LLM-based Few-shot Action System for NPCs in Virtual Reality Games

Fan Wang
Meta Platforms, Inc.
Burlingame, CA, USA
e-mail: fanwang@meta.com

Wen Zhou
Meta Platforms, Inc.
Burlingame, CA, USA
e-mail: zhouwen@meta.com

Rongze Gui
Meta Platforms, Inc.
Burlingame, CA, USA
e-mail: danielgui@meta.com

Jinqiao Li
Meta Platforms, Inc.
Columbus, OH, USA
e-mail: jqli@meta.com

Radoslaw Malicki
Meta Platforms, Inc.
Vancouver, BC, Canada
e-mail: radmalicki@meta.com

Andrey Staroseltsev
Meta Platforms, Inc.
Burlingame, CA, USA
e-mail: staroseltsev@meta.com

***Abstract*—Current trends in the game industry include making games more immersive and realistic through developing games in Virtual Reality (VR) and integrating Generative Artificial intelligence (AI) in Non-Player Characters (NPCs). As Large Language Model (LLM) based conversational NPCs start to emerge and show success in traditional video game medium, we seek to answer the question of "can we leverage LLMs to build believable NPCs that can make logical actions and interact with players naturally in VR?" In this paper, we introduce a design of an LLM-based action system with few-shot learning for NPCs in VR worlds. The use of few-shot learning would allow for rapid adaptation to new games without massive data requirement or expensive model training. We also include an evaluation plan to assess our designed system's performance.**

Keywords-Action Agent; Large Language Model; NPC in Virtual Reality; Agentic Workflow; Retrieval-Augmented Generation.

I. INTRODUCTION

In game worlds, Non-Player Characters (NPCs) are software agents whose actions are not directly controlled by a human player. Player-NPC interaction serves as an important role in enhancing players' experience, by providing companionship and motivating social engagement, giving guidance or helping to advance the game plots, or just adding to the world's ambiance to make the world more dynamic and lifelike [1][2].

In recent years, the rapid advances in Large Language Models (LLM), such as GPT4 [3] and LLama 3 [4] have been impacting the video game field [1]. There is increasing interest from game developers in experimenting and integrating LLM-based AI NPCs in video games [5]. In industry, companies including Inworld AI (partnered with Microsoft Xbox) [6] and ConvAI (partnered with NVidia) [7] are building products that enable developers to build conversation-based AI NPCs that can listen to players and output natural language along with facial expressions. In academics, there is also quite a few previous research proposing design for intelligent NPCs [8][9]. However, most of this work was done on games from a traditional medium (a 2D screen) and the primary function of the LLM-based NPCs is to have conversations. Little has been done with a focus on action planning for NPCs in a Virtual Reality (VR) environment.

VR unveiled a new degree of immersion by enabling the feeling of "presence" for players, 3D environment inputs, and physical character embodiment, which makes VR an ideal medium for realistic player-NPC interaction. Most recently, Meta introduced generative AI NPCs in Meta Horizon Worlds such as Bobber Bay Fishing [10], but those NPCs are primarily conversational. Previous studies have shown that players value the interactivity of physical motion uniquely offered by VR - it would be favored to have NPCs that can provide continuous human-like agency or physical actions, which are triggered by natural interactions (dialogues or gestures) instead of controller buttons or story plots [2][11][12]. We hope to leverage the immersive advantages of VR, examine the potential of LLMs in VR gameplay, and lay the groundwork for future development of believable VR NPCs that:

- Can be invoked by natural player-NPC interaction and understand players' intention;
- Can perceive and understand the 3D VR world environment, plan and execute physical actions intelligently;
- Can be easily extended to different games with a few-shot learning schema.

There are five sections in this paper. In Section 2, we discuss previous work that is related to our research topic. In Section 3, we introduce a design for a scalable LLM-based NPC action system. In Section 4, we present the basic prototyping experiment settings and its preliminary results, as well as a plan for a full, comprehensive evaluation. In Section 5, we briefly summarize the conclusion and future work.

