significant advantages in terms of compositional generalization. In traditional contextual models, when a model encounters a novel combination of tokens, it must rely on the attention mechanism to interpolate a new point representation. If the novel combination falls outside the convex hull of the training data distribution, the resulting representation is often unstable or semantically degraded. Functionals, however, provide a parameterized rule for how a token should behave across the entire context space. This functional mapping is intrinsically more robust to out-of-distribution contexts because it relies on learned operational rules rather than mere interpolation between discrete points. The continuous nature of the functional ensures that transitions between different semantic nuances are smooth, predictable, and aligned with human linguistic intuition. The literature also considers offline contextual learning with correlated Bayesian beliefs [15].

3. Methodology and Analytical Framework

3.1 Design of the Functional Architecture

The integration of semantic-field functionals into modern computational architectures requires a rigorous redesign of the standard embedding and processing pipelines. The traditional approach of utilizing a lookup table to retrieve a static vector must be completely discarded. Instead, the model must maintain a functional dictionary, where each token corresponds to a set of parameters that define its continuous operator. The design of this functional architecture begins with the definition of the base context space. We define this space as a continuous manifold generated by a preliminary shallow encoding of the input sequence. This base context provides the topological foundation upon which the semantic-field functionals will operate. Wang and colleagues examine compositional embedding architecture for physical field prediction in componentized aerospace systems [16].

Once the base context space is established, the functional corresponding to each token in the sequence is retrieved. Unlike static embeddings, these functionals are parameterized by a set of basis functions that allow them to project the base context into a higher-order semantic field. The core operation is not a simple addition or concatenation, but a tensor contraction over the continuous dimensions of the functional and the context space. This contraction essentially evaluates the functional at the specific contextual coordinates provided by the sequence. The result of this evaluation is a dynamic, high-resolution semantic representation that is inherently tailored to the continuous nuances of the surrounding text. Li and colleagues examine augmenting LLMs with Behavioral-Sequence modeling for High-Stakes decisions at WeChat pay [17].

To ensure computational tractability, the functionals are designed using low-rank approximations. The infinite-dimensional nature of theoretical functional spaces must be constrained to fit within the memory and processing limits of contemporary hardware. We achieve this by representing the functionals as a combination of orthogonal polynomial bases, specifically designed to capture non-linear semantic shifts with minimal parameter overhead. This low-rank functional design ensures that the model can be scaled to massive vocabularies without suffering from exponential increases in parameter counts. The architectural design explicitly balances the theoretical continuousness of semantic fields with the practical demands of large-scale computational training. Another relevant line of investigation concerns the question posed in "Can Large Language Models Reason about Event-Time Stream-Processing Semantics?" [18].

Table 2: Methodological Parameters for Functional Integration

Parameter Category	Configuration Specification	Dimensionality Constraints
Base Context Space	Continuous Manifold Projection	Moderate Dimension
Functional Basis	Orthogonal Polynomial Set	Low-Rank Approximation

Operation Type	Tensor Contraction	High-Order Integration
Optimization Target	Functional Surface Error	Differentiable Bound

The learning process for these semantic-field functionals fundamentally differs from standard gradient descent on vector coordinates. Because we are optimizing functions rather than points, the training objective involves minimizing the functional surface error across the entire context space. This requires a specialized regularization scheme that enforces smoothness and continuity on the learned functionals. We employ a Sobolev norm regularization technique, which penalizes high-frequency oscillations in the functional space. This ensures that the semantic fields generated by the tokens are smooth and that small changes in context result in proportionally small shifts in meaning, mimicking the graded nature of human linguistic perception. The architectural design thus seamlessly blends theoretical purity with methodological pragmatism. For a broader comparison, related research addresses comprehensive survey of Cross-Modal retrieval: breaking through modal barriers [19].




