

# Towards Sentient Companions: A Strategic Framework for Developing Emotionally Persistent, Morally Aware, and Adaptively Intelligent Non-Player Characters Using Large Language Models and Deep Reinforcement Learning

Rajan Kumar, Vaibhav Verma, Bishwajit Andia, Sakshi Singh

Department of Computer Applications  
Quantum University, Roorkee, Uttarakhand, India

**Abstract-** Non-Player Characters (NPCs) in Interactive Digital Environments have traditionally been limited by their artificial script, repetitive loops of actions, and primarily lack in psychological depth. The standards players have for NPCs are increasingly advanced, and the role of creating narrative is becoming a key quality indicator for games, making a level of memory, emotions, moral reasoning and strategic adaptations a requirement at this level. The authors introduce a comprehensive framework for research and development of four aspects of next-generation NPC intelligence: Emotional Memory and Relational Persistence, Moral Fatigue and Ethical Exhaustion, Adaptive Villain Intelligence, and Crowd Behavioural Authenticity. This research builds on the theories and computational paradigms of Large Language Models (LLMs) [4,12], Deep Reinforcement Learning (DRL) [13] and hybrid Behaviour Tree architectures [11] to create the theoretical and engineering underpinning needed to develop NPCs that go beyond the traditional role of a background character. Experimental comparisons conducted by [8] show that DRL-based antagonists can give rise to significantly higher player satisfaction and immersion scores when compared to Finite State Machine (FSM) based ones. In addition, the LLM-based crowd agents had better behavioural understandability scores according to the related crowd simulation studies. The paper ends with a few recommendations for researchers who wish to put forth original scientific work in the fast-changing field of Game Artificial Intelligence.

**Keywords—** Non-Player Character Intelligence, Emotional Memory, Moral Fatigue, Deep Reinforcement Learning, Large Language Models, Adaptive Villainy, Crowd Simulation, Behaviour Trees, Game AI.

## I. INTRODUCTION

If you've ever played a role-playing game and had to confront a character who seemed to remember your actions and respond to your choices, who morphed over time after having some moral dilemmas, you've seen a glimpse of the future of Non-Player Character (NPC) design. That initial sensation and how to consistently and reliably replicate it is exactly what this research paper is about. The ability to create believable game characters is not only a matter of art but also a question of architecture, as [11] state, requiring AI systems that are flexible, context-aware,

and tractable. NPCs have been around since the beginning of digital games. They range from the ghosts in Pac-Man to the residents of Cyberpunk 2077's Night City. They can be ghosts in Pac-Man, or the people of Night City in Cyberpunk 2077. However, for all but the past few decades they have been incredibly hollow, acting out pre-written dialogue, adhering to scripted dialogue pathways, and starting any game over again the next time a player logs off. Early NPC architectures were based on a Finite State Machine (FSM) and simple rule-based systems which create "predictable and repetitive" agents when a player has understood the

underlying pattern, as [11] state. What has made the situation worse is the surge of story-telling ambition for today's games. More than just winning, players want to feel seen, challenged, and emotionally connected, as [3] state. An antagonist that is pursuing the player in the same way each encounter is more of a mechanical challenge than an antagonist. The NPC that doesn't remember that you saved her life three quests ago feels more like furniture than a friend. In this paper, I will look at four aspects of next generation NPC intelligence that, together, form a complete model for the creation of realistic animated characters:

- Emotional Memory and Relational Persistence — NPCs that have a memory of the player across sessions and develop significant, dynamic relationships.
- Moral Fatigue and Ethical Exhaustion — NPCs that have been through a series of moral compromises; as humans have done. With adaptive villain intelligence, the player's weaknesses are learned by the villains and they adjust their tactics based on that knowledge.
- Crowd Behavioural Authenticity — Populations which behave realistically in the social scene as opposed to being background scenery.

All three pillars build on existing research in cognitive and moral psychology and on state-of-the-art research in computational Large Language Models (LLMs) [4] and Deep Reinforcement Learning (DRL) [13] [14]. The idea here is that the believable NPCs cannot be built using either a purely psychological approach, or a purely computational approach. domain: psychology without computation is theory without implementation, computation without psychology is experientially sterile. The paper also aims to provide a practical reference for organizing original research papers in the game AI space, including how to define a hypothesis and how to choose the proper methods, and how to select evaluation metrics that are used to measure the believability of an NPC as a human perceptual phenomenon and not only as a technical measure [9].

