Observations On Artificial Intelligence Applications in Non-Player Characters in Games

Xintong Liu *

School of Computing and Data Science, Xiamen University Malaysia, Sepang, 43900, Malaysia

* Corresponding Author Email: DMT2309230@xmu.edu.my

Abstract. This paper surveys the application of artificial intelligence (AI) in NPC behavior modelling development in video games. The motivation for this work comes from the desire to have more interactive, context-based NPC behavior in video games. It is a well-known fact that NPCs used to behave in predefined ways. But now, with the application of AI technologies such as large language models (LLM), reinforcement learning (RL), and imitation learning (IL), NPCs can now behave in response to the player's behavior and changes in the environment, thereby offering more natural interactions. The paper is organized into three main sections: NPC behavior evolution, NPC and player interaction, and application of AI approaches such as LLM, RL, and IL to make NPC more interactive. The contribution of this paper is to provide an overview of how AI is being applied in NPC behavior modelling in video games, making them more intelligent and emotionally engaging to players, and improving their gaming experience. The paper also addresses some of the challenges associated with current work and possible future works on refining NPC behavior such as hybrid models and multimodal learning.

Keywords: Large Language Model; Deep Learning; Reinforcement Learning; Imitation Learning; Non-player characters.

1. Introduction

In recent years, the development of artificial intelligence has changed a lot in many fields. The gaming field is also affected by it. From the most primitive NPCs can only respond in one way to the current ability to produce more complex emotions. Using AI can make games more interesting for players.

In the past, NPCs could only follow rules and behavioral patterns designed in advance. NPCs are not very flexible and diverse. But now with the help of reinforcement learning, imitation learning, deep learning, etc., the behavior and decision-making of NPCs have become increasingly intelligent. NPCs can make more human responses to the environment.

Reinforcement learning allows NPCs to repeat their trial-and-error experiments and perfect their decisions in simulation environments; Imitation learning allows NPCs to learn from players' behaviors and offer interactive situations almost the same as human beings; and deep learning allows NPCs to have a deep understanding of situations and make more various and realistic behaviors in games. Not only do these technologies enhance NPCs' intelligent decision-making capability, but also, they bring better gaming experience for users.

Whether NPCs are intelligent has become a significant criterion for evaluating a game. This article will explore the specific applications of AI in games, with a focus on analyzing how, through aspects such as the behavior simulation and emotional responses of intelligent NPCs, the interactivity between players and the game world can be enhanced.

2. NPCs' Basic Information, Evolution and Interaction

NPCs are characters controlled by a computer in the game. They interact with the Player Character (PC) in the game. The design and behavior of NPCs have an essential impact on the game experience because they are not merely the background element but often play multiple roles such as advancing the story, providing quests, facilitating interactions and conducting transactions.

2.1. History and Evolution of NPCs

In the early days of gaming, NPC design mostly relied on pre-defined script sequences and behavior trees, which outlined how characters would respond to specific player actions. These static algorithms worked well for uncomplicated interactions, for instance, in shooting games, NPCs were programmed to adhere to a pre-defined action tree. And the game adopts core parameters such as movement range, attack frequency to control NPCs [1]. However, this design was limited in its ability to capture the complexity and randomness of human actions. In The Witcher 1, the NPCs' reactions were based mainly on the player's decisions and the progress of the tasks, without any actual dynamic adaptation. Though players could undertake tasks and engage in conversations, such interaction tended to be linear and not essentially change the attitudes or actions of the NPCs.[2] With increasingly sophisticated game worlds, scripted NPC behaviors often became predictable and repetitive. It decreases the player immersion and feelings [3]. The biggest shortcoming of traditional NPC design was that it was very predictable, which resulted in boring gameplay. It caused the game to lack interesting and adaptive challenges. This caused developers to search for more flexible design methods, behavior-based and constraint-based designs in particular, to get NPCs to react more naturally to player behavior and to stay more believable in open worlds [3].