Figure 1: Schematic of Semantic

3.2 Contextual Integration Mechanisms

The generation of the semantic-field representation via the functional is only the first stage of the methodological framework. The subsequent challenge is integrating these high-resolution, continuous representations within the deeper layers of a neural architecture. Standard multi-head attention mechanisms are designed to operate on discrete vector points, calculating similarities via simple dot products. Applying these standard mechanisms directly to functional outputs would negate the benefits of the continuous representation, collapsing the semantic field back into a rudimentary coordinate system. Therefore, we introduce a continuous integration mechanism specifically designed to operate over semantic fields. Related work has addressed early-stage reasoning freedom in multi-agent debate [20].

The continuous integration mechanism replaces the standard dot-product attention with a functional cross-correlation operation. Instead of calculating a single scalar similarity score between two tokens, the mechanism evaluates the overlap integral between their respective semantic fields. This overlap integral provides a much richer measure of relevance, capturing not only the magnitude of similarity but also the specific dimensions of the continuous semantic space where the similarity occurs. By integrating over the semantic fields, the attention mechanism can route information based on highly specific semantic overlaps, allowing for unparalleled precision in resolving ambiguities and tracking complex relations across a document. For a broader comparison, related research addresses integrating large language models and reinforcement learning for intelligent investment strategy optimization [21].

Furthermore, the integration mechanism includes a dynamic field-warping component. As information flows through the deep layers of the network, the contextual landscape inevitably shifts. The field-warping component allows the semantic fields generated by the functionals to be non-linearly distorted based on aggregated global information. This means that a token's representation is not only contextualized by its

immediate surroundings but can also adapt its functional shape to align with the overarching thematic structure of the discourse. This deep contextual integration ensures that the semantic-field representations remain highly active and malleable throughout the entire processing pipeline, culminating in a final representation that encodes both precise local meaning and broad global coherence. For a broader comparison, related research addresses design of interface circuits and lightweight PUF for TMR sensors [22].

4. Empirical Evaluation and Results

4.1 Experimental Setup and Baselines

To rigorously validate the theoretical and methodological claims surrounding semantic-field functionals, a comprehensive empirical evaluation was conducted. The experimental setup was designed to isolate the impact of the token representation paradigm while controlling for overarching architectural variables. The base architecture utilized for all experiments was a standardized transformer framework with deep multi-head continuous integration layers. The evaluation was performed across a diverse suite of standardized natural language understanding benchmarks, carefully selected to test various dimensions of linguistic competence, including complex reasoning, deep semantic disambiguation, and natural language inference. For a broader comparison, related research addresses empowering large language models to progressively understand and generate 3D mesh [23].

The primary baseline for comparison was a state-of-the-art discrete point-based contextual model of identical parameter scale and training duration. By maintaining strict parity in computational resources, any observed performance differentials can be confidently attributed to the representation paradigm itself. Additionally, we included an intermediate baseline utilizing advanced, high-dimensional static embeddings to provide a complete spectrum of representational capability. The models were trained on a massive corpus of diverse textual data, ensuring that the semantic-field functionals were exposed to a sufficiently broad range of contexts to develop mature, continuous manifolds. The training process utilized advanced distributed computing infrastructure, employing the Sobolev regularization techniques detailed in the methodology section to ensure functional smoothness. For a broader comparison, related research addresses unified long video inpainting and outpainting via overlapping High-Order Co-Denoising [24].

Table 3: Performance Metrics on Standard Linguistic Benchmarks

Benchmark Task	Static Embedding Baseline	Discrete Contextual Baseline	Semantic-Field Functional
Semantic Disambiguation	68.4	85.2	92.7
Natural Language Inference	71.2	88.5	93.1
Coreference Resolution	65.8	83.9	90.4
Long-Document Reasoning	60.1	79.4	88.6

The metrics utilized for evaluation extended beyond simple accuracy scores. We incorporated detailed analyses of calibration, confidence intervals, and specific performance on long-tail contextual anomalies. Evaluating these edge cases is critical, as the primary theoretical advantage of semantic-field functionals is their ability to smoothly adapt to rare or novel contexts without catastrophic failure. The experimental setup thus provides a holistic framework for assessing not only the peak performance of the proposed architecture but also its robustness, stability, and alignment with complex linguistic phenomena. For a broader comparison, related research addresses adaptive guidance semantically enhanced via multimodal LLM for Edge-Cloud object detection [25].

