

# EPHYRA Whitepaper

## ECA: A Cognitive Architecture for Autonomous Non-Player Agents in Dynamic Virtual Environments

### Abstract

This whitepaper details a novel computational framework named the EPHYRA Cognitive Architecture (ECA), designed to address the intelligence limitations of Non-Player Character (NPC) behavior in current interactive digital entertainment. Traditional NPCs, which primarily rely on pre-programmed scripts like Finite State Machines (FSMs) or Behavior Trees (BTs), exhibit highly predictable behaviors and lack contextual memory and emotional depth. This constrains the immersiveness and narrative emergence of virtual worlds. The EPHYRA architecture aims to catalyze a paradigm shift from "script executors" to "autonomous intelligent agents" by introducing a multi-layered cognitive model for Non-Player Agents (NPAs). This architecture integrates a hierarchical memory system, an emotion computation model based on the OCC theory, a goal-oriented deliberative planner, and a dynamic dialogue generator based on Large Language Models (LLMs). This paper will elaborate on the design principles of the architecture, its core technical implementations, and its intended applications, and propose a framework for evaluating its effectiveness. Its core contribution lies in constructing a computational framework that enables agents to produce believable, coherent, and intrinsically motivated behaviors, and to collaborate with Procedural Content Generation (PCG) systems to create truly dynamic and evolving worlds.

### 1. Introduction

#### 1.1 Problem Statement: The Paradigm Limitations of Current Non-Player Agents

In the fields of computational narrative and interactive entertainment, Non-Player Characters (NPCs) are a core element for building the world, driving the plot, and enhancing the player experience. However, despite significant advances in graphical fidelity and physics simulation, the Behavioral Intelligence of NPCs has largely stagnated within a deterministic, reactive paradigm. Their underlying technologies, such as Finite State Machines (FSMs), Behavior Trees (BTs), and Utility Systems, essentially pre-map an agent's behavior space onto a discrete set of designer-defined rules.

This paradigm leads to several intractable constraints:

- **Limited and Predictable Behavior:** The set of agent responses is closed and finite. Players can fully predict their behavior patterns after a limited number of interactions, which significantly reduces the world's sense of mystery and long-term engagement value.
- **Contextual Amnesia:** The vast majority of NPCs lack persistent memory of past interactions. They cannot form long-term, dynamic cognitive models of the player, leading to a breakdown in narrative continuity and character relationship development.
- **Lack of Intrinsic Motivation and Autonomy:** Traditional NPCs are passive entities

whose actions are driven by external triggers (the player's presence, specific events). They do not possess intrinsic goals, needs, or motivations, and thus cannot engage in meaningful autonomous activities independent of the player. The social ecosystem they form is essentially static.

- Breaks in Immersion: When a player's actions go beyond the scope of a pre-set script, the NPC's response fails or becomes illogical, immediately breaking the "Suspension of Disbelief" and pulling the player out of their immersive state.

## **1.2 The Core Proposition and Research Goals of EPHYRA**

We propose the core proposition of EPHYRA: by endowing Non-Player Agents (NPAs) with a cognitive architecture that simulates the core functions of the human mind, a qualitative leap from pre-scripted narratives to Emergent Narratives, and from passive interaction to active co-existence, can be achieved.

To validate this proposition, the following goals have been set:

- Design and formalize a multi-layered cognitive architecture that integrates perception, memory, emotion, and deliberative decision-making processes.
- Implement a hierarchical dynamic memory system that can support context-relevant short-term memory and semantics-based long-term episodic memory retrieval.
- Integrate a computational emotion model that allows an agent's emotional state to evolve dynamically based on events and interactions, and in turn, regulate its cognition and behavior.
- Develop a dialogue generation subsystem based on Large Language Models (LLMs) that is conditioned on the agent's internal cognitive state (memory, emotions, intentions) to generate contextually consistent and personalized natural language dialogue.
- Explore the synergistic mechanisms between this cognitive architecture and Procedural Content Generation (PCG) systems to achieve agent-driven dynamic world evolution.

## **2. EPHYRA Cognitive Architecture (ECA)**

The EPHYRA Cognitive Architecture (ECA) is a modular software framework inspired by mental models from cognitive science and classic agent architectures from the field of artificial intelligence (such as the BDI model: Belief-Desire-Intention).