Nowadays, AI and behavioral simulation techniques enable NPCs to react to changes in the game environment and player behavior. Ubisoft has partnered with NVIDIA to bring the power of generative AI and large language models (LLMs) into the game in order to enable NPC communications with the player using natural language. Therefore, NPC can comprehend what the player said in voice or text and respond appropriately with proper context and emotion [4].

While playing The Witcher 3: Wild Hunt, Artificial Intelligence (AI) system is used in the game. Game NPCs can recognize the player's behavior based on visual, auditory and combat experience. For example, enemies can discover the player's hideout by the sound (if they are hiding) or chase the player by sight [5]. Compared with the traditional preset dialogue tree, the above methods make NPCs' response more diverse and personalized so the game is more interactive and interesting.

2.2. NPC-player Interactions

NPC-player interaction is a very important aspect in modern open-world games and contributes significantly to the game's dynamics and immersion. The following are some standard AI ways of NPC-player interaction and how such methods can be utilized in order to experience more natural NPC-player interactions.

2.2.1 Context-based Response System

There are several games that use an event-response system to control NPC responses, depending on the player's behavior, the world, and advancement through a mission. NPCs' responses are based on variables like the player's level of threat, friendliness, or carry. Such behavior is typically utilized through behavior trees or finite state machines (FSM), where NPCs perform pre-defined actions based on the world and mission. In The Witcher 3: Wild Hunt, moral choices affect NPC attitudes. For instance, helping or ignoring a trapped villager in the "Heart of Stone" quest leads to different outcomes, ranging from gratitude and rewards to hostility and retaliation later in the game [5].

2.2.2 Dialogue System

Advances in technology have enabled more and more NPCs in games to interact with players via complex dialogue systems. The choices made by the players not only affect the content of the conversation but also can turn the direction of the plot and actions of the NPCs. In "Cyberpunk 2077", players make choices through the dialogue tree to interact with NPCs, and those choices will directly impact the progression of the game's missions and the moods of the NPCs. For example, if the player chooses to say threatening words to some NPCs, they will end the conversation with the player and initiate a combat [6]. Such interaction will mostly use dialogue trees and natural language processing

(NLP) technology. Employing NLP, the NPC is able to identify voice input from the player and react according to it.

2.2.3 Dynamic Interaction Between the World and the Environment

Many modern games use a dynamic world so that NPCs and the environment and for the players more varied. NPCs not only respond to the players but also bond with other entities in the game world (such as weather, enemies, etc.). This is primarily achieved through sophisticated AI techniques such as behavior trees and environment perception systems, which enable NPCs to respond accordingly to changes in the world.

2.2.4 Generative Dialogue and Personalized Interaction

Generative dialogue systems are increasingly employed in more games in order to enable NPCs to generate more natural and customized responses based on what the player does. The system sometimes employs LLMs and deep learning, NPCs use these technologies to decipher the player's input and generate appropriate responses. Generative AI is also used to provide NPCs with more free and diverse conversation options, so they can be more realistic in their interactions with players [7].

2.2.5 Emotional and Memory Systems

In order to give the NPCs more “personality” and “emotions”, modern games use the emotional simulation and memory systems. The games use the system to let the NPCs remember the player and interact with the player differently depending on the past experience.

The two types of systems often involve reinforcement learning and memory retrieval to allow the NPCs to store information about the player’s choices and make different responses when interacting with the player later in the game.

In the game in-the-last-of-us-part-II, the player plays as Ellie and meets various NPCs, one of them is a good friend of Ellie- Dina. The experience of the Ellie and Dina, conversations, and the choices made by the player while playing, create the emotion of Ellie and Dina’s relationship. The choices made by Ellie whether to be cruel or to sympathize with the enemies, determine the behavior of Dina towards Ellie in the later scenes.