## II. LITERATURE REVIEW

The evolution of NPC intelligence research lines up with the evolution of AI generally, and its periods of high promise and low realization that it still has a long way to go to match human intelligence. The first NPC systems were based on some primitive constructs, like hard-coded finite state machines, rule based dialogue trees and basic pathfinding algorithms. As described at length in [11], such methods yielded agents that seemed to be intelligent – but were only as such when the player was unable to understand the underlying automaton. More modular and readable way to NPC decision-making was brought by the introduction of Behaviour Trees (BTs) in the middle of 2000s [11]. BTs enabled the designers to create complex behaviours using simple reusable leaf nodes. The Far Cry series and Halo 2 are two games that make heavy use of BT-based architectures. However, BTs are still very much scripted; they represent only those situations that were explicitly predicted by their designers and are unable to reason about situations that fall outside of their design scope. Deep Learning took the paradigm shift more fundamentally. With the breakthrough of [14] achieving super-human level performance on the game of Go with a reinforcement learning agent, the design of NPCs was re-thought from first principles. This was expanded on by [16] who showed that multi-agent reinforcement learning could be used to generate behaviours of impressively high tactical complexity in a more strategically interesting environment, StarCraft II. Could a similar approach create a villain that could learn each individual player's "tactics" and exploit them, if the AI could master the game itself to the level of grandmaster? The introduction of Stable and Sample Efficient Training by Proximal Policy Optimization (PPO) [13] has revolutionized the field of game AI, becoming a de facto standard. In addition to the Unity ML-Agents framework [8], PPO enables academic researchers to train adaptive NPC agents without having to rely on industrial level hardware.

More recently, the application of Large Language Models to NPC design has opened a whole new level of possibilities: the natural, context-aware,

generative conversation. Large Scale Language Models have been shown to display remarkable few-shot generalisation by [4] and further enhanced by GPT-4's significantly greater contextual coherence by [12]. These models, when applied to NPC dialogue, can produce next-to-natural sounding and situationally appropriate answers, which is a huge step up from static dialogue trees. The results of research on the PsyMem framework have also shown that adding structured psychological indicators to the LLM output can greatly enhance the consistency of NPC behaviour and its depth [9]. The psychologic grounds that are pertinent to such work are also very deep. In fact, as established by [2] using the Constructionist Model of Emotional Memory (CMEM), emotions are not pre-formed states, but constructions from perceptual, interoceptive and conceptual inputs. This is of great significance for the simulation of emotional memory in NPCs. The taxonomy of episodic, semantic, and procedural memory was offered by [15] which was the foundation of the architectural approach used in Section 3. Narrative memory layer in this framework is based on narrative identity theory introduced by [10]. In moral psychology, [5] formulated Moral Foundations Theory as a multi-faceted theory of moral thinking that aligns well with computational modelling. The five moral foundations were expanded upon by [6] in the empirical basis. The rationale and theoretical basis for the moral fatigue model described in Section 4 are offered in these works. The social force model proposed by [7] is still a reference for simulating crowd behavior, particularly in terms of the spatial navigation of simulated pedestrians. The affective computing dimension is based on the OpenFace 2.0 toolkit [1] that offers validated approaches for detection of facial action points and synchronisation of facial expressions [3] provide a psychological critique of NPC emotional realism, claiming that believability of expression, not expression of fidelity, is the main consideration in creating believable characters. This is because the specific combination mentioned in this paper – emotional memory, moral fatigue, adaptive antagonism and crowd authenticity, in one integrated phenomenon – is missing in the existing literature and needs to be addressed.

### III. TECHNICAL PARADIGMS IN AUTONOMOUS AGENT MEMORY AND AFFECTIVE PERSISTENCE

Discuss with any novelist what makes a character seem real, and he will probably mention memory. If the character can't remember what has just happened to them, they're not a character, they're a snapshot of them in one moment in time. The same also holds true for NPCs. As [15] has found in his taxonomy of memory systems, in order to have genuine relational intelligence, one needs to have more than one type of memory: episodic memory (of events), semantic memory (of the world), and the affective colouring (of the emotions). Traditional NPC architectures have been built on perhaps fairy-fingal reputation systems, where ratings are increased or decreased based on player actions [11]. They have a shallow depth and are still functional. They can't convey the emotional atmosphere of the relationship: the scars of betrayals of the past, the heat of shared success, the guarded optimism of the possibility of reconciliation. The depth of this skill needs a new kind of memory architecture that they call "multidimensional."