### **2.1 Architecture Overview**

ECA is composed of three core subsystems: the Perception Subsystem, the Cognitive Core, and the Behavior Generation Subsystem. The data flow begins with perception, is processed by the Cognitive Core to form behavioral intentions, and is finally translated into concrete actions in the virtual world by the Behavior Generator.

### **2.2 Perception Subsystem**

The function of this subsystem is to translate raw, low-level data streams from the virtual world engine (such as position coordinates, collision events, sound signals) into

structured, symbolic representations that the agent's Cognitive Core can understand. The types of information it processes include:

- Object and Entity Recognition: Identifying specific players, other agents, and key items within the field of view.
- Event Detection: Capturing discrete events, such as "an attack occurred," "an item was picked up."
- State Monitoring: Tracking environmental states (time, weather) and its own physiological states (health, stamina).
- Linguistic Information Parsing: Utilizing a Natural Language Processing (NLP) module to extract semantic information such as intent, entities, and emotional polarity from dialogue.

The output is a set of updated Beliefs, which constitute the agent's representation of the current world state.

## **2.3 Cognitive Core**

The Cognitive Core is the central nervous system of ECA, responsible for processing beliefs, updating internal states, and making deliberative decisions.

### **2.3.1 Hierarchical Memory System**

To overcome "contextual amnesia" and simulate a more biologically plausible memory mechanism, we have designed a hierarchical system that clearly distinguishes between different types of memory.

#### **Memory Type Distinction:**

- Episodic Memory: Stores autobiographical, spatio-temporally contextualized events ("what I experienced"). For example, "At timestamp T, at location L, player A attacked me with a fireball, which triggered my 'anger' emotion with a value of 0.8." This type of memory is crucial for forming an individual history and for contextualized decision-making.
- Semantic Memory: Stores non-personal, factual knowledge about the world ("what I know"). For example, "Fireballs cause fire damage," "Player A's class is a Mage." Semantic memory forms the basis of the agent's understanding of the world's rules and concepts and is generally more stable.

#### **Memory Processing Flow:**

- Ingestion & Importance Scoring: Not all events from the Perception Subsystem are recorded. An evaluation function,  $I = f(\text{salience}, \text{emotional\_impact}, \text{goal\_relevance})$ , calculates an importance score for each event. Only events exceeding a threshold  $\theta_I$  are encoded as episodic memory traces.
- Encoding & Storage: Episodic memory traces are encoded into structured data containing (timestamp, event\_vector, participants, location, emotion\_vector) and stored in a long-term storage based on a vector database. Semantic knowledge is

stored as (subject, relation, object) triplets in a knowledge graph.

- **Abstraction & Generalization:** To prevent episodic memory from growing infinitely, an offline process periodically reviews episodic memories. Through cluster analysis, the system can abstract new semantic knowledge from multiple similar episodic memories. For example, multiple episodes of "Player A gave a healing potion" can be generalized into a new piece of semantic knowledge: "(Player A, has\_trait, helpful)."
- **Decay Strategy:** Each episodic memory trace is associated with an activation value  $A$ . This value decays exponentially over time  $t$ :  $A(t) = A_0 * e^{(-\lambda t)}$ , where  $\lambda$  is the forgetting constant. When a memory is successfully retrieved or associated with a new event, its activation value is boosted. Memories with an activation value below a threshold  $\theta_A$  will be archived, thus simulating the process of forgetting and maintaining retrieval efficiency.
- **Semantic Retrieval Mechanism:** When the agent needs to make a decision, the current context (including recent events, dialogue content, current goal) is encoded into a query vector  $V_q$ . The system retrieves the most relevant memories by calculating the cosine similarity between  $V_q$  and each memory vector  $V_m$  in the episodic memory store:

$$\text{Similarity}(V_q, V_m) = (V_q \cdot V_m) / (\|V_q\| \|V_m\|)$$

The top  $K$  most relevant memories retrieved will serve as contextual information, fed into the deliberative planner and dialogue generator.