3. NPC Behavior Modeling

3.1. LLM-Based NPC Behavior Modeling

LLMs have changed NPCs' engagement with game players. Most conventional NPCs relied on preprogrammed responses and static dialogue trees, limiting interaction depth and realism. LLMs like Claude, GPT-4 and LLaMA can enable NPCs to write contextual replies that are dynamic, which raises the level of immersion for game players. In addition, LLMs can eliminate repetitiveness in games, especially when players go through multiple endings [8].

Transformer based model with self-attention mechanism, process one element at a time and capture the long-distance relationship naturally. NPCs can understand complicated dialogue and context. NPCs replies also sound more natural. LLMs are pretrained on gigantic corpus of text so that they can provide relevant contextual responses.

LLMs are very good at conversation dialog generation, fill in missing texts, answer questions, they are perfect for training models for NPC behaviors in computer games. There are several advantages of LLMs.

Firstly, Contextual Intelligence and Interactive Responses. The game player’s input is taken into account, and NPC responses are generated based on context, game player history, and environment.

Aslan stated in his research that they embedded an LLM into the game NPC dialogue system to make the dialogue more natural and response adaptation [9].

Next, Personalization and Character Depth. By utilizing LLMs, NPCs can have unique personalities, histories, and emotional reactions, providing more personalized and substantial interaction. This also enables players to bond with the character, impacting gameplay and storyline

outcome. In the game Where Winds Meet, players can chat with some specific NPCs, who generate messages through LLMs. After chatting, the relationship between the player and the NPC will change according to the text the player inputs [10].

After that, Cross-Platform Communication. LLMs provide cross-platform communication between NPCs and users across various platforms, e.g., in-game worlds and third-party apps like Discord. This allows for wider coverage for NPC communication, enabling ongoing communication beyond the game world. A prototype system showed the implementation of NPCs with LLMs that communicate with users both within a game based on Unity and on Discord, with coherent conversations across platforms [11].

3.2. Reinforcement Learning and Deep Learning-Based NPC Behavior Modeling

NPC behavior modeling based on reinforcement learning (RL) and deep learning (DL) improves the interactivity function in virtual game characters, making them more natural and intelligent. The latest technologies allow for the optimization and dynamic adaptation of behavioral strategies for NPCs, enabling them to demonstrate more complex and adaptive behaviors while they learn during their interaction with both the environment and users. As a result, this NPC interaction evolution creates a higher level of immersion for users, breaking free from the limitation caused by traditional NPC behaviors [12].

With reinforcement learning, the NPCs learn to improve their way by trial-and-error learning. DRL (Deep Reinforcement Learning) which uses deep learning (convolutional and recurrent layers) to work with high-dimensional inputs and to decide either by approximating the Q values or by directly optimizing the policy. Deep Q networks (DQN) provide better decision-making capabilities in face of deep environments by estimating the Q-value functions. Policy gradient methods optimize policy functions so that NPCs can choose the optimum actions in huge, continuous spaces of actions [13].

There are a lot of advantages to NPC behavior modeling with reinforcement learning and deep learning.

Dynamic and adaptive decisions. NPCs will adapt their behaviors in real-time based on feedback, providing more natural and contextually appropriate responses. A simple example is that, in a brawler game, NPCs will adjust their play style based on the user's play style rather than playing based on a pre-determined sequence.

Capability to generalize and variability. Deep learning enables us to deal with larger and more dynamic worlds for RL to operate in. NPCs with experience on large data sets exhibit variability in pattern of behavior which enriches the dynamics of games and make NPCs better under most circumstances.

Recent works like 2025 work "Dynamic NPC AI Using Reinforcement Learning for an Improved Gaming Experience" focus on these possibilities in deep learning and reinforcement learning to mimic sophisticated and advanced emotional and intellect responses on these non-player characters, and then extend their immersion and interaction capability [12].