#### 1. Architecting Relational Intelligence

Drawing on the PsyMem framework and the psychological work of [15] and [10], we can identify five distinct layers that together constitute a psychologically grounded memory system for NPC relational intelligence:

Memory Layer	What It Stores	Psychological Basis	Data Representation
Episodic Memory	Specific shared events (e.g., player saved the NPC)	<sup>[15]</sup> episodic memory theory	Timestamped event log with context tags
Semantic Memory	Facts about the player and game world	Declarative knowledge systems <sup>[15]</sup>	Key-value knowledge graph
Affective Memory	Emotional quality of past interactions	Affective computing; <sup>[2]</sup> CMEM	Decay-weighted arousal score (-1.0 to 1.0)

Memory Layer	What It Stores	Psychological Basis	Data Representation
Narrative Memory	The life story connecting events meaningfully	Narrative psychology; [10]	Semantic narrative index
Predictive Memory	Anticipated future player actions	Anticipatory cognitive models	Probabilistic state projection

The critical innovation within this architecture is the affective memory layer. [2] argues compellingly that emotional states are not discrete pre-packaged responses but contextually constructed phenomena. Applied to NPC design, this means that an NPC’s emotional response to a player action should be modulated by its current state, its accumulated history, and the conceptual categories it has learned to apply to interpersonal events.

The PsyMem framework — which supplements LLM-generated text with twenty-six psychological indicators drawn from established models like the Big Five personality taxonomy [9] — operationalises this insight computationally. These indicators govern everything from how an NPC weights different categories of memory (some characters are more forgiving by disposition) to how quickly emotional valence decays over time. A naturally resilient NPC might recover from a betrayal within a few interactions; a deeply sensitive one might carry that wound indefinitely.

A particularly important implementation mechanism is what researchers term memory alignment training. Without explicit fine-tuning to align LLM outputs with the NPC’s established memory history, these models exhibit what [9] describe as a tendency towards contextual plausibility at the expense of historical consistency — generating responses that feel appropriate in isolation but contradict the NPC’s established past. Memory alignment training anchors the model’s outputs to its accumulated relationship data, ensuring coherence across sessions.

## 2. The Constructionist Model of Emotional Memory (CMEM)

In order to be able to create a memory system that is psychologically rich, a model of how emotional memories function in humans is necessary. The Constructionist Model of Emotional Memory (CMEM) (Barrett 2017) gives just this basis and its implications for NPC design are profound. The central principle of CMEM is that the emotions are not remembered from storage but are actively reconstructed when they are recalled. This reconstruction is based on three concurrent sources of input: exteroceptive (what the agent is sensing in his or her current environment); interoceptive (physiological signals of arousal); and conceptual (whether the agent can categorize sensations in terms of discrete emotion concepts, e.g., anger, joy, fear, or shame).

- Exteroceptive input: perceptions of the environment that set the context for what might be happening emotionally in the moment.
- Interoceptive input — internal signals of arousal, analogous to the agent’s physiological state.
- Conceptual knowledge — acquired categories that group sensations into named emotion types [2].

For game design, CMEM generates several powerful implications. First, it explains why NPC relationships should feel dynamic rather than static: the emotional tone of a memory can shift based on what has happened since it was formed [2]. Second, it suggests that NPCs should demonstrate variable memory vividness: the most emotionally charged interactions should be remembered most clearly, while mundane ones fade — a prediction consistent with the arousal-based memory enhancement effects documented in human memory research [15]. Third, it enables what researchers call emotional homeostasis: the tendency of relational systems to actively work towards stability, manifesting in NPC behaviours such as pre-emptively addressing tension or offering gestures of reconciliation. Implementing CMEM computationally requires building a mapping between the NPC’s stored interaction history, its current affective state (represented as a point in valence-arousal space, as established by [2]), and the conceptual emotion

categories it has learned to apply. The result is an NPC whose memory is not a database query but a living reconstruction — and whose responses feel accordingly more human.

- Authority/Respect — respect for authority and tradition [6].
- Purity/Sanctity — the sensitivity to degradation and moral contamination [5].

#### IV. MODELLING MORAL FATIGUE AND ETHICAL EXHAUSTION IN SYNTHETIC AGENTS

A player can get a weird feeling when playing a game where one of the NPCs is made to do something bad. The greater the sophistication of the game the more severe this discomfort is likely to be. Detroit: Become Human and The Last of Us Part II were the ones to cause genuine emotional controversy among players because of the way they let the player's character feel the pain of violence, but also the other characters within the game. These games suggest that the NPCs should have their own moral burden of compromise, but that's rarely done as much as it should be.