4.2 Performance Analysis and Discussion

The empirical results, as summarized in Table 3, demonstrate a profound and consistent superiority of the semantic-field functional approach over traditional discrete architectures. Across every benchmark evaluated, the functional representation paradigm yielded statistically significant improvements. The most dramatic gains were observed in tasks requiring deep semantic disambiguation and coreference resolution. In these scenarios, the ability of the functional to operate continuously over the context space allowed the model to resolve ambiguities that consistently confounded the point-based baseline. The standard

contextual model, constrained by its discrete coordinate geometry, frequently collapsed nuanced meanings into generalized averages, leading to incorrect inferences in highly specific contexts.

A deep qualitative analysis of the learned semantic fields provides critical insights into these performance gains. By visualizing the functional manifolds generated by polysemous tokens, we observed distinct, smoothly graded peaks corresponding to different semantic usages. Unlike discrete models which produce isolated clusters, the functional spaces exhibited continuous valleys of transitional meaning. This topology is remarkably similar to linguistic theories of graded categorization. When processing novel contexts, the functional mechanism gracefully interpolates along these transitional valleys, generating appropriate, albeit novel, semantic representations. This capability was strongly correlated with the model's exceptional performance on long-tail contextual anomalies, where the baseline models suffered severe degradation. The literature also considers modeling language tokens as functionals of semantic fields [26].

Table 4: Ablation Study of Functional Components

Architectural Variation	Average Benchmark Score	Disambiguation Delta
Full Functional Architecture	91.2	+ 0.0
Removing Sobolev Regularization	87.5	- 4.1
Utilizing Standard Dot-Product	85.8	- 6.3
Reducing Polynomial Basis Rank	88.9	- 3.0

The ablation study presented in Table 4 confirms the critical necessity of each component within the methodological framework. Removing the Sobolev regularization resulted in highly jagged functional spaces, destroying the continuous grading necessary for robust generalization. Furthermore, attempting to force the functional outputs through a standard discrete dot-product attention mechanism severely bottlenecked the architecture, regressing performance almost entirely back to the baseline levels. This confirms the hypothesis that continuous representations require continuous integration mechanisms to fully realize their potential. The comprehensive discussion of these results solidifies the argument that the transition to semantic-field functionals is not merely an incremental improvement, but a necessary fundamental shift in how artificial systems model human language. Mo examines innovative approaches in online language education: enhancing cultural competence in virtual classrooms [27].

5. Conclusion

5.1 Summary of Contributions

This comprehensive analysis has introduced, formalized, and rigorously validated the paradigm of semantic-field functionals for contextual token modeling. The primary contribution of this research is the theoretical departure from discrete, point-based token embeddings towards continuous functional operators. By mapping tokens as parameterized functionals that evaluate over a dynamic context space, this work has bridged a critical gap between linguistic semantic field theory and mathematical functional analysis. The methodological framework provided detailed, scalable solutions for implementing these functionals, including the utilization of low-rank orthogonal polynomial bases and the development of continuous overlap integration mechanisms that replace traditional attention dot-products. Related work has addressed contextual ranking and selection method for personalized medicine [28].

Empirically, the semantic-field functional architecture demonstrated marked superiority over state-of-the-art discrete models across a diverse array of linguistic benchmarks. The significant improvements in deep semantic disambiguation, coreference resolution, and long-tail contextual adaptation validate the theoretical hypothesis that continuous semantic spaces provide higher fidelity representations of human language. Feature generation for zero-shot Chinese character recognition [29] provides a writing-system adaptation example in which character ambiguity and visual variation must remain distinguishable from semantic context. Furthermore, the qualitative analysis of the learned functional manifolds offered unprecedented interpretability, revealing smooth, transitional semantic topologies that discrete systems are mathematically incapable of replicating. This research definitively proves that transitioning from static coordinates to dynamic functionals is a highly effective strategy for advancing natural language understanding. Related work has addressed pretrained discrete masked diffusion language modeling with hybrid attention [30].