3.3. Imitation Learning-Based NPC Behavior Modeling

Imitation Learning (IL) is nowadays a distinct way of providing a non-player character (NPC) behavior simulation in computer games. Traditional methods for modeling NPC behaviors are typically grounded in hand-designed rules and state machines. Despite this approach being efficient and simple, it is plagued with a rigidity and non-adaptability, which preoccupies deriving various and realistic NPC behaviors. By contrast, NPC modeling grounded on imitation learning produces more realistic and dynamic outputs in terms of behaviors by replicating expert behaviors such that the NPCs can adapt to varying player interactivity levels and change their behaviors based on contextual factors.

Imitation learning is a way to learn behavioral strategies via demonstrations by experts, and is well adapted for producing complex interactive behaviors. In modeling NPC behaviors, imitation learning is commonly comprised of two parts: behavior cloning (BC) and inverse reinforcement learning (IRL). In behavior cloning, the model is learned by mimicking the expert's input-output pairs (e.g., the player

control and NPC response), duplicating the expert's behavior directly. Whereas in inverse reinforcement learning, it interprets the expert's behavior and environment feedback to infer the potential reward function, and then utilizes this reward function to learn the NPC's decision-making strategy [14]. Behavior cloning is implemented with deep learning models such as convolutional and recurrent neural networks, with training based on vast datasets sampled from actual player data or simulated expert behaviors.

Here are the advantages of NPC behavior modeling with imitation learning.

Firstly, Realistic NPC Behaviors. Imitation learning makes NPC behaviors more lifelike by simulating human interaction patterns, allowing NPCs to exhibit natural emotional changes and logical responses in various situations. This improves the quality of interactions between players and NPCs.

Second is dynamic Response and Context Adaptation. Imitation-learning NPCs adjust their behaviors according to either player actions or dynamic environments, achieving more active and environment-aware reactions. NPCs in Red Dead Redemption 2 react to the actions of the player: a favorable chat can evoke thanks and a closer friendship, whereas harmful actions like robbing a store can result in hostility and wariness during future confrontations. By doing so, NPCs can deliver personable, dynamic reactions, encouraging immersion and interactivity in the game [15].

4. Limitations and Future Outlook

4.1. Current Constraints High Computational Requirements

The training and inference of the LLMs are compute-intensive. Running the LLMs in-game at real-time, and especially on weaker game clients, can lead to subpar game performance, such as lag and high-power requirements. Reinforcement learning and Deep learning suffer similarly.

Non-Realism: Even with advances in modern technology, true human-like realism is a challenge. Outputs by LLMs, although relevant in a given context, can be less than fully nuanced and less emotional than human interaction. In a few instances, the computer-generated speech can sound robotic or fall short on natural flow as found in human speech.

Integration Challenge: Integrating new NPC behavior models into existing game engines and pipelines is not an easy task. Programmers in a game have to make their code compatible with a number of game systems, such as graphics, physics, and user interface.

4.2. Future Perspective Hybrid Methods

A potential future direction for NPC behavior modeling is to integrate multiple methods. Hybrid behaviors can utilize LLM to infer player intention based on speech and then initialize the response actions based on a series of RL algorithms. The dual approach utilizes the best elements of both methods to achieve more complete and intelligent NPC behaviors.

Boosting Processing Power: With the development of new hardware technologies, the computation-intensive nature of current NPC behavior models can be overcome. Speaking of the next generations of GPUs and special-purpose AI chips, such as NPUs, to efficiently process the computation of LLM and RL models, will be in the making. In addition, software-level solutions, including model compression and quantization methods, that will also be part of the optimization, will be used to reduce the above-computational requirements of the models.

Adding to Realism with Multimodal Learning: Non-playable characters (NPCs) in the future will increase the realism by utilizing multimodal learning. This means that in addition to dialoguing with players through text, NPCs will also learn from other visual, sound, and so on information from the game environment. For instance, non-playable characters can make use of the visual clues (such as the body gestures or facial expressions of the character player) and sound clues (such as the tones) to have a deeper understanding of the players' emotional states and react more rationally. By using different types of information, non-playable characters can communicate with players in a more

organic and natural way, and the boundaries between face-to-face interactions in the virtual world and communications in the physical world will be erased.