In clinical psychology, moral distress is when the person sees the morally right thing to do, but is unable to do it because of external factors [6]. This neatly translates to what many NPCs experience—the values they're designed to have are being overwritten by the player's actions or game's plot. The character realism is greatly improved if NPCs develop a moral agency which continually reacts to the social and contextual information, rather than only reacting to specific trigger conditions as [3] argue.

##### 1. The Psychology of Moral Distress in Games

Moral Foundations Theory (MFT) is an approach developed by [5] and empirically elaborated in [6] to model the moral landscape an NPC lives in in a computational way. MFT suggests that moral reasoning is based on five psychological foundations which can be breached by player actions to various extents:

- Harm/Care — ability to be sensitive to and responsive to suffering [5].
- Fairness/Reciprocity — being sensitive to justice and proportionality [5].
- Ingroup/Loyalty — commitment to the social group and protection of its members [6].

Each NPC may have a different weighting on each of these foundations, depending on their character. The knight NPC may focus on Authority and Ingroup, the rogue scholar may focus on Harm and Fairness. Importantly, each violation builds up over time, resulting in ethical exhaustion (Em). For this set of reasons, we propose the following mathematical model, inspired by the moral injury literature reviewed by [6]:

$$Em = \sum (w_i \times \delta_i) \text{ for } i = 1 \text{ to } n$$

The formula is Em = sum of the NPC's level of moral exhaustion (accounts for their character and disposition) x the weight of moral foundation i (accounts for their moral character and disposition) x the degree of violation to moral foundation i caused by player actions, (accounted for by the term  $\delta_i$ ) When the NPC's behavioural profile starts to vary from defined threshold values, it changes in predictable ways, in line with the literature on moral disengagement [5]:

Em Threshold	Behavioural Manifestation	Dialogue Tone	Relational Impact
0.0 – 0.25	Normal moral functioning; full empathic engagement	Warm, principled, engaged	Neutral to positive
0.25 – 0.50	Visible moral strain; reduced altruistic responses	Cautious, questioning	Slight withdrawal
0.50 – 0.75	Moral disengagement begins; cynicism emerges [6]	Bitter, deflective, dark humour	Significant relational distance
0.75 – 1.00	Severe ethical exhaustion; breakdown or abandonment	Nihilistic, hostile, or silent	Relationship rupture

## 2. Measuring and Simulating Ethical Decline

Careful design of measurement instruments is necessary in order to translate this psychological model into a working game system. Based on the reviewed experimental methods in moral

psychology [6] that can be applied to the Em calculation, we suggest four basic variables that influence the calculation of Em:

Variable	Definition	Measurement Method	Impact on Em
Utilitarian Proportion [6]	Ratio of self-sacrifice to goal-oriented acts forced on NPC	Logged decision tree outcomes	Drives base exhaustion rate
Affective Valence [2]	Pleasantness/unpleasantness of forced choices	Sentiment analysis of player-NPC interactions	Modulates decay rate
Mnemonic Arousal [15]	Emotional intensity of the decision moment	Physiological proxy from NPC model state	Increases memory salience
Time Pressure [5]	Constraint on NPC deliberation time	Decision window logging	Amplifies exhaustion multiplier

The variable time pressure should be noted especially. The findings of moral psychology research are consistently that under time pressure agents move from deliberative analytical reasoning to affect-driven responses [5]. If this is the case for an NPC, then the more it faces threat in the same high-pressure situation, the more "default" it becomes and the less "reflective" it is. A game that is intentionally forcing the NPCs to operate in a morally challenging situation ought to make them more inclined to act impulsively over time, whether it be as a result of the player's actions or not. This opens up a great design possibility: if a companion NPC turns bad, it can be a story arc all on its own, that the player has to fight against. It's emotionally true because of its parallels with human experiences of moral injury and burnout, studied in the occupational psychology literature[6].

## V. COMPUTATIONAL ARCHITECTURES FOR ADAPTIVE VILLAINY

What makes the most indelible character in any tale is the one who is the most personal, more than the

most powerful. A bad guy who is able to target the player's weaknesses, understand their strategy and do just that is much more terrifying and interesting than one who just goes for the draw. Adaptive villain intelligence is the goal of making an enemy who is not only a challenge, but also adapts to the unique player they are facing. The distinction between traditional boss AI and other AI is that it's built on specific patterns, which can be balanced, tested, and tuned. Emergent complexity, which is notoriously hard to control, can be found in a villain that truly adapts [11]. Today, however, there are tools to overcome this challenge and the results as shown empirically in [8] are worth the effort greatly.