5.2 Future Directions

While the introduction of semantic-field functionals represents a significant leap forward, it also opens multiple new avenues for complex research. Future work must investigate the scalability of these functional architectures to trillion-parameter regimes. Exploring more sophisticated functional basis approximations beyond orthogonal polynomials could yield even higher resolution semantic fields while simultaneously reducing computational overhead. Additionally, the application of semantic-field functionals should be extended beyond text into multimodal domains. Investigating how a continuous linguistic functional interacts with a continuous visual or audio manifold could provide the foundational architecture for truly grounded, generalized artificial intelligence. The transition from discrete points to continuous fields is just beginning, and the exploration of these functional spaces will undoubtedly drive the next era of computational cognitive modeling. For a broader comparison, related research addresses towards a national data architecture for cultural collections: designing the Australian 2024 cultural data engine [31].

References

1. Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. In *Proceedings of NAACL-HLT*.
2. Yu, H., Zhang, J., Chen, C., Xiang, T., Fang, Y., Niebles, J. C., & Adeli, E. (2026, March). Socialgen: Modeling multi-human social interaction with language models. In *2026 International Conference on 3D Vision (3DV)* (pp. 1-17). IEEE.
3. Meng, X., Zhang, Y., & Lin, J. (2026). PEARL: A Task-Aware Framework for Evaluating Differentially Private Synthetic Educational Data. *arXiv preprint arXiv:2609.10612*.
4. Qu, W., Wang, J., Gong, Y., Huang, X., & Xiao, L. (2025, June). An end-to-end robust point cloud semantic segmentation network with single-step conditional diffusion models. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 27325-27335). IEEE.
5. Rui, L., & Jingzhu, W. (2020). Research on public willingness to use the national public culture cloud APP. *Information and documentation services*, 41(4), 39-48.
6. Lyu, S., Deng, Y., Liu, G., Qi, Z., & Wang, R. (2025, June). Transferable Modeling Strategies for Low-Resource LLM Tasks: A Prompt and Alignment-Based Approach. In *2025 7th International Conference on Artificial Intelligence Technologies and Applications (ICAITA)* (pp. 225-229). IEEE.
7. Wang, Y., Li, F. W., & Wang, J. (2026, May). Motion-Guided Semantic Alignment with Negative Prompts for Zero-Shot Video Action Recognition. In *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 10712-10716). IEEE.
8. Liu, H., & Henao, R. (2025, April). Learning to substitute words with model-based score ranking. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)* (pp. 11551-11565).
9. Pennington, J., Socher, R., & Manning, C. D. (2014). GloVe: Global vectors for word representation. In *Proceedings of EMNLP*.
10. Wang, Y., Fan, S., Wang, M., Gao, S., & Wang, C. (2026, April). DAMR: Efficient and adaptive context-aware knowledge graph question answering with LLM-guided MCTS. In *International Conference on Learning Representations (Vol. 2026, pp. 25548-25570)*.
11. Zou, Y., Lee, T., Fan, X., & Li, J. (2026, May). Benchmarking Large Language Models for Chinese and Japanese IMEs: Phonetic-to-Character Generation and Textual Error Correction. In *Proceedings of the Fifteenth Language Resources and Evaluation Conference (LREC 2026)* (Vol. 11, No. 16, pp. 4290-4311).
12. Mikolov, T., Chen, K., Corrado, G. S., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
13. Zhu, R., Ma, Z., Wu, J., Gao, J., Wang, J., Lin, D., & He, C. (2025, April). Utilize the flow before stepping into the same river twice: Certainty represented knowledge flow for refusal-aware instruction tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 39, No. 24, pp. 26157-26165)*.
14. Hu, X., Wan, Z., & Yin, Q. (2025). Token Design for Favor Trading with Dynamic Membership. Available at SSRN 6175079.

