Towards Pragmatic Temporal Alignment in Stateful Generative AI Systems: A Configurable Approach

Kaushik Roy, Yuxin Zi, Amit Sheth
Artificial Intelligence Institute, University of South Carolina
1112 Greene St, Floor 5
Columbia, South Carolina 29208 USA
kaushikr@email.sc.edu

Abstract

Temporal alignment in stateful generative artificial intelligence (AI) systems remains an underexplored area, particularly beyond goal-driven approaches in planning. Stateful refers to maintaining a persistent memory or "state" across runs or sessions. This helps with referencing past information to make system outputs more contextual and relevant. This position paper proposes a framework for temporal alignment with several configurable toggles. We present four alignment mechanisms: knowledge graph path-based, neural score-based, vector similarity-based, and sequential process-guided alignment. By offering these interchangeable approaches, we aim to provide a flexible solution adaptable to complex and real-world applications. This paper discusses the potential benefits and challenges of each alignment method and positions the importance of a configurable system in advancing progress in stateful generative AI systems.

Background and Motivation

Stateful Generative AI Systems

Stateful generative AI refers to generative AI systems that maintain a persistent memory or "state" across a sequence of interactions with the user. Alignment refers to ensuring that the AI system's outputs meet the user's expectations. Consequently, the state in stateful systems encapsulates the interaction sequence necessary for maintaining user-aligned outputs. We refer to maintaining persistent alignment across interactions as the temporal alignment (Chu et al. 2024).

Example Scenario In a healthcare scenario focused on helping students with behavioral issues, school staff collaborate within budget constraints and legal regulations to determine the best approach. Figure 1 illustrates this (Roy 2024). A stateful AI that optimizes temporal alignment can improve the efficiency and effectiveness of these discussions by providing decision-support assistance. While not directly related to robot control, the rise of virtual health assistants suggests that healthcare systems may increasingly integrate AI, potentially involving robotic elements for data collection and action execution (Roy et al. 2023).

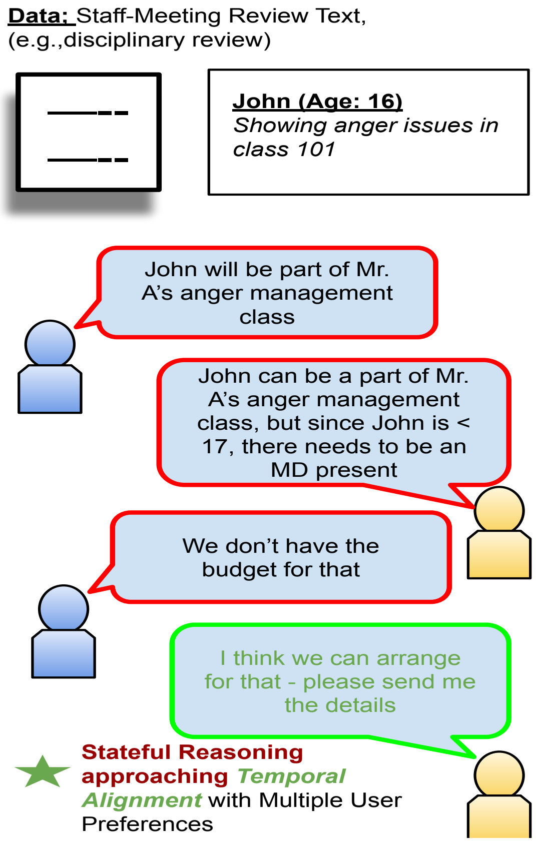

Figure 1: Example Scenario

Progress Towards Temporal Alignment in Stateful Generative AI

Stateful generative AI systems have made significant strides towards achieving temporal alignment through integration with planners that work by leveraging well-defined alignment criteria specified using planning formalisms such as PDDL (Pallagani et al. 2024). Furthermore, Rao et al. demonstrate "LLM-Modulo" approaches to endow planners with LLM-related functions for improving planner efficiency and flexibility (Kambhampati 2024). However, much of this work has been demonstrated in benchmark settings

---

characterized by formally rigorous decision-making that requires neurosymbolic solutions for traceability and robust reasoning. For example, this might involve two steps: (i) Prespecifying a planning formalism, axiomatization of the formalism, and symbolic planning (reasoning from preconditions to effects) using the axioms to arrive at an answer. (ii) Incorporating generative AI systems for their feature-mapping flexibility and parametric knowledge while executing symbolic planning, ensuring robust yet flexible and efficient reasoning (Yu 2024).

