TAME THE BALROG: TASK-ADAPTIVE MODULAR EMERGENCE FRAMEWORK FOR GAME AGENTS

Anonymous authorsPaper under double-blind review

000

001

002

004

006

008 009 010

011 012

013

014

015

016

018

019

021

023

025

026

027

028

029

031

033

035

037

038

040

041

042 043

044

046

047

048

051

052

ABSTRACT

Interactive games have proven to be key benchmarks for advancing Artificial Intelligence (AI), requiring capabilities like long-term planning, exploration, and adaptation to stochastic environments. While Large Language Models (LLMs) have achieved notable results across many domains, they struggle in complex gaming environments like those in the BALROG benchmark. The absence of adaptive frameworks that can dynamically configure themselves based on environmental characteristics, limits the progress of AI in games. To this end, we introduce the Task-Adaptive Modular Emergence (TAME) framework, which employs genetic algorithms to evolve environment-specific structures from modular components, enabling significant performance improvements of LLMs across diverse domains. TAME discovers high-performing configurations by selecting between baseline and hierarchical structures, selectively incorporating specialised modules, and fine-tuning each component through systematic mutations. Evaluating TAME across the BALROG benchmark, we find that the emergent modular structures discovered by TAME significantly enhance LLM performance, raising average progression scores of Gemini 2.0-Flash from 27.15% to 34.77%. Moreover, these structures demonstrate transferability across models. Directly employing TAME discovered structures for Gemini-2.0-Flash to a population of Gemini-2.5-Pro, we achieve new state-of-art performance on BALROG. This transferability suggests that TAME identifies fundamental structural principles for game-playing agents that adapt their cognitive architecture to match task demands.

1 Introduction

Large language models (LLMs) have achieved remarkable growth across a wide range of tasks, from general language understanding (Hendrycks et al., 2020) and code generation (Wang et al., 2024a; Pan et al., 2025; Hong et al., 2024), to recent breakthroughs including mastering the ARC reasoning benchmark (Chollet, 2024; Chollet et al., 2024) and performing at gold-medal level on International Mathematical Olympiad (Chervonyi et al., 2025). However, these models struggle significantly in interactive decision-making environments that require sequential actions, state awareness, and long-term planning (Liu et al., 2024; Klissarov et al., 2025).

Interactive games have historically served as major testbeds for artificial intelligence, with examples including Atari (Mnih et al., 2013), Starcraft (Team, 2019), or GrantTurismo (Team & Digital, 2022). Those successes predominantly emerged from reinforcement learning (RL) approaches specifically engineered for each domain, often requiring millions of training episodes and domain-specific reward shaping. While LLMs hold considerable promise on the possibility of zero-shot generalisation across games through their vast pretraining experience, e.g., game wikis, strategy guides, and gameplay discussions, they fail to translate this latent knowledge effectively. This performance gap is clearly illustrated in the BALROG benchmark (Paglieri et al.), a suite of diverse games traditionally employed in RL research, where even state-of-the-art LLMs achieve only partial success in the simpler games and barely progress with more challenging ones.

Notably, structured agentic frameworks have emerged as a dominant approach to enhancing LLM capabilities in other complex domains. In software development, frameworks like SWE-Agent

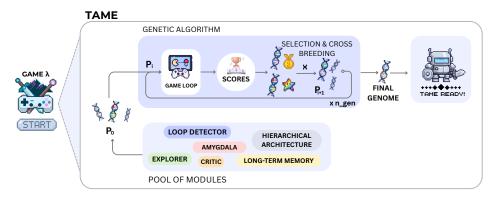


Figure 1: TAME framework overview. TAME runs a genetic algorithm to generate a population of genomes representing different module combinations, fine-tunes them, and selects the genome achieving the highest score. Icons were generated with Google Gemini Pro 2.5, 2025.

(Yang et al., 2024), MetaGPT (Hong et al., 2024), and CodeCor (Pan et al., 2025) enable LLMs to decompose complex programming tasks through specialised modules that maintain project context or check for domain-specific challenges such as correct compilation or indentation. Similarly, in scientific research, systems like ChemCrow (Bran et al., 2023), MathCoder (Wang et al., 2024a) and the AI Scientist (Lu et al., 2024) augment LLMs with relevant skills for their respective applications, including hierarchical planning to tackle multi-step experimental design, or automatic peer-reviewing. There are also recurrent efforts to improve long-term memory management in agentic frameworks, with solutions that prioritise either speed and cost efficiency, such as Jarvis-1 (Wang et al., 2024b), or performance, such as A-mem (Xu et al., 2025). Yet, despite presenting similar challenges, no agentic framework has been applied to games. Unlike math or code, gaming environments are partially observable, evolve quickly, and demand long-horizon strategies that rely on persistent memory and planning, capabilities that remain difficult for LLMs.

To address this gap, we introduce the TAME (Task-Adaptive Modular Emergence) framework, a genetic-driven approach for LLMs to automatically discover and configure effective agentic structures for diverse gaming environments. TAME consists of a series of human-designed modules that enable different capabilities that might be relevant in games. Examples of that capabilities include hierarchical planning, exploration or long-term memory. However, unlike in other domains, games can encode very diverse dynamics, e.g., Nethack is a game where exploration, long-term planning and memory are of paramount importance, while none of those skills help in a game like TextWorld. Thus, TAME undergoes an evolutionary process, iteratively exploring the best structure for a specific game. Figure 1 illustrates this process. Each candidate structure is encoded through a vector that represents which modules are activated, along with the hyperparameters and the prompts selected for this candidate. Through successive generations, TAME employs mutation and cross-over operations on the genomes of the selected candidates to discover increasingly effective structures, balancing performance and diversity in their selection.