### 2.3.2 Affective State Module

We adopt the Ortony, Clore, and Collins (OCC) model as the theoretical basis for emotion computation. This model defines emotions as valenced reactions to events, agents' actions, and objects' attributes. For example, an event that is congruent with the agent's goals will elicit "joy," while an incongruent event will elicit "distress." A negative event attributable to another agent will elicit "reproach" or "anger." The emotional state is a continuously varying vector that is not only an output of behavior (e.g., facial expressions) but, more importantly, serves as an internal heuristic that modulates cognitive processes, such as influencing attentional allocation, memory retrieval bias, and decision-making risk preference.

### 2.3.3 Motivational & Deliberative Subsystem

This module endows the agent with autonomy.

- **Motivational Layer:** Each agent is endowed with a set of basic motivations based on its character archetype (e.g., a simplified model based on Maslow's hierarchy of needs: survival, safety, belonging, esteem). These are the long-term, intrinsic "Desires" that drive behavior.
- **Deliberative Planning Layer:** Based on current beliefs, emotional state, and long-term motivations, the deliberative planner is responsible for generating one or

more specific, executable short-term "Intentions." We employ techniques such as Goal-Oriented Action Planning (GOAP) to enable the agent to autonomously plan a sequence of actions that can achieve its most pressing current goal. For example, a guard agent with a low "safety" motivation might generate the intention to "patrol a weak area" or "interrogate a suspicious person."

## **2.4 Behavior Generation Subsystem**

This subsystem translates the "intention" output from the Cognitive Core into concrete implementations in the virtual world.

- **Action Scheduler:** Decomposes the planned action sequence (e.g., GoTo(LocationA) -> Interact(ObjectB) -> Say(DialogueC)) into atomic commands executable by the engine. It employs a variation of a dynamic behavior tree, whose structure and parameters can be modified in real-time by the output of the Cognitive Core.
- **Dialogue Synthesizer:** This is the key to achieving natural interaction. We utilize a Large Language Model (LLM) fine-tuned for a specific character or world setting. Unlike simple text generation, our LLM is conditioned on the agent's complete cognitive state. The prompt engineering includes relevant episodes retrieved from the memory module, the current emotional state vector, and the intention being executed. Through this "Chain-of-Thought" style of conditioned generation, we ensure that the dialogue is not only fluent and natural but also logically, emotionally, and narratively consistent with the agent's internal mental state.

## **3. Implementation Methods and Technical Innovations**

### **3.1 LLM Dialogue Generation Conditioned on Cognitive State**

Through Instruction Fine-tuning of a base LLM, we enable it to understand and respond to input formats containing structured cognitive states (e.g., [Emotion: Joy=0.8, Anger=0.1], [Memory: "Player saved me"], [Intention: "Express gratitude"]). This method transforms dialogue generation from ungrounded text continuation into a well-founded cognitive expression.

### **3.2 Acquiring Complex Behaviors via a Hybrid Learning Paradigm**

To achieve high believability and effectiveness in agent behavior (especially in complex scenarios like combat), we adopt a two-stage learning paradigm:

- **Behavioral Cloning:** Using player data collected from the "Destiny of Gods" test, we pre-train the agent's behavior policy network through Imitation Learning, allowing it to master the basic tactics and operational patterns of human players.
- **Online Policy Optimization:** Building on the pre-trained model, we use Reinforcement Learning (such as Proximal Policy Optimization, PPO) for extensive self-play in a simulated environment. This allows the agent to explore and optimize its strategies, and even discover super-human level solutions.

### **3.3 Agent-driven Procedural Content Generation (Agent-driven PCG)**

We establish a bidirectional feedback loop between the cognitive architecture and the PCG system. The agent's internal state (especially unmet motivations and newly generated goals) can act as a trigger, requesting the PCG system to generate new content. For example, a merchant agent whose "wealth" motivation is frustrated due to a disrupted trade route can trigger a dynamic quest generator to create a "caravan escort" or "monster cleanup" quest. The existence of this quest, in turn, becomes part of the world state, perceived by other agents, thus triggering new social dynamics.

#### **4. Application and Evaluation Framework**

To conduct a rigorous empirical validation of ECA's effectiveness, we have designed an operational evaluation framework that includes a control group, explicit metrics, and standardized statistical analysis.

