

CADA: A Contextual Adaptive Dialogue Agent Integrating Dynamic Feedback for Enhanced Conversational AI

Herbert Wanga¹

¹Department of Mathematics and Information Technology University of Iringa, Tanzania

Publication Date: 2026/01/08

Abstract: Conversational AI models have revolutionized human-computer interaction, yet challenges persist in achieving seamless, context-aware, and adaptive dialogues. This paper proposes and evaluates a novel hybrid framework designed to bridge two critical gaps: limited contextual awareness and inadequate real-time user feedback integration. The framework synthesizes multimodal contextual analysis with a dynamic, reinforcement learning-based feedback loop. I present a methodological implementation using a modified Transformer architecture augmented with a contextual memory module and a reward model trained on human preferences. Evaluation on a custom dataset simulating educational and customer service dialogues shows a 28% improvement in response appropriateness and a 32% increase in user satisfaction scores compared to a baseline GPT-3.5-turbo fine-tuned model. Key findings highlight the importance of real-time adaptation and transparent feedback mechanisms in fostering trust. The paper concludes with a critical discussion on ethical implications, specifically bias amplification in feedback loops, and provides recommendations for future research in scalability and cross-cultural generalization.

Keywords: *Conversational AI, Contextual Understanding, User Feedback, Reinforcement Learning from Human Feedback (RLHF), Adaptive Learning, Ethical AI.*

How to Cite: Herbert Wanga (2025) CADA: A Contextual Adaptive Dialogue Agent Integrating Dynamic Feedback for Enhanced Conversational AI. *International Journal of Innovative Science and Research Technology*, 10(12), 2727-2730.
<https://doi.org/10.38124/ijisrt/25dec1454>

I. INTRODUCTION

Conversational AI models, such as chatbots and virtual assistants, have become ubiquitous in domains like education, healthcare, and customer service [15], [5]. Despite their widespread adoption, these systems often struggle with contextual ambiguity, misinterpretation of user intent, and static response generation [13]. Two critical, interconnected gaps hinder their effectiveness: (1) limited contextual awareness, particularly in dynamic or culturally nuanced interactions, and (2) inadequate mechanisms for incorporating real-time user feedback for continuous adaptation [10].

While prior work has addressed these areas in isolation, for instance, using memory networks for context [6] or reinforcement learning for feedback [14], few studies have proposed an integrated, trainable architecture that dynamically adjusts its contextual interpretation based on iterative user signals. This paper fills this gap by introducing a novel, implementable framework that combines multimodal context processing with a closed-loop feedback mechanism. Contribution:

- A detailed architectural design for a context-and-feedback-integrated conversational agent.
- An empirical evaluation demonstrating significant improvements in accuracy and user satisfaction.
- A critical analysis of the ethical risks, particularly bias, inherent in feedback-driven systems.

II. RELATED WORK & CRITICAL SYNTHESIS

➤ *Contextual Understanding:*

Advances include multimodal integration [11], [9] and cultural/emotional embeddings [3]. However, methods like hierarchical memory networks [6] often operate statically after training, lacking the ability to update contextual priors based on direct user interaction. Furthermore, as [15] critique, these models frequently exhibit bias and poor performance for underrepresented languages due to homogenized training data.

➤ *User Feedback Mechanism*

Strategies range from explicit feedback [7] to Reinforcement Learning from Human Feedback (RLHF) [14]. RLHF has proven powerful for alignment but is often resource-intensive and operates on delayed, aggregated

feedback batches [10], limiting real-time adaptability. A significant, under-discussed gap exists between collecting feedback and contextualizing it, using feedback to refine how context itself is understood in future interactions.

➤ *Synthesis and Identified Gap*

The literature reveals a parallel development of context modeling and feedback learning. My framework posits that these are not separate modules but should be co-optimized: user feedback should directly inform and refine the model's contextual reasoning capabilities in an online manner.

III. PROPOSED FRAMEWORK AND METHODOLOGY

➤ *Architecture Overview:*

The Contextual Adaptive Dialogue Agent (CADA) is designed to generate context-aware and reliable conversational responses by tightly integrating contextual reasoning with adaptive learning. As illustrated in Figure 1, the framework processes user input through a structured pipeline that captures not only the textual content of a query but also its associated sentiment and dialogue history.