We validate TAME through extensive experiments on the BALROG benchmark. Our results show that TAME-discovered configurations improve overall progression scores by 28% compared to the baseline LLM performance, including progress in the most challenging domains such as Nethack. Moreover, we show that architectures discovered by TAME exhibit strong transferability: genomes evolved using Gemini-2.0-Flash directly enhance the performance of Gemini-2.5-Flash-Lite and Gemini 2.5-Pro without additional adaptation. In the case of Gemini-2.5-Pro achieving new state-of-the-art results on BALROG. Through ablation studies, we further demonstrate the individual contributions of our long-term memory system and the effectiveness of the genetic adaptation.

We summarise our contributions as follows: (1) we introduce TAME, the first emergent agentic framework that enables LLMs to evolve modular structures tailored to gaming environments; (2) we propose a novel and effective long-term memory system that combines embedding-based retrieval with LLM-augmented semantic memory (3) we demonstrate TAME's strong performance on the BALROG benchmark, with a new state-of-the-art system; (4) we show that TAME-discovered architectures transfer effectively, enhancing models without additional adaptation.

2 RELATED WORK

Prompting and Memory. Hallucinations remain a key challenge for LLMs (Kalai et al., 2025), which can be mitigated through prompting techniques like chain-of-thought (Wei et al., 2022) and step-by-step reasoning. Limited context windows (Brown et al., 2020) is another limitation, driving development of memory systems. Retrieval-Augmented Generation (RAG) (Lewis et al., 2021) combines LLMs with external document retrieval to reduce hallucinations without retraining. HiAgent (Hu et al., 2024) manages hierarchical memory using subgoals, dividing into "working memory" and "cross-trial memory" with LLM-based observation summarisation, while Park et al. (2023) balance memory retrieval using recency, importance, and relevance scores. Jarvis1 (Wang et al., 2024b) stores task names, plans, and observation sequences using embedding (CLIP) for encoding and retrieval. A-mem (Xu et al., 2025) introduces structured memory notes with timestamps, keywords, and embeddings, establishing inter-memory connections through LLM calls.

LLMs as hierarchical planners. In TWOSOME (Tan et al., 2024), LLMs score actions based on observations, allowing RL agents to leverage world knowledge for improved decisions. MaestroMotif (Klissarov et al., 2024) uses LLMs to generate reward functions for skills, while LLM-Augmented Hierarchical Agents (Prakash et al., 2023) use LLMs to inject commonsense priors for more efficient policy learning. Jarvis1 (Wang et al., 2024b), consists of planner and controller, enhanced by multimodal memory system. An important limitation of Jarvis-1 is the necessity of human-crafted goals based on specific skills, limiting its application in games with emerging tasks.

Agentic frameworks. Agentic frameworks are systems that enable Large Language Models to act as autonomous agents capable of reasoning, planning, and interacting with external tools and environments. Recent work, such as AGENTBREEDER Rosser & Foerster (2025), shows that optimising frameworks provides superior multi-agent performance on reasoning, mathematics, and safety benchmarks. Moreover, multiple works show improvements in scientific discovery (Lu et al., 2024) and software development (Yang et al., 2024) through these structured frameworks.

Evolutionary Strategies. Recent work incorporates LLMs into evolutionary frameworks. Lehman et al. (2022) propose "evolution through large models" using LLMs as evolutionary operators. Evo-Prompt (Guo et al., 2025) employs LLMs for crossover and mutation in genetic algorithms to discover diverse prompts, while Rainbow Teaming (Samvelyan et al., 2024) mutates adversarial prompts to populate MAP-Elites archives systematically. DOMiNO (Zahavy et al., 2022) balances quality-diversity trade-offs using Lagrange multipliers. Eureka (Ma et al., 2023) shows evolutionary optimisation over reward code benefits from human initialisation.

3 TAME FRAMEWORK

We introduce TAME (Task-Adaptive Modular Emergence framework), a novel agentic framework designed for dynamic LLM adaptation across diverse gaming environments. Inspired by Eureka (Ma et al., 2024), which shows that human priors significantly improve LLM-based evolutionary optimisation performance, TAME begins with an initial population \mathcal{P}_0 comprising diverse modular structure configurations, each encoding different combinations of human-crafted modules and hyperparameters. The framework's modular architecture consists of six core components illustrated in Figure 2: hierarchical goal decomposition (comprising a Meta-Controller, Low-Level Executor, and Completion Validator), Long-Term Memory, Critic, Loop Detector, Amygdala, and Explorer. TAME's evolutionary process operates iteratively: in each generation, the framework evaluates all new members $p \in \mathcal{P}_i$ on the target game and selects candidates for the next generation based on two criteria: (1) the top N performers by absolute score, and (2) M additional diverse solutions that achieve at least a fraction α of the best performer's score. This dual selection strategy balances exploitation of successful structures with exploration of the relevant solution space. Each candidate's genome is represented as a vector encoding active modules, hyperparameters (e.g., memory decay rates, exploration-exploitation trade-offs), and module-specific prompts. After every iteration of TAME, the genomes of the selected candidates undergo mutation and crossover operations to generate the new members of the subsequent population \mathcal{P}_{i+1} . Through successive generations, TAME discovers increasingly effective structures tailored to each game's requirements. The remaining of this section details the genetic algorithm and the design and functionality of each modular component in TAME.

