Performant LLM Agentic Framework for Conversational AI

Anonymous Author(s) Anonymous Institution(s) Emails: *anonymous@domain*

Abstract—The rise of Agentic applications and automation in the Voice AI industry has led to an increased reliance on Large Language Models (LLMs) to navigate graph-based logic workflows composed of nodes and edges. However, existing methods face challenges such as alignment errors in complex workflows and hallucinations caused by excessive context size. To address these limitations, we introduce the Performant Agentic Framework (PAF), a novel system that assists LLMs in selecting appropriate nodes and executing actions in order when traversing complex graphs. PAF combines LLM-based reasoning with a mathematically grounded vector scoring mechanism, achieving both higher accuracy and reduced latency. Our approach dynamically balances strict adherence to predefined paths with flexible node jumps to handle various user inputs efficiently. Experiments demonstrate that PAF significantly outperforms baseline methods, paving the way for scalable, real-time Conversational AI systems in complex business environments.

Index Terms—Machine Learning, Agentic Workflow, LLM Agent, Agentic, Voice AI, LLM Alignment, Agentic Framework

I. INTRODUCTION

Graph-based workflows are central to numerous business processes across industries such as education, legal, healthcare, and customer support. These workflows represent decisionmaking steps as nodes, and connections between them as edges. The rise of Conversational AI within these spaces introduces new challenges. Autonomous agents, powered by large language models (LLMs), are increasingly being used to navigate these workflows, enabling the automation of complex business processes (1). Each node in the workflow contains specific instructions or prompts that guide the agent's speech generation and certain actions to trigger. Nodes can be classified into several types, including Start Nodes, which define the root and entry point of a workflow; End Nodes, which signal the termination of the workflow; and generic Nodes, which serve as intermediate decision points containing speech instructions for the LLM to converse with users in predefined ways. Additionally, Transfer Nodes in Conversational AI workflows allow for transitioning the conversation to another autonomous or human agent. Edges between nodes may include logical conditions that dictate the agent's transitions, ensuring workflows are executed accurately.

Figure 1 illustrates how tasks such as determining health care eligibility can be broken down into nodes, edges, and conditions. For example, a healthcare provider might use such a workflow to efficiently filter out patients without the required insurance, reducing the burden on human agents. However, workflows like these can rapidly grow in complexity. As shown in Figure 2,



Fig. 1. Example illustration of an Agentic workflow for a healthcare call center use case, where the Agent needs to route calls based on different conditions.

adding just a few additional conditions to the conversation flow can drastically increase the number of nodes and edges, making the workflow more difficult to manage and execute effectively.

Although LLMs such as GPT and LLAMA are built on autoregressive decoder-based transformer architectures optimized for natural language generation, they are not inherently designed to handle structured, multi-step processes with extensive context (2; 3). Existing approaches have been to add a planning phase, where the LLM would take time to orchestrate the action, and then proceed to the generation tasks (4; 5). However, this approach is not optimal to the Conversational AI use case, as it would increase the overall latency by doubling the number of queries needed. Tasks such as managing end-toend customer service requests with non-standard return policies, performing outbound sales calls that involve dynamic CRM updates, or redirecting users to appropriate departments after a sequence of filtering questions require precision, alignment, and low-latency responses. These limitations force businesses to oversimplify workflows, sacrificing accuracy and operational efficiency—an outcome that is contrary to their objectives.

The challenges inherent in adapting LLMs to graph-based



Fig. 2. Example illustration of an Agentic workflow for an internet service company helping callers troubleshoot connection issues. This workflow demonstrates how a more complex use case can have more conditions, nodes, and edges.

workflows underscore the need for new approaches that can accurately and efficiently execute workflows while respecting real-world constraints such as latency. While adding more reasoning steps could theoretically improve accuracy, such methods are impractical for Conversational AI applications where rapid response times are critical.

To address these challenges in the current Conversational AI space, this paper introduces the **Performant Agentic Framework (PAF)**, a novel solution for efficient graph traversal that balances accuracy and latency in real-world applications. By leveraging both traditional decision-making logic and mathematical methods for next-node selection, PAF enables agents to execute workflows with greater precision and speed. Our experiments demonstrate that PAF significantly outperforms baseline and traditional methods in both accuracy and latency, as evidenced by higher alignment scores and reduced response times.