II. RELATED WORK

Most NPCs in conventional computer games make actions based on a pre-defined action system heuristics. Traditional AI planning algorithms are generally used to power such action systems. One of the most prominent approach is Goal-Oriented Action Planning (GOAP), where NPC / Agents are driven with specific goal and heuristics. Earlier researches had been focusing on these algorithms, including some examples of application and optimization in hide-and-seek games [13][14]. Researches also focused on improve the naturalness and believably for NPCs, through techniques like dynamic reputation system and Observer-Orient-Decide-Act (OODA) theories [15][16]. Recent researches also explored machine-Learning techniques to enable more intelligent NPC actions, including behavior tree design, and cognitive and emotional models (FatiMA and PSI) [17][18].

After the LLM surge in 2022, researches also discussed enriching NPC through Generative-AI technologies. Some of the LLM-based AI Agents explore VR game settings to enrich conversational and behavioral experiences in NPC-human interaction [2][9][19]. Few researches have touched on LLM-based action system. The State-of-the-Art (SOTA) gameplay agent with action is VOYAGER, which created a single Generative-AI powered agent player in a Mincraft game [8]. The work used a Retrieval-Augmented Generation (RAG) powered iterative prompting method to prompt a GPT-4 model and execute actions in gameplay. However, generative action-system for interactive NPCs in VR are not widely researched.

III. PROPOSED DESIGN

The proposed design introduces a novel action system for LLM-based NPCs in VR games. This framework leverages dynamic world perception and game knowledge injection to enable intelligent, context-aware interactions between NPCs and players. The system is designed to extend across multiple types of games using few-shot learning, supporting interactions through natural inputs from player actions or in-game events. As illustrated in Figure 1, the system has the following key components:

1. Dynamic World Perception Injection:

- The system continuously gathers real-time data from the game world, including:
  - **Objects:** Items available in the environment.
  - **Events:** Ongoing activities or situations.
  - **State:** Current game or player status (e.g., health, location).
  - **Location:** Spatial data of game elements or characters.
- This perception data is filtered and ranked to ensure only the most relevant information influences NPC behavior.

2. Game Knowledge Injection:

- Game-specific rules, available functions, and object interaction logic are injected into the framework.
- An *available action list* is dynamically updated to ensure NPC behavior aligns with the current game context and mechanics.

3. NPC Memory Injection:

- Based on the world perception and game rules, the relevant information from following memories are also retrieved for NPCs to plan the next actions:
  - NPC's knowledge about itself: identity, personality, backstory.
  - NPC's short-term memory within the same session.
  - NPC's long-term memory.
- These memories are also accumulated and summarized for developing reflections.

4. Prompt Construction:

- The core of the action system is prompt engineering, which combines:
  - **World Perception Data:** Extracted objects, events, and states.
  - **Game Rules and Knowledge:** Embedded functions and interactions.
  - **NPC Data:** NPC's persona, short-term and long-term memory.
  - **Few-shot Learning Examples:** Pre-selected or dynamically chosen examples.
  - **User Commands:** Inputs from the player or game system.
- Multiple formats, such as JSON, YAML, and XML, are used to structure these prompts.
- An *embedding-based example selector* helps identify the most relevant few-shot learning examples to tailor NPC behavior.
- The prompt is also augmented with *chain-of-thought* [20] for improving NPC's planning logic.

5. Model Zoo Utilization:

- The system leverages a collection of LLMs, such as LLaMA-3 models (8B and 70B) [4], to generate structured outputs.
- These models are queried to predict an appropriate action plan or natural language responses for NPCs.

6. Structured Output:

- The framework provides actionable outputs in a structured format to ensure consistency across various game scenarios. Key outputs include:
  - **Action or No Action Decision:** Determining if an NPC should respond to a given scenario.
  - **Function Call List:** A sequence of game functions to be executed.
  - **Function Chaining:** Enabling complex actions to unfold over time.
  - **Natural Language Response:** Generating appropriate conversational dialogue between NPCs and players.
- NPC will execute actions based on this structured output, and the outcome resulted from these actions are fed back to NPC's memory as a reflection.