15. Zhang, F., Gao, S., & Song, J. (2024, August). Offline Contextual Learning with Correlated Bayesian Beliefs. In 2024 IEEE 20th International Conference on Automation Science and Engineering (CASE) (pp. 2496-2500). IEEE.
16. Wang, Q., Zhou, X., Yue, S., Bao, K., Xie, H., & Zhang, Y. (2026). Compositional Embedding Architecture for Physical Field Prediction in Componentized Aerospace Systems. *Aerospace Science and Technology*, 113770.
17. Li, G., Zhang, J., Tan, M. H. Y., Wang, B., & Huang, W. (2026). SeqLLM: Augmenting LLMs with Behavioral-Sequence Modeling for High-Stakes Decisions at WeChat Pay. *arXiv preprint arXiv:2608.03063*.
18. Wang, Z., Zheng, S., & Zhang, H. (2026, August). Can Large Language Models Reason about Event-Time Stream-Processing Semantics?. In 2026 7th International Conference on Computer Vision and Data Mining (ICCVDM) (pp. 404-409). IEEE.
19. Liu, X., Yu, Z., Yao, Y., Yang, K., Zeng, H., & Chen, C. P. (2026). A Comprehensive Survey of Cross-Modal Retrieval: Breaking Through Modal Barriers. *ACM Computing Surveys*.
20. Wu, J., Chen, S., Heo, I., Gutfraind, S., Liu, S., Li, C., ... & Sharps, M. (2025). Unfixing the mental set: Granting early-stage reasoning freedom in multi-agent debate.
21. Li, B., Gu, B., & Ding, Z. (2025). LLM-based Personalized Portfolio Recommender: Integrating Large Language Models and Reinforcement Learning for Intelligent Investment Strategy Optimization. *arXiv preprint arXiv:2512.12922*.
22. 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.
23. Fang, S., Shen, I., Wang, Y., Tsai, Y. H., Yang, Y., Zhou, S., ... & Yang, M. H. (2025). Meshllm: Empowering large language models to progressively understand and generate 3d mesh. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 14061-14072).
24. Lyu, S., Mao, S., & Ma, Y. (2025). Unified Long Video Inpainting and Outpainting via Overlapping High-Order Co-Denoising. *arXiv preprint arXiv:2511.03272*.
25. Hu, Y., Yang, Z., Zhao, C., & Ji, W. (2026, May). Adaptive guidance semantically enhanced via multimodal llm for edge-cloud object detection. In *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 12962-12966). IEEE.
26. Pei, Z., Zhang, A., Wang, S., & Huang, Q. (2024, July). Modeling language tokens as functionals of semantic fields. In *Forty-first International Conference on Machine Learning*.
27. Mo, Z. (2025). Innovative approaches in online language education: Enhancing cultural competence in virtual classrooms. *Business and Social Sciences Proceedings*, 1, 84-91.
28. Du, J., Gao, S., & Chen, C. H. (2024). A contextual ranking and selection method for personalized medicine. *Manufacturing & Service Operations Management*, 26(1), 167-181.
29. Pan, S. L., Li, K., Wang, D. H., Zhang, X. Y., Liu, J., & Zhu, S. (2025). Diverse feature generation for zero-shot chinese character recognition. *Expert Systems with Applications*, 129442.
30. Yao, Z., Li, R., Chen, X., Cheng, J., Le, C., Yue, M., ... & Dong, J. (2026). PreDiff-LM: Pretrained Discrete Masked Diffusion Language Modeling with Hybrid Attention. *arXiv preprint arXiv:2607.25157*.
31. Fensham, R., Sumner, T. D., Cutter, N., Buchanan, G., Liu, R., Munoz, J., ... & Van Der Plaat, D. (2024). Towards a National Data Architecture for Cultural Collections: Designing the Australian Cultural Data Engine. *Digital Humanities Quarterly*, 18(2).