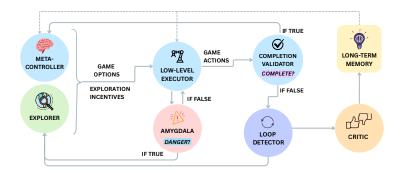


Figure 2: Fully enabled TAME modular structure: an epsilon-greedy mechanism selects between Meta-controller (providing options toward game objectives) and Explorer (options towards exploring the environment). Low-level Executor proposes actions, while Amygdala checks for danger and prioritises survival. Completion Validator checks option completion while the Loop Detector identifies stuck states, and the Critic summarises key actions that led to the outcome of the option. Long-term memory saves successful and failed trajectories and adds them to LLM context.

3.1 NOTATION AND FUNCTIONAL INPUTS

Let $\mathcal{G} = \{g_1, \dots, g_n\}$ the set of decomposed options towards the game objective, A the action space, O the observation space, S the state space, S the set of option summaries, and S, space of natural text. Each state S at time S

$$s_t = (o_t, I, h_{i-1}, g_i, \tau_i, M_i^{\pm}, F_{\text{recent}}, F_i, F_{i,act}, F_{i,obs})$$

where $o_t \in O$ is the current agent's observation and I provides game context (objective and available actions), $h_{i-1} \in H$ summarises the previous option's outcome, $g_i \in \mathcal{G}$ is the active option with termination condition τ_i, M_i^\pm contains the top successful k_{succ} and failed k_{fail} options. F_{recent} stores the last m action-observation pairs, while $F_i, F_{i,actions}$, and $F_{i,obs}$ maintain action-observation, action-only, and observation-only trajectories for the current option respectively.

Each component within s_t can serve as a functional input to any module within the framework. Through early experimentation we noted that the selection of appropriate functional inputs has a significant impact on performance, and it is up to the evolutionary process to find the most appropriate inputs for each module. For the initial population generation, we hand-crafted inputs that we consider most relevant for each module, providing the genetic algorithm with informative human priors. This is further detailed in Section 3.2.

3.2 GAME ADAPTATION

Given the diverse dynamics and requirements across different gaming environments, we propose a genetic optimisation approach that automatically explores diverse agentic architectures to adapt the underlying LLM to the specific game characteristics. To that end, we first encode the agentic structure descriptors into genomes. Each genome is constructed from three core components:

- Modules: TAME modules (See Section 3.3) are parametrised as a set of binary values. Such values indicate if a module is active (1) or not (0). The set of modules include: Hierarchy (enables or disables hierarchical goal decomposition, includes Meta-controller, Low-level executor and Completion Validator), Long-Term Memory (stores past experiences), Critic (summarises key actions toward the current option), Amygdala (activates survival mode), Loop Detector (detects looping behaviour), Explorer (controls exploration strategies).
- Hyperparameters: A set of continuous values encoding: Long-Term Memory Time Decay Factor λ , which sets the option priority decay rate; Long-Term Memory Similarity Threshold τ_{cos} , which specifies the cosine similarity cutoff for storing memories; Exploration Parameter ϵ , which controls the epsilon-greedy exploration; and Language Model Temperature.

• **Prompts**: The last component of the genome includes the prompts used by the modules of the agentic structure.

For any game G the initial population \mathcal{P}_0 consists of four predefined genomes:

$$\mathcal{P}_0 = \{\mathbf{g}_{\text{basic}}, \mathbf{g}_{\text{hierarchical}}, \mathbf{g}_{\text{default}}, \mathbf{g}_{\text{full}}\}$$

Where g_{basic} is genome corresponding to a baseline structure (no modules activated) with a single prompt as in Paglieri et al. (BALROG). $g_{hierarchical}$ corresponds to TAME[Hierarchical+Long-Term Memory], a genome with the hierarchical and long-term memory modules activated. $g_{default}$ - refers to TAME[Full Structure], a genome with all the modules activated that employs prompts proposed by Claude-Sonnet-4. Finally, g_{full} is TAME[Full Structure] is also a genome with all the modules active, but with our own engineered prompts. See Appendix O for full detail.

Genetic Operations. Given the distinct nature of the components within the genome (i.e., binary or continuous values, or prompts) TAME employs distinct crossover and mutation strategies depending on the type of variable or representation that is handled. For binary variables, we use a parent-based probabilistic flipping mechanism that incorporates an inheritance bias. Continuous variables are handled through Gaussian perturbation for mutation and linear interpolation for crossover. Finally, for prompt optimisation, we adopt the EvoPrompt methodology by Guo et al. (2025), which enables crossover and mutation tailored to LLM-based prompts (availabe in Appendix N).

Genome Evaluation. Each genome is evaluated using the average game progression across n_{ep} episodes as a fitness function. Moreover, we embed the genome representation and calculate the minimum distance to genomes already existing in the population as a score of diversity, alowing for keeping population of best scoring and most diverse genomes.

Genetic Algorithm. Our genetic algorithm iterates through four steps: 1) Parent Population Selection: TAME chooses parents based on wheel selection 2) Reproduction: Parents for reproduction are chosen from parent population with p_{single} and $1-p_{single}$ referring to the probability of single-parent or two-parent, respectively. Genetic operations are applied, represented by mutation+crossover or mutation alone (based on the number of parents) 3) Fitness Evaluation: calculates genome's performance score and diversity based on average game progression and embeded genome representation, respectively. 4) Population Pruning: the population is trimmed to maintain a maximum of N+M individuals — N highest-performing plus M most diverse genomes (subject to achieving a factor α of the performance of the best genome within the population). Most diverse genomes are those with largest minimum distance of their embedding with respect to the already existing embeddings in the population. Further detail and discussion can be found in Appendix G, together with pseudocode and hyperparameters in Appendix G.3 and C, respectively.