This design framework creates a cohesive system where NPCs can perform dynamic actions and engage in meaningful dialogues based on real-time game events. The use of few-shot learning ensures adaptability to various game types with minimal retraining, offering a scalable solution for enriching player-NPC interactions across VR worlds.

IV. EXPERIMENTS & EVALUATION PLAN

To validate the effectiveness of our proposed LLM-based NPC action system, we conducted a basic prototyping experiment, and also planned a comprehensive set of enriched experiments across several domains. These experiments focus

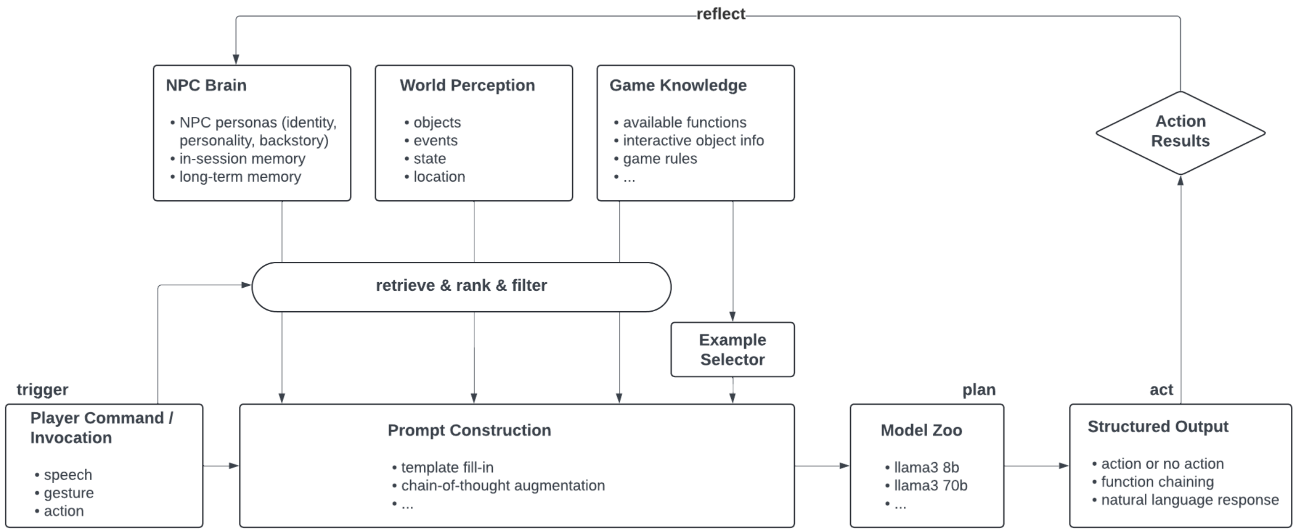

Figure 1. Overview of few-shot LLM-based NPC action system design.

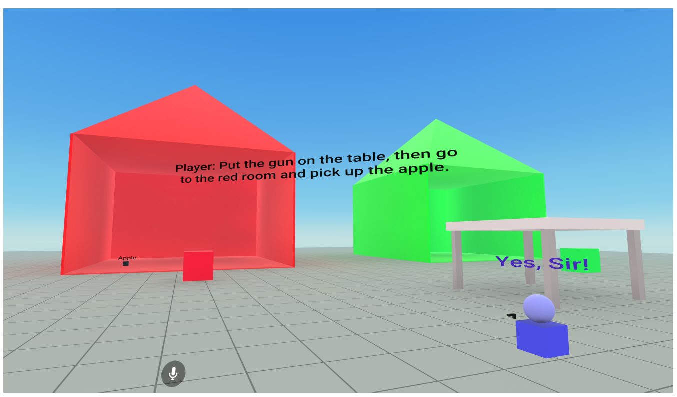

Figure 2. Example test setup for the basic prototyping experiment. The end-to-end playtest is done in Meta Horizon World with a Meta Quest headset [21].

on assessing the model's adaptability, natural invocation, and ability to generate accurate responses in real-time dynamic VR environments.