### 1. Evolutionary Strategy and Reinforcement Learning

Deep Reinforcement Learning (DRL) is the computational method of developing adaptive villain intelligence. The paper [14] showed that RL agents can be made to perform better than humanly possible in strategically complex environments by self-play and optimisation based on reward. They then extended this result to multi-agent systems by showing that coordination in a multi-agent system

could emerge spontaneously based on the reward shaping alone, without explicit programming of individual strategies. [16] For NPC antagonist design, it takes a State-Action-Policy-Reward (SAPR) approach. The villain, at each time step, sees the current state of the game (player's position, health, weapons, etc.), chooses a move based on its current policy, and gets a reward signal based on the result. We chose the Proximal Policy Optimization (PPO)

algorithm [13] which is a sample efficient and stable policy gradient algorithm that is natively supported by the Unity ML-Agents toolkit [8] and is suitable for this application. Perhaps the most important design decision in the whole system is the design of the reward function [13]. We want to suggest a multi-component reward function whose form is as follows:

Reward Component	Description	Signal Type	Design Rationale
Tactical Success <sup>[13]</sup>	Successfully predicting and countering player movement	Positive	Rewards reading the player's habits
Resource Attrition	Forcing the player to consume defensive resources	Positive	Creates pressure without guaranteed defeat
Meta-Adaptation <sup>[14]</sup>	Successfully adjusting after player changes tactics	Strong Positive	Rewards genuine learning over exploitation
Predictability Penalty <sup>[8]</sup>	Being caught in repeated identical patterns	Negative	Prevents the agent from becoming boring
Damage Received	Taking damage from the player	Negative	Encourages self-preservation and tactical care

The reward aspect of the Meta-Adaptation is the most theoretically meaningful. We put pressure on the villain to learn how to learn—meta-learning, also called model-agnostic meta-learning as discussed in [14]—by providing a strong positive signal for changing its strategy when the player changes it for the better. Anyone who cares to develop a good meta-learning skill can start to adjust to a new player's style in a matter of minutes after meeting them. [8] empirically compares FSM-based and DRL-based villain agents and consistently shows better results in terms of player experience for the latter. Players who label DRL-based villains with terms such as 'feels like it knows me,' 'unpredictable in a satisfying way,' and 'actually threatening' have a clear reason to believe they are playing with a genuine evil character. Those who use words such as 'feels like it knows me,' 'unpredictable in a satisfying way,' and 'actually threatening' have a good reason

to think that they are playing with a real evil character. The distinction is not only a technical one, it's also perceptual and emotional, which is the key in immersive game design [3].

## 2. Hybrid Models and Strategic Depth

However, there are a number of known limitations for pure DRL, in particular its difficulty of ensuring high level narrative coherence [11]. A DRL demon could have a very clear idea of how to play a certain weakness at the micro level and be making quite unstrategical decisions at the macro level. The answer, as outlined in the game AI literature [11] [8] is a hybrid architecture, where strategies are detached from any specific implementation of tactics. We suggest a three-level model combining the benefits of Behaviour Trees, DRL and LLM [4] [12]:

Architecture Layer	System Used	Scope	Examples of Decisions
Strategic (Macro) — [11]	Behaviour Tree (BT)	High-level encounter management	Ambush vs. direct assault, retreat timing, resource use
Tactical (Micro) — [13]	Deep RL (DRL/PPO)	Real-time execution and adaptation	Attack timing, dodge patterns, target prioritisation
Narrative (Meta) — [4]	LLM with memory context	Personality and relationship with player	Taunts, dramatic escalation, backstory-relevant behaviour

In this hybrid, the Behaviour Tree is used as the villain's brain, deciding who should play what high-level gambit depending on the situation of the encounter. After the gambit is chosen, the DRL agent is responsible for the detailed execution by an adaptive approach [13] on a moment-by-moment

basis. The LLM layer makes sure that the villain's speech and actions are consistent with a character's personality and its past interactions with the player [4] [12].