The reliance on LLM-based systems to execute graph-based workflows has seen significant research attention, particularly in developing frameworks that aim to balance accuracy, latency, and alignment with predefined workflows. Below, we discuss prominent related works and their limitations.

Agentic Framework Serving as examples, LangChain (6) and LangGraph (7) streamline graph-based workflows by utilizing function calls and prompt chaining. While effective for simple tasks, their reliance on keyword-based triggers often results in alignment errors, especially in workflows with hundreds or thousands of nodes. These frameworks lack robustness for real-world applications where actions must be dynamically triggered at various points in conversations. Furthermore, their reliance on LLM-generated triggers leads to unreliability in critical workflows, where adherence to predefined paths is essential for compliance and business logic (6; 7). Additionally, limitations in LLM context windows further exacerbate their inefficiency in retaining relevant information across extended workflows, introducing hallucinations and context drift during execution (8).

Conversational AI Conversational AI has been a key focus for Natural Language Processing. The existing Conversational AI solutions emphasize the need for multi-modality, guardrails, and advanced tuning to enhance dialogue quality (9). Prior approaches to the Voice AI space have been proven to work in a sandbox conversational setting (10), but lack the consistency and accuracy required for production use. As suggested, LLMs miss certain abilities to maintain performance in a dynamic conversational setting, unable to handle numerous tasks conditionally while reducing hallucinations and staying within context (8; 9; 11).

MetaGPT and SOP Translation MetaGPT leverages Standardized Operating Procedures (SOPs) to structure workflows, enabling agents to replicate domain-specific expertise. However, its reliance on iterative planning and validation increases latency, making it unsuitable for real-time applications. For example, as noted in (12), the planning phase requires additional LLM calls, which adds computational overhead. While MetaGPT is effective for SOP alignment, it struggles with unusual user inputs and extended workflows, leading to significant context drift. Its dependence on domain-specific fine-tuning also hinders generalizability, limiting its use in broader applications (12; 13).

Comparison and Our Contributions Existing frameworks have made valuable contributions but are hindered by issues such as context drift, high latency, and alignment errors. PAF addresses these limitations by replacing LLM planning phases with a mathematical decision-making approach, combining vector-based node selection and specialized prompt engineering. Unlike previous methods, PAF reduces context size while improving accuracy, making it a scalable and production-ready solution for navigating graph-based workflows.

III. PERFORMANT AGENTIC FRAMEWORK (PAF)

PAF is a framework designed for Agentic workflows, enabling agents to navigate graph-based structures composed of nodes and edges to execute predefined workflows. It is comprised of two components: Basic PAF and Optimized PAF, each tailored to address specific challenges in workflow execution.

A. Basic PAF

Problem Formulation PAF enables agents to operate by following nodes connected by logical edges. During each generation turn, the agent follows the nodes in the graph according to the logical conditions specified as outcomes of the node. If a condition is met, the agent navigates to and executes the instructions of the next node in the graph.

Our PAF involves leveraging LLM as a Judge for identifying the Agent's location in the map dynamically per each generation as follows:

Algorithm I LLW as a judge for Node Identification	Algorithm	1	LLM	as	а	Judge	for	Node	Identification
--	-----------	---	-----	----	---	-------	-----	------	----------------

Input: ConversationHistory, NavigationMap, LastestIdentifiedNode

Output: UpdatedLatestIdentifiedNode

Step 1: Format Input for the LLM

Construct a prompt using ConversationHistory

Add a contextual anchor by traversing from the StartNode to *LastestIdentifiedNode* and collect all first layer children Nodes in the map, e.g., "You were previously on Node {LastestIdentifiedNode} with options to navigate to in the map {Path} each with instructions being..."

If LastestIdentifiedNode is unavailable, use: "This is the start of the task {task}, proceed to Node 0."

Step 2: Query the LLM

Send a question to the LLM:

"Based on your latest responses, where are you currently in the navigation map?"

Step 3: Process the Response

Parse the response to identify the node mentioned by the LLM.

Validate the identified node against *NavigationMap*. **Step 4: Return the Result**

	Output	the	validated	node	as
Updc	tedLatestId	entified	Node.		