3.3 MODULAR BLUEPRINTS

As anticipated through Section 3, inspired by how Eureka improved its ability to find better reward functions by starting from a human-crafted set of prior, TAME incorporates a set of human-crafted modules that target essential capabilities for agents in interactive games. The remaining of this section details such components. We remind the reader than s_t^{mc} , s_t^{lc} , s_t^{cv} , s_t^c , s_t^a , s_t^{ld} used in sections below, represent subsets of s_t selected by the genetic algorithm for corresponding modules.

3.3.1 HIERARCHICAL PLANNING

The hierarchical module consists of three main components illustrated in Figure 2: Meta-Controller suggests sequence of options, Low-Level Executor performs a sequence of actions towards each option, and Completion Validator judges if an option has been completed successfully or failed.

Meta-Controller. The Meta-Controller decomposes the game objective into a more manageable sequence of options. Specifically, it implements $\pi_{high}: S \to \mathcal{G}$, mapping the current state to an ordered sequence of options:

$$g = \pi_{high}(s_t^{mc}) = \text{LLMprompt}_{high}(s_t^{mc}) = (g_1, g_2, \dots)$$

$$\tag{1}$$

Each option g_i consists of the fields: name, description, prerequisites, success conditions, penalty component, progress indicators, estimated priority.

Low-level Executor. This system implements $\pi_{low}: S \to A$, producing an action sequence based on the current state information provided s_t^{le} :

$$a = \pi_{low}(s_t^{le}) = \text{LLMprompt}_{low}(s_t^{le}) = (a_1, a_2, \dots)$$
(2)

where the length of the sequence is decided by the Low-level Executor.

Completion Validator. The Completion Validator implements the binary classifier $\varphi: S \to \{0,1\}$, determining whether an option has been completed:

$$C_i = \varphi(s_t^{cv}) = \text{LLMprompt}_{\varphi}(s_t^{cv}) \in \{0, 1\}. \tag{3}$$

Here $C_i = 1$ indicates successful termination. For details on the hand-crafted LLM prompts used as initial seeds we refer the reader to Appendix O.

3.4 Long-term memory

TAME implements a novel memory system that seeks to leverage the cost and speed efficiency of embedding-based systems like Jarvis-1 (Wang et al., 2023) while achieving a performance closer to more complex systems like A-mem (Xu et al., 2025). To that end, our system adopts Jarvis-1's storage framework, maintaining option information including name, description, prerequisites, success conditions, progress indicators, penalty components, and observation sequences. We extend this with two key additions: (1) Critic Ilm-generated summaries highlighting key success/failure actions, and (2) success/failure classification labels obtained from Completion Validator. This enhancement provides actionable guidance for future tasks requiring a single LLM call while avoiding A-mem's computational overhead of three LLM calls for memory and link creations, and evolutions. We now provide further detail of how TAME's long-term memory works:

Critic. The Critic module is a function $\rho: S \times \{0,1\} \to \mathcal{T}$, mapping the state and recent option outcome from Completion Validator to text:

$$h_i = \rho(s_t^c, C_i) = \text{LLMprompt}_{\rho}(s_t^c) \tag{4}$$

where the text aims to summarise the key factor that led to the success or failure of the option.

Creation of the memory. Each memory entry is defined as:

$$M_i = \{g_i, C_i, o, h_i\} \tag{5}$$

where g_i is the option, $C_i \in \{0,1\}$ is the output of the Completion Validator, o is the observation sequence towards current option, and h_i is the option summary from the Critic. Then, the memory structure is implemented as follows: each memory entry is stored as a vector embedding of the above, enabling efficient similarity-based retrieval. The embedding function ϕ transforms each entry:

$$\mathbf{e}_i = \phi(M_i) \in \mathbb{R}^d \tag{6}$$

Similarity-Based Filtering. Following the "importance" scoring approach from the generative agents framework (Park et al., 2023), we prevent storage of repetitive experiences. A new memory M_{new} is stored only if:

$$\max_{M_i \in \mathcal{M}} \frac{\mathbf{e}_{new} \cdot \mathbf{e}_i}{||\mathbf{e}_{new}|| \cdot ||\mathbf{e}_i||} \le 1 - \tau_{cos}$$
(7)

where τ_{cos} is a set constant. This method effectively filters out frequently repeated actions (e.g., "chop wood" in Crafter) that provide limited learning value.

Long-Term Memory Retrieval Mechanism. In order to address the limited context window size, we only extract $k_{succ} + k_{fail}$ best scoring memories at each option execution. Inspired by Retrieval-Augmented Generation (RAG) (Lewis et al., 2021), we enable access to past experiences through the following steps:

• Query Encoding: The new option name (n_{g_i}) and description (d_{g_i}) are embedded into the same vector space as stored memories using sentence embeddings:

$$\mathbf{q}_i = \phi(n_{q_i}, d_{q_i}) \in \mathbb{R}^d \tag{8}$$

• *Temporal Decay*: Following A-mem (Xu et al., 2025), we prioritise recent experiences using exponential decay:

$$w(e_t) = \exp(-\lambda \cdot t) \tag{9}$$

where t is the time elapsed since memory creation.