5. Conclusions

This paper discusses the use of artificial intelligence in games. It describes NPC behavior modeling and interaction. It discusses the use of AI technology such as LLM, imitation learning and deep reinforcement learning. It enables the NPCs to interact dynamically with the environment and the actions of the player and improve the interactivity and immersion of the game.

This paper starts by discussing NPC from being simple pre-scripted behaviors to complex AI behaviors. This gives a point where an old game and a new game differs in their approach to NPC modeling. The review then proceeds to applying LLMs, DRL and imitation learning for NPC behavior modeling and discusses the advantages of each method along with the technology involved in it.

LLMs allow NPCs to generate dialogue that is dynamic and context dependent which improves the player engagement and immersion from transformer. Deep learning and reinforcement learning allows NPCs to adapt and optimize their behavior strategies dynamically with the help of Q-learning and gradient method. Furthermore, NPC modeling via imitation learning generates more natural and dynamic behavioral performance by using behavior cloning (BC) and inverse reinforcement learning (IRL) to imitate master behavior. The article also discusses the current limitations of NPC behavior models such as being computationally expensive and lack of realism in NPC dialogue. The current prospects like hybrid models and multimodal learning are discussed as the method to upgrade the NPC behavior further and make the game more enjoyable.

The significance of this paper is that it is an overall review of the work of AI in determining the NPC behavior, which gives a better idea of how the AI technology is changing the player experience and helps in building the next generation of games.

References

[1] Rivera G, Hullett K, Whitehead J. Enemy NPC design patterns in shooter games. In: Proceedings of the First Workshop on Design Patterns in Games, 2012 May: 1-8.

[2] CD Projekt Red. The Witcher [Video Game]. CD Projekt Red, 2007.

[3] Černý M, Gemrot J, Brom C, Bída M. Using behavior objects to manage complexity in virtual worlds. IEEE Transactions on Computational Intelligence and AI in Games, 2015, 7(4): 355–368. DOI: 10.1109/TCIAIG.2014.2365367.

[4] Ubisoft, NVIDIA. Generative AI and large language models in NPC interactions. Ubisoft, 2024. Available from: https://www.ubisoft.com

[5] CD Projekt Red. The Witcher 3: Wild Hunt [Video Game]. CD Projekt Red, 2015.

[6] Badowski A. Cyberpunk 2077 [Video Game]. 2020.

[7] Inworld AI. Generative AI and NPC interactions in gaming. Inworld AI, 2024. Available from: https://www.inworld.ai

[8] Brown T, Mann B, Ryder N, Subbiah M, Kaplan J D, Dhariwal P, et al. Language models are few-shot learners. Advances in Neural Information Processing Systems, 2020, 33: 1877–1901.

[9] Aslan P. Understanding LLMs for game NPCs: an exploration of opportunities. 2024.

[10] Everstone Studios. Where Winds Meet [Video Game]. Everstone Studios, 2024.

[11] sSong L. LLM-driven NPCs: cross-platform dialogue system for games and social platforms. arXiv, 2025. Available from: https://arxiv.org/abs/2504.13928

[12] Senanayake C, Campus H. Dynamic NPC AI using reinforcement learning for an enhanced gaming experience. 2025.

[13] Korkiakoski M, Sheikhi S, Nyman J, Saariniemi J, Tapio K, Kostakos P. An empirical evaluation of AI-powered non-player characters' perceived realism and performance in virtual reality environments. arXiv preprint, 2025. arXiv:2507.10469.

[14] Ng A Y, Russell S. Algorithms for inverse reinforcement learning. In: ICML, 2000 Jun, Vol. 1, No. 2: 2.

[15] Rockstar Games. Red Dead Redemption 2 [Video Game]. Rockstar Games, 2018.