A. Basic Prototyping Experiment Setup

We implemented the basic skeleton in the system design and did a prototyping experiment to prove the feasibility of this design, using Llama models with H100.

1. Data Generation and Scenario Design:

- We generated around 120 game scenarios with varying environment settings and user commands to test the prototype action system.
- An example scenario is (as shown in Figure 2):
  - **Environment:**
    - You **SEE** green target, red target, tree, green room, red room, and player.
    - You **HAVE** a gun.
    - You are **NEAR** a table.
  - **User Command:** Put the gun on the table, then go to the red room and pick the apple.

2. Prototype Evaluation and Prompt Iteration:

- We iteratively tested the prototype system with multiple prompt designs to improve the performance.
- We explored different formats such as **JSON**, **YAML**, and **XML** to evaluate their impact on response generation and structured outputs.

3. Few-shot Learning Evaluation:

- To improve response quality, we tested different numbers of few-shot examples that cover different in-game scenarios, such as invalid but semantically correct command.
- Preliminary results showed that more relevant examples improve system performance, validating the importance of example selection.

B. Enriched Experiments Plan

In addition to basic prototyping interaction scenarios, we planned more complex experiments aligned with the paper's goal of achieving natural, multi-game adaptive interactions:

- **Cross-Game Adaptability:** The system is planned to be tested on scenarios from different types of games (e.g., puzzle games, adventure games, and role-playing games) to evaluate how well the LLM-based NPCs adapt without retraining.
- **Multi-Step Tasks:** Evaluations included sequences requiring NPCs to perform chained actions (e.g., collect, carry, and deliver items) to measure function chaining effectiveness.
- **Real-Time Inputs and Dynamic Changes:** NPCs will be exposed to changing environments mid-action (e.g., new objects appearing or targets moving) to test real-time adaptability.

- **Voice Input Integration:** In alignment with VR settings, we will test voice-based user inputs, assessing how accurately the system processes and acts on these commands.

C. Evaluation Metrics

The evaluation plan focuses on key metrics relevant to the system's performance across different environments and interaction scenarios:

- **Accuracy:**
  - Measured by how accurately the NPCs interpret user commands and execute the intended actions.
  - We already have preliminary proof-of-concept results from different prompt formats for the Basic Experiment setting:

TABLE I. MODEL ACCURACY WITH DIFFERENT PROMPT FORMATS.

Format	*Model*	*Accuracy*
JSON	llama3-70B	86.40%
JSON	llama3-8B	70.91%
YAML	llama3-8B	76.40%
XML	llama3-8B	69.10%

- **Response Latency:** Evaluated based on how quickly the NPCs respond to commands, ensuring timely interaction in VR environments.
- **Action Quality:** Assessed by measuring whether multi-step actions were executed correctly and in sequence.
- **Natural Invocation Success:** A measure of how accurately the system responds to natural language inputs without requiring predefined triggers or extensive training.
- **Adaptability Across Games:** Evaluated by measuring how well the system performs in multiple game genres with minimal prompt reconfiguration.
- **User Study Results:** We will engage human subjects participating in multiple playtest sessions with a VR headset and give feedback through a survey. The survey will be focusing on understanding the participants' subjective satisfaction/affection ratings on 1) their interaction with the NPCs (including invocation), and 2) the intelligent level of NPC's action planning.

V. CONCLUSION & FUTURE WORK

This paper presented a design for an LLM-based few-shot action system for NPCs in VR games. We did a proof-of-concept experiment with the very basic implementation, and the results revealed acceptable accuracy as well as an emphasis on the importance of example selection. We have also proposed a comprehensive experiment and evaluation plan to be done after the full implementation, which will allow us to assess the effectiveness of the system in handling the complex, dynamic, and natural interactions across various VR game types.