• Memory score: We combine similarity and recency through weighted sum:

$$score(\mathbf{q_i}, \mathbf{e_j}) = w_{similarity} \cdot sim(\mathbf{q_i}, \mathbf{e_j}) + w_{recency} \cdot w(\mathbf{e_j})$$
 (10)

• Stratified Retrieval: The system retrieves top-k successful and failed memories:

$$M_i^+ = \text{Top-}k_{succ}(\text{score}(\mathbf{q_i}, \mathbf{e}_j) : \varphi_i = 1)$$
(11)

$$M_i^- = \text{Top-}k_{fail}(\text{score}(\mathbf{q_i}, \mathbf{e}_i) : \varphi_i = 0)$$
(12)

where φ_i indicates success/failure of memory j.

• Context Integration: All $k_{succ} + k_{fail}$ retrieved memories are integrated into the modules prompts (modules including memory are decided by genetic algorithm).

This process gives access to both effective strategies and failure patterns, allowing for informed decision-making. Visualisations of retrieval patterns are shown in Appendix E.

3.5 SKILL-SPECIFIC MODULES

On top of the hierarchical structure, we identify *survival* and *exploration* as two key components in many video games. Moreover, we identify looping behaviour as a significant LLM limitation. All three modules are illustrated in Figure 2.

Explorer. Let explorer be defined as a function $\pi_{explorer}: S \to \mathcal{G}_{explore}$, where $\mathcal{G}_{explore}$ is the set of exploration-oriented options, and:

$$g^{exp} = \pi_{\text{explorer}}(s_t^e) = \text{LLM}_{\text{explorer}}(s_t^e)$$
(13)

where $g^{exp} = \{g_1^{\exp}, g_2^{\exp}, \dots, g_k^{\exp}\}$ is the sequential exploration plan, and each g_i^{\exp} is structured identically to regular options but prompted for discovery rather than game goal completion.

• Exploration Strategy We implement an ϵ -greedy exploration strategy where the Metacontroller selection becomes:

Controller
$$(s_t) = \begin{cases} \pi_{\text{explorer}}(s_t^e) & \text{with probability } \epsilon_t \\ \pi_{\text{high}}(s_t^{mc}) & \text{with probability } 1 - \epsilon_t \end{cases}$$
 (14)

with
$$\epsilon_0 = 0.1, \epsilon_t = 0.99 \times \epsilon_{t-1} = 0.99^t \times \epsilon_0$$

Amygdala. Let $\sigma: S \to \{0,1\}$ be the amygdala function mapping observations to binary classification of danger assessment:

$$D_i = \sigma(s_t^a) = \text{LLM}_{\text{amygdala}}(s_t^a)$$
(15)

At each Low-Lever Executor step, if $D_i = 1$, the system immediately activates a "survival option" (see Appendix F.1 for details); otherwise, normal execution continues.

Loop Detector. The loop detector implements $\psi: S \to \{0,1\}$, detecting repetitive behavior in recent execution history:

$$L_i = \psi(s_i^{ld}) = \text{LLMprompt}_{\psi}(s_i^{ld}) \in \{0, 1\}. \tag{16}$$

Where $L_i = 1$ means looping behaviour is detected.

4 EMPIRICAL EVALUATION

We evaluate our method through three key experiments. First, we benchmark our genetic algorithm with Gemini-2.0-Flash against the SOTA systems on the BALROG benchmark. Second, we demonstrate the transferability of TAME's selected genomes across different Gemini models without additional training. Third, we compare our memory system against Jarvis-1 and A-mem baselines.

4.1 TAME RESULTS

This section compares the baseline and TAME's performance on the BALROG benchmark. Baseline scores are obtained by evaluating the BALROG repository with Gemini-2.0-Flash following the original author's methodology. TAME scores represent the best performance achieved selected by our genetic algorithm (Section 3.2). We run genetic algorithm through $n_{gen}=4$ iterations, with each iteration producing $n_{child}=5$ children. Through empirical evaluation we notice that gives sufficient performance gains. Number of episodes per each child evaluation is adapted from BALROG.

Environment	Baseline[2.0-Flash] (†)	TAME (†)	Episodes	$\mid \Delta$ Full Pop. Score(\uparrow)
Average	$27.16\% \pm 2.12\%$	$34.78\%\pm2.22\%$	-	+12.18%
babyai	$58.00\% \pm 6.98\%$	$72.0\% \pm 6.35\%$	50	+6.69%
babaisai	$30.83\% \pm 4.22\%$	$41.67\%\pm4.50\%$	120	+10.00%
textworld	$32.55\% \pm 6.95\%$	$32.55\% \pm 6.95\%$	30	+26.47%
crafter	$29.09\% \pm 4.51\%$	$39.09\% \pm 4.9\%$	10	+19.09%
minihack	$12.50\% \pm 5.23\%$	$22.5\%\pm 6.6\%$	40	+10.00%
nle	$0.00\% \pm 0.00\%$	$0.85\% \pm 0.47\%$	5	+0.85%

Table 1: Average game progression of baseline LLM and TAME (with Gemini-2.0-Flash base). Results also include the scores with all modules activated and human-crafted prompts (full) to illustrate the impact of the genetic adaptation algorithm on the framework.

As shown in Table 1, TAME consistently outperforms the baseline achieving relative gain of \sim 28%. TAME improves performance in five out of six games, with the same score on TextWorld, as it is the environment where original BALROG's paper approach was chosen by the genetic algorithm. Notably, while the baseline model cannot achieve any noticeable progress on Nethack (the hardest game) TAME achieves 0.85% average score (note that the best model scores 1.8% on Nethack). In order to access the performance gain of genetic algorithm, we compare it against human-crafted Full Structure. We notice the average performance gain of 12.18%, illustrating the importance of adaptation in games. Detailed per-task results are provided in Appendix I along with an analysis of module activations in Appendix H. The final genomes returned by genetic algorithm are available in Appendix P with further results of the performance of initial population \mathcal{P}_0 in Appendix M.