AI Architecture	Strategic Coherence	Tactical Adaptability	Personality Consistency	Suitability
FSM [11]	High (scripted)	Very Low	Moderate	Low
Behaviour Tree [11]	High	Low	Moderate	Medium
Deep RL [13]	Low	Very High	Low	High (micro)
Hybrid BT + DRL + LLM	High	Very High	High	Optimal

## VI. MEASURING CROWD BEHAVIOURAL AUTHENTICITY IN OPEN WORLDS

When playing Open-World, the crowd is the world. The population reacts to you in sensible ways, depending on the context; the dialogue of people around you is contextually appropriate; collective behaviour during a crisis is believable – these are all things that can't be replaced by a photo-realistic render. However, as noted in [3] it has been a 'low priority technical problem': NPCs have minimal

budgets for behaviour, the pathfinding for collision avoidance is based on social force paradigm as in [7], and there is a generic ambient conversation that plays every 30 seconds. The outcome is what we can term as the uncanny valley of population, that is, a population that seems to be there, but is missing. But it takes more than this to move forward and rethink the meaning of 'authentic' crowd behaviour, and create more sophisticated measurement frameworks to differentiate it from mere 'graphical fidelity'.

### 1. Proposing a Believability Scoring Framework

The first hurdle to making crowds more authentic is the accurate measurement of the crowd. Current evaluation measures are mostly geared towards graphic fidelity or pathfinding efficiency (which are proxies for believability) that, as shown by [3] show,

do not correlate well with how players perceive NPC believability. We suggest a Believability Scoring Framework based on four dimensions which we believe to be psychologically relevant:

Dimension	Definition	Measurement Method	Reference
Understandability	Can the player 'read' what a crowd member is doing and why?	Player identification accuracy in user studies	[3]
Context-Awareness	Does the crowd member's dialogue fit the current situation?	Semantic relevance scoring (LLM-evaluated)	[4], [12]
Spatial Autonomy	Does the NPC observe realistic social distancing and navigation?	Social-force model simulation metrics	[7]
Emotional Synchrony	Do facial cues and body language match the NPC's speech and state?	Affective computing benchmarks	[11]

One notable result, in line with the study of [3] is that the crowd-motion generated by AI can perform better than crowds recorded by humans, on the dimension of Understandability. This is opposite to what you might think but it is true: Understandability is not accuracy, but clarity in communication. The purpose of the movement can be clearly sent by AI-generated movements, whereas human movements may not be easily understood or interpreted in isolation. This is a significant insight into the design philosophy because the object of the simulation of crowds is not to create human movements with the utmost realism. It should be to convey behaviour with as much clarity as possible – realism for readability as [3] state.

### 2. The Social Sandbox and Embodied AI

Aside from the measurement, the more profound conceptual change that needs to happen is from viewing a crowd as ambient decoration to viewing a crowd as a social sandbox, an interactive system where the actions of the player have real social consequences, which ripple through the population. The minimal Relational Cognition Layer consists of a lightweight module for each NPC that keeps track of the NPC's recent emotional state and has a simple opinion about the player based on observed behaviour, while also providing the ability to generate the NPC's response from the context. The use of few-shot generalisation capabilities of LLM shown by [4] and extended by [12] for crowd dialogue generation is a great stride. Instead of

having a list of pre-cropped ambient lines, the LLM-generated crowd NPCs can respond in a relevant way to the player's current behaviour and tone. This can be added to game engines through the use of platforms like Convai in Unreal Engine 5, for instance [9]. The non-verbal aspect of the dimension is provided by Embodied AI. The previous research in the field of affective computing has proven that the first clue to human's emotional state is through non-verbal communication, such as facial expression, body posture, gestural pattern, and proxemic behaviour [3]. In order to achieve a serious crowd authenticity system, the channels must be synchronised with validated tools like OpenFace 2.0 for facial action unit detection and using tools that blend the animation channels according to the sentiment, such as the OpenFace tool's blend shapes. The spatial aspect of crowd behaviour is modelled using the social force model [7] which has equations for pedestrian behaviour that have been validated, such as lane formation, density dependent flow and obstacle avoidance.

The aspirations of the framework outlined in this paper put significant pressures on the hardware and software infrastructure. In this section, a practical analysis of some of the technical decisions that researchers would have to make if they wanted to actually create these systems is presented, based on the comparative study of the performance of the various games engines documented in the literature [8].

### 1. Game Engine Selection

The selection of the game engine isn't just a choice of taste: it is a choice that affects different aspects of the research. Meanwhile, [8] introduced Unity as the new main platform for academic reinforcement learning research by releasing their ML-Agents toolkit, which is a native implementation of PPO [13] and contains a simplified API for agent-environment interaction. By contrast, Unreal Engine 5 delivers better graphic accuracy, essential for research that relies on the accurate generation of facial expressions and accurate crowd simulation on a large scale.