At the core of CADA is the Contextual Reasoning Engine (CRE), which functions as a multimodal encoder. It jointly represents user text, sentiment signals, and extracted dialogue-history embeddings to build a rich contextual understanding of each interaction. The CRE is coupled with a Gated Context Memory, a dynamic memory module that stores and selectively retrieves relevant past interactions, cultural cues, and user preferences. While this memory is initially seeded using structured knowledge graphs [1], it continuously evolves during interaction, allowing the system to adapt to changing user behavior and conversational context.

The contextual representations produced by the CRE are passed to the Adaptive Response Generator, which integrates this information into a large language model to produce candidate responses. These responses are further examined by a Confidence Estimator, which evaluates uncertainty and supports downstream decision-making. Together, these components ensure that generated responses remain coherent, relevant, and contextually grounded.

CADA further incorporates an Adaptive Feedback Loop (AFL) that enables continuous improvement through interaction. The AFL is implemented as a reinforcement learning actor-critic framework, where the reward signal is derived from a composite of explicit user ratings, implicit behavioral cues (such as query reformulation or follow-up corrections), and a context-consistency check performed by the CRE. This feedback loop allows the system to refine its response strategies over time while maintaining alignment with contextual constraints.

Before delivery, candidate responses are ranked by a Response Optimizer, which selects the most appropriate output based on relevance, confidence, and contextual alignment. The final Agent Response is then presented to the

user, combining accurate textual content with emotionally informed signals to support natural and engaging dialogue.




Fig 1 CADA Framework Architecture Diagram.

The figure illustrates the end-to-end workflow of the Contextual Adaptive Dialogue Agent, from user interaction and contextual reasoning to adaptive learning and response optimization.

➤ *Implementation & Training*

A prototype of the CADA framework was implemented using a pre-trained T5-base model as the backbone for response generation. The CRE's gated memory was realized as a differentiable key-value store, enabling efficient retrieval and update of contextual representations. The Adaptive Feedback Loop was trained using Proximal Policy Optimization (PPO).

Training was conducted in two phases. Phase 1 involved supervised fine-tuning on a combination of the Taskmaster-1 dataset for customer service dialogues and the Wizard of Wikipedia dataset for knowledge-intensive conversations. Phase 2 focused on reinforcement learning from feedback (RLHF), where simulated user interactions were generated using a learned preference model following [11]. This was complemented by a small-scale human evaluation involving 50 participants, which was used to calibrate and validate the reward signals.

➤ *Evaluation Dataset*

A Multi-Domain Dialogue Evaluation Set (MDDES) was constructed comprising 1000 conversation threads across education and customer service scenarios, annotated for contextual appropriateness, helpfulness, and user satisfaction.

IV. RESULTS AND ANALYSIS

CADA was compared against two baselines: (1) a fine-tuned GPT-3.5-turbo model (Baseline-FT), and (2) the same model with a simple, non-adaptive feedback scorer (Baseline-FB).

Table 1 CADA Against Two Baselines

Metric	Baseline-FT	Baseline-FB	CADA (Mine)	% Improvement (vs. Baseline-FT)
Contextual Appropriateness	3.2	3.5	4.1	+28%
Helpfulness Score	3.4	3.6	4.2	+24%
User Satisfaction	3.1	3.4	4.1	+32%
Avg. Turns to Resolution	4.8	4.3	3.5	-27%

➤ *Analysis:*

CADA significantly outperformed both baselines. Qualitative analysis showed CADA excelled in scenarios requiring long-term context tracking (e.g., remembering a student's learning goal across a session) and adapting tone based on implicit negative feedback (e.g., a user rephrasing a question). The superior performance of CADA over Baseline-FB indicates that the integration and contextualization of feedback, not just its collection, are crucial.

V. ETHICAL CONSIDERATIONS AND LIMITATIONS

➤ *My Framework Introduces Specific Ethical Challenges:*

- Bias Amplification in Feedback Loops: The AFL risks reinforcing existing biases if initial user feedback is skewed. For example, in my pilot, polite but incorrect responses initially received high ratings, which the model began to replicate. I mitigated this by incorporating a fairness regularization term in the reward function, penalizing responses that deviated from verified knowledge bases in sensitive domains (e.g., healthcare).
- Transparency and Explainability: The dynamic memory update process is a "black box." I am developing a simple dashboard to show users which past interactions most influenced the current response.
- Limitations: The current prototype's computational overhead limits real-time deployment for large user bases. Furthermore, my simulated feedback, while cost-effective, cannot fully capture the nuance of human reactions.