This design is particularly effective in production AI systems as it separates the generation tasks from other downstream modules, like Text-to-Speech (TTS). This modular approach optimizes latency, enabling parallel processing by downstream services such as a TTS service. Compared to implementations where prompts are added in a single body, Basic PAF achieves lower error rates by using a step-by-step logic tree, reducing the need for additional validation iterations through customized testing schemas (14; 15).

Algorithm 2 Basic Agentic Framework

Input: ConversationHistory, NavigationMap, LatestIdentifiedNode

Output: UpdatedLatestIdentifiedNode

Step 1: Initialize LLM Instructional Message

Construct an instructional prompt for the LLM agent. Add *ConversationHistory* to the prompt in a formatted structure.

Include instructions based on *LatestIdentifiedNode*, e.g., "You are currently on Node {LatestIdentifiedNode}."

Constructed navigation prompt by traversing the NavigationMap and collecting all first layer children nodes' instructions on the {Path} from StartNode to LatestIdentifiedNode.

Step 2: Query the LLM

Send the query to the LLM:

"Based on the navigation map and your current node, respond to the user question: {user question}."

Step 3: Process LLM Output in a Streaming Loop

while LLM agent streams output do

(a) Identify Current Node:

Use Algorithm 1 to determine the node the agent selects.

(b) Update LatestIdentifiedNode:

Set *LatestIdentifiedNode* to the node identified in Step (a).

(c) Trigger Actions:

Execute any actions related to the newly identified *LatestIdentifiedNode*.

(d) Update NavigationMap:

Modify NavigationMap as needed based on the new LatestIdentifiedNode.

end while

Step 4: Return the Updated Node

Output UpdatedLatestIdentifiedNode = LatestIdentifiedNode.

B. Optimized PAF

While Basic PAF offers significant improvements (shown later in the experiment section), it faces bottlenecks when workflows involve numerous nodes (e.g., 50 nodes with 4 conditions each). These bottlenecks arise as the agent struggles to differentiate between semantically similar prompts on different paths of the graph. Optimized PAF addresses this with vector-based scoring to reduce the size of the context window and improve logical adherence.

The heart of optimized PAF is the **Vector Node Search**, which evaluates nodes using embedding models with a confidence threshold as follows:

Optimized PAF leverages vector-based reasoning, incorporating both semantic direction and magnitude through the dot product. Notably, when comparing different metrics to use as a similarity score, we found that the dot product is particularly

Algorithm 3 Vector-Based Node Search

Input: NavigationMap, LatestIdentifiedNode, Threshold, LatestAgentResponse

Output: *UpdatedLatestIdentifiedNode*

Step 1: Vectorize Instructions and Agent Response

Compute vector representations for:

(a) LatestIdentifiedNode.

(b) LatestIdentifiedNode's children nodes.

(c) LatestAgentResponse from the LLM.

Step 2: Compute Similarity Scores

Compare the vector representation of LatestAgentResponse against:

(a) Vector for *LatestIdentifiedNode* instruction.

(b) Vectors for instructions of child nodes.

Compute similarity scores using a suitable metric (e.g., dot product similarity).

Step 3: Identify the Best Matching Node

Select the node with the highest similarity score.

If the score exceeds the *Threshold*, assign the corresponding node as *UpdatedLatestIdentifiedNode*.

Step 4: Update and Return

if A node is identified with a score above the *Threshold* then

UpdateLatestIdentifiedNodetoUpdatedLatestIdentifiedNode.Return UpdatedLatestIdentifiedNode.

else

Return false (Use LLM as a Judge Approach).

end if

effective for Conversational AI applications. This finding aligns with research by (16), which demonstrates the advantages of using dot product as a vector score over cosine similarity, where cosine similarity may produce ambiguous results by ignoring magnitude. This is particularly relevant when dealing with over-fitted domain jargon, where it is critical for the agent to differentiate between subtly distinct expressions that hold drastically different implications. This approach aligns well with emerging models like OpenAI's text-2-vec-3-small (17), which are tuned to reflect confidence alongside semantic direction.

IV. EXPERIMENT

To evaluate the effectiveness of PAF, we designed experiments to compare the performance of PAF with existing graph traversal and node selection methodologies. These experiments are designed to assess the latency, accuracy, and alignment of the framework across various workflows, particularly in Conversational AI applications.

A. Experiment Setup

The experiments utilize a simulated dataset generated to mimic real-world workflows.