## VII. TECHNICAL INFRASTRUCTURE AND IMPLEMENTATION CONSIDERATIONS

Evaluation Criterion	Unity 2021.3 LTS	Unreal Engine 5	Research Implication
Physics Simulation	Higher resource consumption	More stable at scale	Critical for large crowd systems <sup>[7]</sup>
Rendering Fidelity	Better for simpler scenes	Superior for emotional facial detail	Relevant for CMEM & embodied AI <sup>[2]</sup>
ML Integration	Native ML-Agents / PPO <sup>[8]</sup>	External APIs (PyTorch)	Impacts speed of RL training cycles
LLM Integration	Moderate (REST API approach)	Strong (Convai platform native)	Relevant for dialogue research <sup>[4]</sup>
Hardware Demands	Lower threshold	High GPU/VRAM requirements	Influences institutional feasibility

Evaluation Criterion	Unity 2021.3 LTS	Unreal Engine 5	Research Implication
	(accessible to most labs)		

In fact, Unity's ML-Agents toolkit [8] offers the simplest development workflow for research that relies on reinforcement learning prototyping, including the adaptive villain architectures introduced in Section 5. In cases where emotional study or large-scale crowd simulation with OpenFace 2.0 pipeline [1] is required, the rendering of Unreal Engine 5 is worth the added resource usage.

## 2. Regional Research Infrastructure

IIT Roorkee's Machine Vision and Intelligence Lab has made key contributions to the field of Emotion AI and Affective Computing, creating datasets that classify personality traits and can be directly applied to the PsyMem framework detailed in Section 3. The University of Petroleum and Energy Studies (UPES) Dehradun has a dedicated Centre of Excellence (CoE) for AI driven technologies equipped with high end GPUs and TPUs clusters with the ability to provide the DRL training pipelines needed for adaptive villain research [8]. The researchers at these institutions are urged to seek opportunities to collaborate with game industry members to gain access to real player interaction logs, a type of data that isn't widely used in academic NPC research and could be a game-changer for the development of player-adaptive systems as both [14] and [16] indicate.

## VIII. DISCUSSION

All four pillars discussed in this paper are a vision of NPC intelligence that is qualitatively different from what is being done via traditional game AI design. However, it is important to note that this framework is designed to achieve certain ends, and it also poses some questions.

The most important unanswered question is: how scalable is this? Such architectures as described here, such as the multidimensional memory system based on Tulving's (1985) classification, and the LLM-based

dialogue system of [4] are high-cost. For each NPC in a scene, running full memory reconstruction via CMEM is at this time problematic. Practical implementation will be a delicate project that will probably be restricted to those NPCs that the player is interacting with so that they can be set up in a way that is truly believable, while other crowd NPCs are set in a more lightweight manner, in line with the proposed scoring for believability from Section 6.

There are also some design conflicts to be carefully managed by practitioners. It is possible that such an adaptive villain, with his ability to "learn" from his environment, as noted in [14] and [16], might be unfair and be the one to feel like he's been punished. Players must feel that the villain's adaptation is taking advantage of some aspect of the player's play that they could have avoided. If the player feels the AI is random, it is worse than if it was scripted [11].

On the other hand, if an NPC's moral fatigue occurs at a rate that's far too rapid, the player might feel penalised for playing the game normally. As with the Em thresholds, the rate of moral decay is not a technical issue but a design issue, and must be carefully tested with players from a variety of demographics — something that seems to be acknowledged by the methodological guidelines of [3].

The most profound philosophical question is how to make an NPC really feel the impact of a moral choice. The systems presented here emulates the so called behavioural and physiological correlates of moral fatigue, emulating it under the framework of the MFT model of [5] but in absence of any possible subjective experience. But it is important to recognise this honest difference between simulation and feeling distress, as both technical and ethical reflection on what we are creating [2]

## IX. CONCLUSION AND FUTURE DIRECTIONS

Tomorrow's NPCs will not be scripted. They will not be predictable. They will never forget. The framework outlined in this paper suggests a means to create characters that truly remember their history of interaction with the player, that sense the cumulative burden of ethical compromise modelled as in the MFT model of [5], that learn and adapt to their particular foe through the DRL approach of [14] and [13], and that live in game worlds that have the social realism of a believable crowd [7] [1] It's not science fiction. The constituent technologies are LLMs [4] [12] DRL [13] [16] Affective Computing [1] and Hybrid architectures [11] [8] are sufficiently advanced to be incorporated in these systems today. Whereas, rather than a lack of tools, there's a lack of integrating framework: a coherent theoretical explanation of what NPC intelligence should look like, based on both cognitive psychology and computational architecture. This provides several avenues for future research:

- Construction of a unified benchmark set of NPC believability across the four dimensions at the same time, so that the NPCs from different studies can be compared in the same framework as suggested in Section 6.
- Meta-learning for transfer learning: Investigate transfer learning methods that allow the villain AI to learn well from a set of players, and then apply this knowledge to a different set of players, with minimal additional training (drawing on the foundations of [14]).
- Longitudinal view of the emotional and narrative effects of moral fatigue on NPCs (as encouraged by [3]), rather than snapshots of the process (which tend to be misleading because of the different contexts and content of the encounters compared).
- Exploration of procedural narrative generation systems for generation of dynamically consistent backstories for crowd NPCs, that rely on generative abilities of [4] and [12].
- Design ethics issues (no precedent in the literature), such as how to design moral fatigue systems so that games do not cause second

order emotional distress to players when they are watching NPCs suffer.

For a computer scientist, this is one of the most compelling technical and humanistic frontiers of modern computer science, and it is also one of the most exciting challenges for the NPC, as it could be: the new NPC. The researchers developing work at this nexus, with the psychological depth of [2] and [10] and the computational power of [14] and [4] are not only creating better games, but they are creating better games in the first place. They are investigating the concept of simulating a mind.

## REFERENCES

All references follow APA 7th Edition formatting conventions.

1. Baltrušaitis, T., Zadeh, A., Lim, Y. C., & Morency, L.-P. (2018). OpenFace 2.0: Facial behaviour analysis toolkit. In Proceedings of the 13th IEEE International Conference on Automatic Face & Gesture Recognition (pp. 59–66). IEEE. <https://doi.org/10.1109/FG.2018.00019>
2. Barrett, L. F. (2017). *How emotions are made: The secret life of the brain*. Houghton Mifflin Harcourt.
3. Brom, C., Preuss, M., & Klima, M. (2012). Are video game characters useful models for studying human emotional reactions? *Entertainment Computing*, 3(4), 119–127. <https://doi.org/10.1016/j.entcom.2012.09.003>
4. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., ... Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
5. Haidt, J., & Graham, J. (2007). When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research*, 20(1), 98–116. <https://doi.org/10.1007/s11211-007-0034-z>
6. Haidt, J., Graham, J., & Joseph, C. (2009). Above and below left–right: Ideological narratives and moral foundations. *Psychological Inquiry*, 20(2–

- 3), 110–119.  
<https://doi.org/10.1080/10478400903028573>
7. Helbing, D., & Molnár, P. (1995). Social force model for pedestrian dynamics. *Physical Review E*, 51(5), 4282–4286.  
<https://doi.org/10.1103/PhysRevE.51.4282>
8. Juliani, A., Berges, V.-P., Teng, E., Cohen, A., Harper, J., Elion, C., Goy, C., Gao, Y., Henry, H., Mattar, M., & Lange, D. (2020). Unity: A general platform for intelligent agents. *arXiv preprint arXiv:1809.02627*.  
<https://arxiv.org/abs/1809.02627>
9. Jurafsky, D., & Martin, J. H. (2023). *Speech and language processing (3rd ed. draft)*. Stanford University.  
<https://web.stanford.edu/~jurafsky/slp3/>
10. McAdams, D. P. (2001). The psychology of life stories. *Review of General Psychology*, 5(2), 100–122. <https://doi.org/10.1037/1089-2680.5.2.100>
11. Millington, I., & Funge, J. (2009). *Artificial intelligence for games (2nd ed.)*. CRC Press.
12. OpenAI. (2023). GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.  
<https://arxiv.org/abs/2303.08774>
13. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.  
<https://arxiv.org/abs/1707.06347>
14. Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484–489.  
<https://doi.org/10.1038/nature16961>
15. Tulving, E. (1985). Memory and consciousness. *Canadian Psychology/Psychologie Canadienne*, 26(1), 1–12. <https://doi.org/10.1037/h0080017>
16. Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Choi, J., Powell, R., Ewalds, T., Georgiev, P., Oh, J., Horgan, D., Kroiss, M., Danihelka, I., Huang, A., Sifre, L., Cai, T., Agapiou, J. P., Jaderberg, M., Vezhnevets, A. S., ... Silver, D. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782), 350–354.  
<https://doi.org/10.1038/s41586-019-1724-z>
17. Khaleel Khan Mohammed. "A Survey on Digital Health Care Data Analysis Techniques for Developing Machine Learning Models"