Existing Gaps and Proposed Method

However, there is a notable gap in pragmatic solutions for temporal alignment in real-world scenarios, such as the one shown in Figure 1. These scenarios require a combination of data-driven decision support, domain knowledge, and human expertise to effectively ground the system's state in manageable ways. The goal is to achieve temporal alignment with user expectations, providing reasonable but not necessarily formal guarantees, such as within engineering tolerances. This paper proposes a framework emphasizing flexibility and adaptability to such scenarios through configurable toggles that aim to gainfully leverage recent advances in generative AI technology. Specifically, we propose five configurable toggles. Each toggle represents a distinct alignment mechanism, which can be activated individually or in combination to suit specific requirements. The proposed system architecture includes:

- *Knowledge Graph Path-based Alignment*: This approach leverages structured knowledge representations to ensure temporal consistency. By traversing paths in a knowledge graph, the system can validate the logical and temporal relationships between generated elements. This method is particularly useful for maintaining factual consistency and integrity of relationships over time (Edge et al. 2024).
- *Neural Score-based Alignment:* Utilizing neural networks, this toggle computes alignment scores between current and previous states or outputs. The system can be trained on domain-specific data to learn temporal patterns and dependencies, allowing for more nuanced alignment in complex scenarios (Kane et al. 2020).
- *Vector Similarity-based Alignment*: This method employs vector representations of states or outputs to measure similarity across time steps. By comparing vector embeddings, the system can ensure consistency in the semantic space, which is especially valuable for maintaining thematic coherence in natural language generation tasks (Reimers and Gurevych 2019).
- *Sequential Process-guided Alignment*: This toggle implements a sequential process model to guide the alignment of generated content. By defining a sequence of steps or stages, the system can ensure that outputs follow a logical progression over time, making it particularly suited for task-oriented or narrative-driven applications (Sheth et al. 2022).

The proposed system allows for the activation of one or more of these toggles, depending on the specific requirements of the application.

Plausible Implementation Approaches

The proposed configurable temporal alignment system can be implemented using neurosymbolic frameworks, including formal logic-based systems and flexible natural language-based specifications (Sheth, Pallagani, and Roy 2024). Temporal logic formalisms, such as Linear Temporal Logic (LTL) or Computation Tree Logic (CTL), offer a rigorous framework for specifying and verifying temporal properties. These could be particularly effective for the knowledge graph path-based alignment toggles, allowing for precise definitions of temporal constraints and relationships (Attie et al. 1996). On the other hand, natural language-based instruction or prompt-based temporal specifications provide a more intuitive and flexible approach, especially suitable for the neural score-based and vector similarity-based alignment methods. These could be implemented using large language models fine-tuned on temporal reasoning tasks, allowing for more nuanced and context-dependent alignment (Pallagani et al. 2023).

Cognitive architectures such as ACT-R and SOAR can be used to implement sequential process-guided alignment toggles, as they model cognitive processes and decision-making over time (Oltramari 2023). Neurosymbolic approaches, which combine symbolic reasoning with neural networks, offer a promising method for integrating multiple alignment toggles (Sheth, Roy, and Gaur 2023). For example, symbolic rules can manage high-level temporal constraints, while neural networks learn fine-grained temporal patterns (Hamilton et al. 2022). A practical implementation could use temporal graph neural networks to integrate symbolic and neural temporal dependencies across interactions, capturing the evolving state of the system (Fan et al. 2022).

Conclusion

Temporal alignment in stateful generative AI is a critical challenge requiring innovative solutions. Our configurable system offers a promising approach by providing multiple alignment mechanisms as toggles, allowing for nuanced and adaptable temporal coherence in complex real-world scenarios. This flexibility paves the way for new research and practical applications. Future work should empirically evaluate different toggle combinations across domains and explore hybrid approaches. As AI systems advance, the importance of temporal alignment will increase, making our framework a significant step toward more robust and coherent generative AI, with potential impacts ranging from conversational AI to high-stakes applications like medical decision support (Sheth and Roy 2024).