In the future, we seek to further improve our system, and explore creative solutions for using the latest technology to deliver new level of immersion, interactivity, and engagement for VR game players. Specifically, we will compare our system with other action-generation frameworks such as GOAP to evaluate the strengths and trade-offs of different approaches in game scenarios through user study. Key limitations of the current system include hallucination, and privacy concerns, which we plan to mitigate through introducing response grounding and user identifiable information masking. We also aim to implement a fuzzy memory module to better protect sensitive user data.

REFERENCES

[1] R. Gallotta et al., "Large language models and games: A survey and roadmap," *IEEE Transactions on Games*, 2024.

[2] M. Yin and R. Xiao, "Press a or wave: User expectations for npc interactions and nonverbal behaviour in virtual reality," *Proceedings of the ACM on Human-Computer Interaction*, vol. 8, no. CHI PLAY, pp. 1–25, 2024.

[3] J. Achiam et al., "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.

[4] A. Grattafiori et al., "The llama 3 herd of models," *arXiv preprint arXiv:2407.21783*, 2024.

[5] M. Sas, "Unleashing generative non-player characters in video games: An ai act perspective," in *2024 IEEE Gaming, Entertainment, and Media Conference (GEM)*, IEEE, 2024, pp. 1–4.

[6] "Inworld - Build custom AI applications for gaming," Accessed: May 12, 2025. [Online]. Available: https://inworld.ai/ai-for-gaming.

[7] "ConvAI - Conversational AI Characters," Accessed: May 12, 2025. [Online]. Available: https://convai.com/.

[8] G. Wang et al., "Voyager: An open-ended embodied agent with large language models," *arXiv preprint arXiv:2305.16291*, 2023.

[9] H. Wan et al., "Building llm-based ai agents in social virtual reality," in *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–7.

[10] "Bobber Bay Fishing," Accessed: May 12, 2025. [Online]. Available: https://horizon.meta.com/world/365118142359176/.

[11] Y. Jang and E. Park, "An adoption model for virtual reality games: The roles of presence and enjoyment," *Telematics and Informatics*, vol. 42, p. 101 239, 2019.

[12] M. Ochs, N. Sabouret, and V. Corruble, "Simulation of the dynamics of nonplayer characters' emotions and social relations in games," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 1, no. 4, pp. 281–297, 2009.

[13] E. Long, "Enhanced npc behaviour using goal oriented action planning," *Master's Thesis, School of Computing and Advanced Technologies, University of Abertay Dundee, Dundee, UK*, 2007.

[14] D. A. Suyikno and A. Setiawan, "Feasible npc hiding behaviour using goal oriented action planning in case of hide-and-seek 3d game simulation," in *2019 Fourth International Conference on Informatics and Computing (ICIC)*, IEEE, 2019, pp. 1–6.

[15] J. Mooney and J. M. Allbeck, "Rethinking npc intelligence: A new reputation system," in *Proceedings of the Seventh International Conference on Motion in Games*, 2014, pp. 55–60.

[16] Z. Lin, Z. Zhang, X. Sun, and H.-T. Zheng, "Make npc more realistic: Design and practice of a hybrid stealth game npc ai framework based on ooda theory," in *International Conference on Human-Computer Interaction*, Springer, 2024, pp. 321–328.

[17] M. Y. Lim, J. Dias, R. Aylett, and A. Paiva, "Creating adaptive affective autonomous npcs," *Autonomous Agents and Multi-Agent Systems*, vol. 24, pp. 287–311, 2012.

[18] X. Zhu, “Behavior tree design of intelligent behavior of non-player character (npc) based on unity3d,” *Journal of Intelligent & Fuzzy Systems*, vol. 37, no. 5, pp. 6071–6079, 2019.

[19] H. Wan et al., “Building llm-based ai agents in social virtual reality,” in *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–7.

[20] J. Wei et al., “Chain-of-thought prompting elicits reasoning in large language models,” *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.

[21] “Meta Horizon,” Accessed: May 12, 2025. [Online]. Available: https://horizon.meta.com/.