4.2 Transferability of TAME structures

Next, we evaluate whether architectures evolved with Gemini-2.0-Flash can be effective when transferred to other models. Thus, we use TAME selection to evaluate populations of Gemini-2.5-Flash-Lite and Gemini-2.5-Pro models to exlusively choose between the base configuration from BAL-ROG or the best-performing structure discovered with Gemini-2.0-Flash for each game.

Method	Score (†)	BALROG Rank (↓)
Gemini-2.5-Pro[Transferred]	$47.65\% \pm 2.20\%$	$(1) \uparrow 1$
Grok-4	$43.60\% \pm 2.20\%$	1
Gemini-2.5-Pro[Baseline]	$43.35\% \pm 2.3$	2
Gemini-2.0-Flash[TAME]	$34.78\% \pm 2.22\%$	(4) ↑ 8
Gemini-2.0-Flash[Baseline]	$27.16\% \pm 2.12\%$	(12)
Gemini-2.5-Flash-Lite[Transferred]	$22.76\% \pm 1.73\%$	$(13) \uparrow 10$
Gemini-2.5-Flash-Lite[Baseline]	$11.87\% \pm 1.32\%$	(23)

Table 2: Comparison of TAME against top scoring models in BALROG leaderboard (September 2025). We show how they would rank (in parenthesis) relative to the current leaderboard. Rank improvements are indicated with \uparrow .

From Table 2 presenting the results, we observe that Gemini-2.5-Flash-Lite achieves almost 100% improvement. Detailed analysis in Table 6 in Appendix J demonstrate that TAME's discovered structures successfully transfers in four out of six environments, with only TextWorld and NetHack achieving baseline performance. For Gemini-2.5-Pro, we also observe gains although more moderate. Table 7 in Appendix J shows that the transferred structures significant improvements in

the BabyAI and BabaIsAI environments, which require extensive planning, highlighting the framework's strengths in this domain. However, improvements are not observed in the remaining environments. We hypothesise that Gemini-2.5-Flash-Lite benefits more substantially because it is a non-reasoning model similar to Gemini-2.0-Flash, where we carried the optimisation, whereas Gemini-2.5-Pro is a reasoning-based models. Notably, transferring TAME's discovered genomes to Gemini-2.5-Pro we achieve state-of-art performance above the best model on the BALROG leader-board - Grok-4. Similarly, we see large improvements on the leaderboard for Gemini-2.5-Flash-Lite and Gemini-2.0-Flash with TAME, now occupying rank 13 and 4 from 23 and 12 respectively.

4.3 ABLATION: MEMORY TYPES

We also include ablations to demonstrate the effectiveness of our long-term memory system. In order to test memory, we use the hierarchical structure described in Section 3.3.1, combined with three different memory architectures: Jarvis, TAME-Memory[ours] and A-mem. Both **Jarvis** memory and **A-mem** store the same core elements: g_i (the option, including all information associated with it), $C_i \in \{0,1\}$ (the status indicator), and o (the observation sequence corresponding to the current option). TAME extends Jarvis framework by introducing a critic, as well as a filter for successful and failed trajectories (but does not create links between memories). This requires one additional LLM call compared to Jarvis, but two fewer LLM calls per generation compared to A-mem. Thus, our approach explores a trade-off between the simplicity of Jarvis and the more complex and expensive structure of A-mem.

We notice an improvement compared to Jarvis and A-mem as shown on Table 3, motivating the integration of critic module for memory storage. Moreover, we achieve this while requiring a third of the LLM calls that A-mem employs. Thus allowing our system to iterate faster and with a reduced compute cost. Further details are included in Appendix L.

Environment	Jarivs (†)	TAME-Memory[ours] (†)	A-mem (†)
Average	$17.52\% \pm 1.73\%$	$23.11\%\pm1.75\%$	$21.45\% \pm 1.80\%$

Table 3: Comparison of average game progression across 6 games using different memory types.

5 DISCUSSION AND CONCLUSION

We presented TAME, a genetic framework for evolving LLM-based agents that is both game-agnostic and adaptive. Through genetic mutations and in-game evaluation, TAME configures human-crafted modules for core gaming skills such as exploration, survival, long-term memory, and loop detection. Its novel memory system combines the efficiency of embedding retrieval with the contextual depth of LLM-augmented memory. To our knowledge, this is the first application of such a memory design in this domain.

We evaluated TAME on the well-established BALROG benchmark and find that it consistently enhances the underlying LLMs. Gemini-2.0-Flash improves from 27.16% to 34.78%, while solutions discovered on one architecture transfer training-free to others, with Gemini-2.5-Pro reaching 47.65% and outperforming the state-of-the-art. These results showcase both the generalisability of the core modules and the effectiveness of our genetic approach. We further confirm the importance of long-term memory and adaptive architecture, with our proposed memory system outperforming two existing baselines while remaining more cost-efficient than complex agentic systems.

We note some limitations. We find that TAME provides greater benefits to some games than others, where it defaults to the baseline architecture. We also observed that while TAME improves complex reasoning tasks overall, spatial reasoning remains a weakness. This suggests the potential not only for expanding the set but for genetic discovery of entirely new modules and capabilities, beyond those hand-crafted in this work. Moreover, while transferability proved effective, gains were less pronounced for reasoning models, motivating further study of transfer and emergence across different architectures.