VI. FUTURE DIRECTIONS

➤ *Future Work will Focus on:*

- Scalability and Efficiency: Exploring distillation techniques to reduce the model size while preserving performance.
- Cross-Cultural and Linguistic Validation: Actively testing and refining CADA on low-resource language datasets to address the disparity highlighted by [15]
- Advanced Multimodality: Integrating real-time visual and auditory cues [8] to enrich the CRE's context pool, moving beyond text-centric interaction.

VII. CONCLUSION

This paper demonstrates that enhancing conversational AI requires the co-evolution of contextual understanding and feedback mechanisms. My proposed CADA framework provides a concrete architecture for this integration, yielding significant improvements in dialogue quality and user satisfaction. However, this approach amplifies the need for vigilant ethical design, particularly in managing bias within adaptive loops. I argue that the next generation of conversational AI must be not only more adaptive and accurate but also more transparent and auditable, necessitating continued collaboration between researchers, ethicists, and domain experts.

REFERENCES

- [1]. Abu-Rasheed, H., Weber, C., Zenkert, J., & Fathi, M. (2023). Building contextual knowledge graphs for personalized learning recommendations. 2023 IEEE International Conference on Advanced Learning Technologies (ICALT), 36–40. <https://doi.org/10.1109/ICALT58122.2023.00016>
- [2]. Abu-Rasheed, H., Weber, C., Zenkert, J., & Fathi, M. (2023). Building contextual knowledge graphs for personalized learning recommendations. 2023 IEEE International Conference on Advanced Learning Technologies (ICALT), 36–40. <https://doi.org/10.1109/ICALT58122.2023.00016>
- [3]. Belda-Medina, J., & Calvo-Ferrer, J. R. (2022). Using chatbots as AI conversational partners in language learning. *Applied Sciences*, 12(17), 8427. <https://doi.org/10.3390/app12178427>
- [4]. Belda-Medina, J., & Calvo-Ferrer, J. R. (2022). Using chatbots as AI conversational partners in language learning. *Applied Sciences*, 12(17), 8427. <https://doi.org/10.3390/app12178427>
- [5]. Chauncey, S. A., & McKenna, H. P. (2023). A framework for ethical and responsible use of AI chatbot technology. *Computers and Education: Artificial Intelligence*, 5, 100182. <https://doi.org/10.1016/j.caei.2023.100182>
- [6]. Divekar, R. R., Drozdal, J., Chabot, S., Zhou, Y., Su, H., Chen, Y., et al. (2021). Foreign language acquisition via AI and extended reality. *Computer Assisted Language Learning*. <https://doi.org/10.1080/0958221.2021.1879162>

- [7]. Gedikli, F., Jannach, D., & Ge, M. (2014). How should I explain? A comparison of explanation types for recommender systems. *International Journal of Human-Computer Studies, 72*(4), 367–382. <https://doi.org/10.1016/j.ijhcs.2013.12.007>
- [8]. Hu, D., Wei, L., & Huai, X. (2021). DialogueCRN: Contextual reasoning networks for emotion recognition in conversations. arXiv preprint arXiv:2106.01978.
- [9]. OpenAI. (2023). GPT-4 system card. <https://cdn.openai.com/papers/gpt-4-system-card.pdf>
- [10]. Owoicho, P., Sekulic, I., Aliannejadi, M., Dalton, J., & Crestani, F. (2023). Exploiting simulated user feedback for conversational search. Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, 1–4. <https://doi.org/10.1145/3539618.3591683>
- [11]. Park, S., Li, H., Patel, A., Mudgal, S., Lee, S., Kim, Y.-B., Matsoukas, S., & Sarikaya, R. (2021). A scalable framework for learning from implicit user feedback. arXiv preprint arXiv:2010.12251.
- [12]. Safi, Z., et al. (2020). Technical aspects of developing chatbots for medical applications. Journal of Medical Internet Research, 22(12), e19127. <https://doi.org/10.2196/19127>
- [13]. Stahl, B. C., & Eke, D. (2024). The ethics of ChatGPT. International Journal of Information Management, 74, 102700. <https://doi.org/10.1016/j.ijinfomgt.2023.102700>
- [14]. Tai, T.-Y. (2022). Effects of intelligent personal assistants on EFL learners' oral proficiency. Computer Assisted Language Learning. <https://doi.org/10.1080/09588221.2022.2075013>
- [15]. Zhai, C., & Wibowo, S. (2023). A systematic review on AI dialogue systems for enhancing interactional competence. Computers and Education: Artificial Intelligence, 4, 100134.