##### **4.1 Experimental Design**

- Design Paradigm: A double-blind, between-subjects A/B testing design will be used.
- Participants: N=100 qualified gamers will be recruited as subjects and randomly assigned to the experimental group (N=50, interacting with ECA-NPAs) and the control group (N=50, interacting with NPCs based on traditional scripts and behavior trees). The two versions of the game environment will be identical in terms of graphics, quest framework, etc.
- Procedure: Each subject will participate in one continuous 4-hour gaming session. During this time, the system will automatically record all interaction logs, behavioral data, and in-game events. After the session, subjects will be required to complete a series of questionnaires.

##### **4.2 Evaluation Metrics and Measurement Tools**

###### **Metric 1: Behavioral Complexity**

- Measurement Tool: Behavioral Entropy ( $H(B)$ ). We will discretize the NPA's behavior set into  $B = \{b_1, b_2, \dots, b_n\}$  (e.g., patrol, talk, attack, flee, trade, etc.). Over the entire 4-hour session, we will calculate the frequency  $p(b_i)$  of each behavior.
- Formula:  $H(B) = - \sum_i [p(b_i) * \log_2(p(b_i))]$
- Hypothesis: The  $H(B)$  of the experimental group will be significantly higher than that of the control group, indicating that its behavior patterns are richer and less predictable.

###### **Metric 2: Narrative Emergence**

- Measurement Tool: Emergent Narrative Chain Analysis. Two trained, independent evaluators will code the game logs to identify and count "emergent narrative chains." A chain is defined as a sequence of events initiated by an NPA's autonomous action (not directly triggered by the player), containing at least 3 causally linked steps, and having a measurable impact on the world state or player experience.
- Statistical Period and Baseline: A full analysis will be conducted after the 4-hour session. The number of emergent narrative chains in the control group (baseline) is

expected to be close to zero.

### **Metric 3: Social Structure Complexity**

- Measurement Tool: Dynamic Social Network Analysis. We will model the interactions between NPAs (e.g., dialogue, cooperation, conflict) as a dynamic graph  $G=(V, E)$ .
- Statistical Period and Baseline: A network snapshot analysis will be performed every 30 minutes to calculate network density and average clustering coefficient. We will compare the evolution curves of these metrics over the 4-hour period between the experimental and control groups (baseline).
- Hypothesis: The social network metrics of the experimental group will exhibit more complex and dynamic evolutionary patterns.

### **Metric 4: Player-Perceived Believability and Immersion**

Measurement Tools:

- Agent Believability Scale (ABS): Adapted from existing academic scales, including dimensions such as reactivity, autonomy, and emotional expression.
- Immersive Experience Questionnaire (IEQ): A standardized psychological scale.

### **4.3 Statistical Analysis**

- For continuous variables such as behavioral entropy and questionnaire scores, an independent samples t-test (if data meets normality and homogeneity of variance assumptions) or a non-parametric Mann-Whitney U test will be used to compare group differences.
- For discrete variables such as the count of emergent narrative chains, Poisson regression or a Chi-squared test will be used.
- The significance level ( $\alpha$ ) for all statistical tests will be set at  $p < 0.05$ .

## **5. Limitations and Future Research Directions**

- Computational Overhead: Running hundreds of complex cognitive models in real-time within a single environment poses a significant challenge to computational resources. Future research will focus on model quantization, distributed computing, and selectively simulating fidelity.
- Balancing Authorial Intent and Agent Autonomy: A key design challenge is how to maintain the game's core narrative framework and authorial intent while ensuring a high degree of agent autonomy.
- Controllability and Safety of LLMs: Ensuring that the content generated by LLMs always conforms to the game's world-building, rating requirements, and does not produce harmful information is an "Alignment" problem that requires continuous research.
- Future Directions: Include exploring more complex cognitive functions (such as Theory of Mind), extending the architecture to large-scale multiplayer environments, and conducting longitudinal studies on the formation of long-term emotional bonds between players and agents.

## **6. Conclusion**

The EPHYRA Cognitive Architecture, as articulated in this whitepaper, represents a fundamental challenge to the current paradigm of agent design in interactive entertainment. By simulating key functions of the human mind—memory, emotion, and deliberation—we aim to elevate Non-Player Characters from passive world decorations to active, intrinsically alive participants. This framework not only provides a technical path toward creating unprecedentedly deep immersive experiences and dynamic emergent narratives but also contributes valuable insights to the broader research field of Human-Agent Interaction. We believe that the "agent-centric" design philosophy represented by ECA is the inevitable path toward the next generation of intelligent interactive entertainment.