Overall, TAME establishes a new state of the art in game-playing LLM agents, laying the foundation for more efficient and emergent frameworks.

REFERENCES

- Andres M Bran, Sam Cox, Oliver Schilter, Carlo Baldassari, Andrew D White, and Philippe Schwaller. Chemcrow: Augmenting large-language models with chemistry tools. arXiv preprint arXiv:2304.05376, 2023.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020. URL https://arxiv.org/abs/2005.14165.
- Yuri Chervonyi, Trieu H Trinh, Miroslav Olšák, Xiaomeng Yang, Hoang Nguyen, Marcelo Menegali, Junehyuk Jung, Vikas Verma, Quoc V Le, and Thang Luong. Gold-medalist performance in solving olympiad geometry with alphageometry2. *arXiv preprint arXiv:2502.03544*, 2025.
- François Chollet. Openai o3 breakthrough high score on arc-agi-pub. https://arcprize.org/blog/oai-o3-pub-breakthrough, December 2024. Accessed: 2025-09-12.
- François Chollet, Mike Knoop, Gregory Kamradt, and Bryan Landers. Arc prize 2024: Technical report. arXiv preprint arXiv:2412.04604, 2024. URL https://arxiv.org/abs/2412.04604.
- Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujiu Yang. Evoprompt: Connecting Ilms with evolutionary algorithms yields powerful prompt optimizers, 2025. URL https://arxiv.org/abs/2309.08532.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.
- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Ceyao Zhang, Jinlin Wang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, et al. Metagpt: Meta programming for a multi-agent collaborative framework. In *International Conference on Learning Representations, ICLR*, 2024.
- Mengkang Hu, Tianxing Chen, Qiguang Chen, Yao Mu, Wenqi Shao, and Ping Luo. Hiagent: Hierarchical working memory management for solving long-horizon agent tasks with large language model, 2024. URL https://arxiv.org/abs/2408.09559.
- Adam Tauman Kalai, Ofir Nachum, Santosh S. Vempala, and Edwin Zhang. Why language models hallucinate, 2025. URL https://arxiv.org/abs/2509.04664.
- Martin Klissarov, Mikael Henaff, Roberta Raileanu, Shagun Sodhani, Pascal Vincent, Amy Zhang, Pierre-Luc Bacon, Doina Precup, Marlos C. Machado, and Pierluca D'Oro. Maestromotif: Skill design from artificial intelligence feedback, 2024. URL https://arxiv.org/abs/2412.08542.
- Martin Klissarov, R Devon Hjelm, Alexander T Toshev, and Bogdan Mazoure. On the modeling capabilities of large language models for sequential decision making. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=vodsIF3o7N.
- Heinrich Küttler, Nantas Nardelli, Alexander H. Miller, Roberta Raileanu, Marco Selvatici, Edward Grefenstette, and Tim Rocktäschel. The nethack learning environment, 2020. URL https://arxiv.org/abs/2006.13760.
 - Joel Lehman, Jonathan Gordon, Shawn Jain, Kamal Ndousse, Cathy Yeh, and Kenneth O. Stanley. Evolution through large models, 2022. URL https://arxiv.org/abs/2206.08896.

- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive nlp tasks, 2021. URL https: //arxiv.org/abs/2005.11401.
 - Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, Minlie Huang, Yuxiao Dong, and Jie Tang. Agentbench: Evaluating LLMs as agents. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=zAdUB0aCTQ.
 - Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. The ai scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*, 2024.
 - Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Eureka: Human-level reward design via coding large language models. *arXiv preprint arXiv:2310.12931*, 2023.
 - Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Eureka: Human-level reward design via coding large language models, 2024. URL https://arxiv.org/abs/2310.12931.
 - MiniHack Team. Battle environments. MiniHack Documentation. URL https://minihack.readthedocs.io/en/latest/envs/navigation/battle.html. Accessed: 2025-09-15.
 - Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning, 2013. URL https://arxiv.org/abs/1312.5602.
 - Davide Paglieri, Bartłomiej Cupiał, Samuel Coward, Ulyana Piterbarg, Maciej Wolczyk, Akbir Khan, Eduardo Pignatelli, Łukasz Kuciński, Lerrel Pinto, Rob Fergus, et al. Balrog: Benchmarking agentic llm and vlm reasoning on games.
 - Ruwei Pan, Hongyu Zhang, and Chao Liu. Codecor: An Ilm-based self-reflective multi-agent framework for code generation. *arXiv preprint arXiv:2501.07811*, 2025.
 - Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior, 2023. URL https://arxiv.org/abs/2304.03442.
 - Bharat Prakash, Tim Oates, and Tinoosh Mohsenin. Llm augmented hierarchical agents, 2023. URL https://arxiv.org/abs/2311.05596.
 - J Rosser and Jakob Nicolaus Foerster. Agentbreeder: Mitigating the ai safety impact of multi-agent scaffolds via self-improvement. *arXiv preprint arXiv:2502.00757*, 2025.
 - Mikayel Samvelyan, Robert Kirk, Vitaly Kurin, Jack Zhang, Shimon Jiang, and Jakob Foerster. Minihack the planet: A sandbox for open-ended reinforcement learning research. *arXiv preprint arXiv:2109.13202*, 2021.
 - Mikayel Samvelyan, Sharath Chandra Raparthy, Andrei Lupu, Eric Hambro, Aram Markosyan, Manish Bhatt, Yuning Mao, Minqi Jiang, Jack Parker-Holder, Jakob Foerster, et al. Rainbow teaming: Open-ended generation of diverse adversarial prompts. *Advances in Neural Information Processing Systems*, 37:69747–69786, 2024.
 - Weihao Tan, Wentao Zhang, Shanqi Liu, Longtao Zheng, Xinrun Wang, and Bo An. True knowledge comes from practice: Aligning llms with embodied environments via reinforcement learning, 2024. URL https://arxiv.org/abs/2401.14151.
 - DeepMind Team. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575:350–354, 2019. doi: 10.1038/s41586-019-1724-z.