Dataset Generation: The experiment utilized a synthetic dataset generated to simulate real-world workflows. Each dataset entry contained:

Algorithm 4 Optimized Agentic Framework

Input:ConversationHistory,NavigationMap,LatestIdentifiedNode, Threshold

Output: UpdatedLatestIdentifiedNode

Step 1: Precompute Vectorized Instructions

Compute and store vectorized representations for instructions at all nodes in *NavigationMap*.

Step 2: Format Input for the LLM

Construct a message including:

(a) Formatted ConversationHistory.

(b) Instructions for LatestIdentifiedNode.

(c) A constructed navigation prompt by traversing NavigationMap and collecting all first-layer children nodes' instructions on the {Path} from StartNode to LatestIdentifiedNode.

Step 3: Query the LLM Agent

Send the constructed message to the LLM agent.

Step 4: Process LLM Output in a Streaming Loop

while LLM agent streams output do

(a) Perform Vector-Based Node Search (Algorithm 3).

if a node is successfully identified then

Proceed to Step 5.

else

(b) Perform LLM as a Judge (Algorithm 1).

end if

end while

Step 5: Update Current Node

Set *LatestIdentifiedNode* to the node identified in Step 4.

Step 6: Trigger Actions and Update the Graph

(a) Execute any actions related to the updated *LatestIdentifiedNode*.

(b) Modify NavigationMap as needed.

Step 7: Return the Updated Node

- SystemPrompt: A node navigation map with Agentic logic.
- **ConversationHistory:** Turn-by-turn interactions between the user and the agent.
- **GoldenResponse:** A pre-verified response audited through LLM-As-a-Judge and human evaluation, serving as the ground truth.

Conversations were executed by two agents (a "user" LLM and a "response" LLM), with a random turn length (6–10). Golden responses were derived from the corresponding node's instruction and validated by humans.

B. Framework Performance Evaluation

We evaluated three frameworks under the following metrics:

• Semantic Similarity: Alignment between the generated response and the golden response using OpenAI's text-2-vec-3-small embedding model (17).

- **Total Complete Hit Rate:** Percentage of responses that achieved a similarity score above 0.97.
- Mean and Median Similarity Scores: Overall alignment performance.

Frameworks Tested:

- 1) **Baseline:** A naive approach treating the entire conversation as a single prompt.
- 2) **Basic PAF:** A step-by-step logic tree leveraging LLM-as-Judge (Algorithms 1 and 2).
- 3) **Optimized PAF:** A vector-based approach (Algorithms 1, 3, 4).

C. Hypotheses

- H1: Basic PAF achieves higher average similarity than Baseline.
- **H2:** Optimized PAF achieves higher average similarity than Baseline.
- H3: Optimized PAF achieves higher average similarity than Basic PAF.

A one-sided paired t-test with $\alpha = 0.05$ was used for statistical significance.

D. Experiment Steps

- 1) Simulate 100 conversations using the predefined work-flow.
- 2) Generate responses for each method.
- 3) Compute similarity scores against the golden responses.
- 4) Aggregate metrics such as total hit rate, mean, and median.
- 5) Perform hypothesis testing.

E. Results

TABLE I Result Metrics Across Algorithms

Method	Total Hits	Count above 0.8	Mean	Median
Baseline	0	0	0.391	0.387
Basic PAF	16	14	0.481	0.400
Optimized PAF	35	22	0.594	0.496

 TABLE II

 STATISTICAL COMPARISON RESULTS (ONE-SIDED PAIRED T-TESTS)

Comparison	t-statistic	p-value
Baseline vs Basic PAF	2.9982	0.0020
Baseline vs Optimized PAF	7.3077	0.0000
Basic PAF vs Optimized PAF	4.2494	0.0000

- **H1:** Basic PAF significantly outperforms the Baseline (p = 0.002).
- **H2:** Optimized PAF significantly outperforms the Baseline (p < 0.001).
- H3: Optimized PAF significantly outperforms Basic PAF (p < 0.001).



Fig. 3. Distribution of Similarity Scores for the 3 frameworks (Naive=Baseline, Base=Basic PAF, Optimized=Optimized PAF).