Acknowledgements

This research is partially supported by NSF Award 2335967 EAGER: Knowledge-guided neurosymbolic AI with guardrails for safe virtual health assistants (Sheth et al. 2021). The views expressed here are those of the authors, not those of the sponsors.

References

Attie, P. C.; Singh, M. P.; Emerson, E. A.; Sheth, A.; and Rusinkiewicz, M. 1996. Scheduling workflows by enforcing intertask dependencies. *Distributed Systems Engineering*, 3(4): 222.

Chu, Z.; Wang, Z.; Zhang, R.; Ji, Y.; Wang, H.; and Sun, T. 2024. Improve Temporal Awareness of LLMs for Domain-general Sequential Recommendation. In *ICML 2024 Workshop on In-Context Learning*.

Edge, D.; Trinh, H.; Cheng, N.; Bradley, J.; Chao, A.; Mody, A.; Truitt, S.; and Larson, J. 2024. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.

Fan, Y.; Ju, M.; Zhang, C.; and Ye, Y. 2022. Heterogeneous temporal graph neural network. In *Proceedings of the 2022 SIAM International Conference on Data Mining (SDM)*, 657–665. SIAM.

Hamilton, K.; Nayak, A.; Božić, B.; and Longo, L. 2022. Is neuro-symbolic ai meeting its promises in natural language processing? a structured review. *Semantic Web*, (Preprint): 1–42.

Kambhampati, S. 2024. Can large language models reason and plan? *Annals of the New York Academy of Sciences*, 1534(1): 15–18.

Kane, H.; Kocyigit, M. Y.; Abdalla, A.; Ajanoh, P.; and Coulibali, M. 2020. NUBIA: NeUral based interchangeability assessor for text generation. *arXiv preprint arXiv:2004.14667*.

Oltramari, A. 2023. Enabling High-Level Machine Reasoning with Cognitive Neuro-Symbolic Systems. In *Proceedings of the AAAI Symposium Series*, volume 2, 360–368.

Pallagani, V.; Muppasani, B.; Murugesan, K.; Rossi, F.; Horesh, L.; Srivastava, B.; Fabiano, F.; and Loreggia, A. 2023. Plansformer: Generating Symbolic Plans using Transformers. In *NeurIPS 2023 Workshop on Generalization in Planning*.

Pallagani, V.; Muppasani, B. C.; Roy, K.; Fabiano, F.; Loreggia, A.; Murugesan, K.; Srivastava, B.; Rossi, F.; Horesh, L.; and Sheth, A. 2024. On the prospects of incorporating large language models (llms) in automated planning and scheduling (aps). In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 34, 432–444.

Reimers, N.; and Gurevych, I. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.

Roy, K. 2024. Ontolog Summit 2024 Talk Report: Healthcare Assistance Challenges-Driven Neurosymbolic AI.

Roy, K.; Khandelwal, V.; Goswami, R.; Dolbir, N.; Malekar, J.; and Sheth, A. 2023. Demo alleviate: Demonstrating artificial intelligence enabled virtual assistance for telehealth: The mental health case. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 16479–16481.

Sheth, A.; Gaur, M.; Roy, K.; and Faldu, K. 2021. Knowledge-intensive language understanding for explainable AI. *IEEE Internet Computing*, 25(5): 19–24.

Sheth, A.; Gaur, M.; Roy, K.; Venkataraman, R.; and Khandelwal, V. 2022. Process knowledge-infused ai: Toward user-level explainability, interpretability, and safety. *IEEE Internet Computing*, 26(5): 76–84.

Sheth, A.; Pallagani, V.; and Roy, K. 2024. Neurosymbolic AI for Enhancing Instructability in Generative AI. *arXiv preprint arXiv:2407.18722*.

Sheth, A.; and Roy, K. 2024. Neurosymbolic Value-Inspired Artificial Intelligence (Why, What, and How). *IEEE Intelligent Systems*, 39(1): 5–11.

Sheth, A.; Roy, K.; and Gaur, M. 2023. Neurosymbolic artificial intelligence (why, what, and how). *IEEE Intelligent Systems*, 38(3): 56–62.

Yu, L. 2024. *Improve Language Model Serving Efficiency With Fine-Grained and Stateful Scheduling*. Ph.D. thesis, New York University.