F. Reproducibility

We provide code and data generation scripts in an anonymized repository.¹

V. CONCLUSION

Our approach introduces novel mechanisms for leveraging LLMs to navigate graph-based workflows, replacing the need for extensive planning phases and minimizing error rates. PAF achieves faster response times and greater accuracy in real-world applications by reducing reliance on large context windows and optimizing computational steps.

In summary, PAF resolves key limitations in existing Agentic frameworks by:

- Removing extra iterations for validation and planning, thereby reducing latency.
- Improving alignment through step-by-step logic tree navigation.
- Reducing context window size by focusing on relevant graph information.
- Introducing vector-based scoring of semantic similarity, reducing redundant LLM calls.

VI. FUTURE WORK

While Conversational AI serves as a compelling case study, the PAF framework holds promise for broader applications. Future research will explore:

- Node Weights and Path Rules: Introducing weights and flexible rules.
- **Integration with Different Models:** Experimenting with in-house or domain-specific LLMs.
- **Open-Source Model Improvements:** Tuning embeddings or scoring for domain-specific tasks.

¹https://anonymous.4open.science/r/performant-agentic-framework-F5F6/ README.md

IMPACT STATEMENT

This paper presents work whose goal is to advance the field of Machine Learning and Agentic Workflows. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

REFERENCES

- [1] L. Zhuge, K. Zhang, and M. Wang, "Aflow: Automating agentic workflow generation," *arXiv preprint*, 2023.
- [2] Y. Qiu and Y. Jin, "Chatgpt and finetuned bert: A comparative study for developing intelligent design support systems," *Intelligent Systems with Applications*, vol. 21, p. 200308, 2024.
- [3] X. Shi, J. Liu, and Y. Song, "Bert and Ilm-based multivariate hate speech detection on twitter: Comparative analysis and superior performance," in *International Artificial Intelligence Conference*. Singapore: Springer Nature Singapore, November 2023, pp. 85–97.
- [4] K. Valmeekam, M. Marquez, S. Sreedharan, and S. Kambhampati, "On the planning abilities of large language models - a critical investigation," *Advances in Neural Information Processing Systems*, vol. 36, pp. 75993– 76005, 2023.
- [5] Z. Zhou, J. Song, K. Yao, Z. Shu, and L. Ma, "Isr-Ilm: Iterative self-refined large language model for long-horizon sequential task planning," in 2024 IEEE International Conference on Robotics and Automation (ICRA). IEEE, May 2024, pp. 2081–2088.
- [6] LangChain, "Langchain documentation," https://www. langchain.com, 2023, accessed: Jan 19 2025.
- [7] LangGraph, "Langgraph overview," https: //www.langchain.com/langgraph, 2023, accessed: Jan 19 2025.
- [8] F. Dong and R. Qian, "Characterizing context influence and hallucination in summarization," *arXiv preprint*, 2024.
- [9] X. L. Dong, S. Moon, Y. E. Xu, K. Malik, and Z. Yu, "Towards next-generation intelligent assistants leveraging llm techniques," in *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM, August 2023, pp. 5792–5793.
- [10] A. J. James, N. D. Vangapalli, J. Siripurapu, and Y. R. Chinnamallu, "Integration of voice assistant with chatgpt and dall-e," in 2024 International Conference on Emerging Techniques in Computational Intelligence (ICETCI). IEEE, August 2024, pp. 95–101.
- [11] S. S. Gill and R. Kaur, "Chatgpt: Vision and challenges," *Internet of Things and Cyber-Physical Systems*, vol. 3, pp. 262–271, 2023.
- [12] Z. Gao, Y. Li, and W. Sun, "Metagpt: Sop-driven workflow generation," arXiv preprint, 2023.
- [13] Q. Wang and B. Liu, "Extending metagpt for complex workflow tasks," *arXiv preprint*, 2024.
- [14] F. Li and J. Yuan, "Improving multi-step planning with workflow-aware llms," *arXiv preprint*, 2023.

- [15] S. Reddy and A. Gupta, "Abstractive text summarizer: A comparative study on dot product attention and cosine similarity," *IEEE Explore*, 2021.
- [16] J. Huang and S. Wang, "Dot product vs. cosine similarity in domain-specific retrieval," *IEEE Transactions on Information Retrieval*, 2021.
- [17] OpenAI, "Openai embedding model," https://platform. openai.com/docs/guides/embeddings/, 2025, accessed: Jan 19 2025.