- MiniHack Team. Minihack: Corridor environment, 2024. URL https://minihack.readthedocs.io/en/latest/envs/navigation/corridor.html. Accessed: September 15, 2025.
 - Sony AI Team and Polyphony Digital. Outracing champion gran turismo drivers with deep reinforcement learning. *Nature*, 602:223–228, 2022. doi: 10.1038/s41586-021-04357-7.
 - Ke Wang, Houxing Ren, Aojun Zhou, Zimu Lu, Sichun Luo, Weikang Shi, Renrui Zhang, Linqi Song, Mingjie Zhan, and Hongsheng Li. Mathcoder: Seamless code integration in Ilms for enhanced mathematical reasoning. In *The Twelfth International Conference on Learning Representations, ICLR*, 2024a.
 - Zihao Wang, Shaofei Cai, Anji Liu, Yonggang Jin, Jinbing Hou, Bowei Zhang, Haowei Lin, Zhaofeng He, Zilong Zheng, Yaodong Yang, Xiaojian Ma, and Yitao Liang. Jarvis-1: Openworld multi-task agents with memory-augmented multimodal language models, 2023. URL https://arxiv.org/abs/2311.05997.
 - Zihao Wang, Shaofei Cai, Anji Liu, Yonggang Jin, Jinbing Hou, Bowei Zhang, Haowei Lin, Zhaofeng He, Zilong Zheng, Yaodong Yang, et al. Jarvis-1: Open-world multi-task agents with memory-augmented multimodal language models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024b.
 - Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*, 2022.
 - Wujiang Xu, Kai Mei, Hang Gao, Juntao Tan, Zujie Liang, and Yongfeng Zhang. A-mem: Agentic memory for llm agents, 2025. URL https://arxiv.org/abs/2502.12110.
 - John Yang, Carlos E Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. Swe-agent: Agent-computer interfaces enable automated software engineering. *Advances in Neural Information Processing Systems*, 37:50528–50652, 2024.
 - Tom Zahavy, Yannick Schroecker, Feryal Behbahani, Kate Baumli, Sebastian Flennerhag, Shaobo Hou, and Satinder Singh. Discovering policies with domino: Diversity optimization maintaining near optimality. *arXiv* preprint arXiv:2205.13521, 2022.

A LLM USAGE DECLARATION

We employed LLMs to assist us in the writing of this paper. Our writing pipeline consisted of one of the authors first writing a draft paragraph, then using an LLM to assist in polishing the writing and grammar, and finally having other authors review and provide the final version to the text. Additionally, we employed foundational models to assist us in creating illustrations.

Finally, throughout the course of this research we used LLM-powered search engines like Perplexity in addition to traditional alternatives such as Google Scholar and conference proceedings while gathering relevant literature.

B BALROG GAME DETAILS

The BALROG framework incorporates six distinct gaming environments, each designed to evaluate specific aspects of agentic reasoning (Figure 3):

BabyAI BabyAI is a grid-based environment with different difficulty levels. The agent is presented wth five different tasks.

TextWorld TextWorld offers a text-based exploration environment where agents interact exclusively through natural language commands. There are three different tasks.

Crafter Crafter simulates a Minecraft-inspired survival environment where progression is measured through 22 distinct achievements.

BabaIsAI BabaIsAI presents a rule-based puzzle environment where agents must navigate grid-based scenarios. There are 40 different tasks.

MiniHack MiniHack represents a task-oriented version of the classic NetHack (Küttler et al., 2020) game, evaluating agents across eight challenges testing different skills.

NetHack Learning Environment (NLE) NLE implements the complete NetHack roguelike game, presenting the most comprehensive challenge within the benchmark. This environment simultaneously evaluates navigation, survival instincts, long-term strategic planning, resource management, and exploration skills within an unpredictable, dynamically evolving game state.

Skills	BabyAI	TextWorld	Crafter	Baba Is AI	MiniHack	NLE
Navigation	V	~	V	~	V	~
Exploration	V	~	V	~	✓	~
Resource Management	X	~	V	×	✓	~
Complex Credit Assignment	X	×	V	~	✓	~
Deducing Env. Dynamics	X	×	X	~	~	~
Long-term Planning	X	X	X	V	~	~
Turns to Complete Time to Master for Humans	10 ¹ Seconds	10 ² Minutes	10 ³ Hours	10 ² Hours	10 ² Hours	10 ⁴ –10 ⁵ Years

Figure 3: Game environments overview. Adapted from BALROG (Paglieri et al.)

C HYPERPARAMETERS SELECTED

This Appendix details the choice of hyperparameters in our methodology. Table 4 details the values and descriptions.

The genetic algorithm optimises four hyperparameters: τ_{\cos} , λ , ϵ_0 , and T. We establish lower bounds of 0 for τ_{\cos} , λ , and ϵ_0 , effectively disabling these components when not beneficial to performance. Upper bounds were determined in order to maintaini sufficient search space for optimisation.

The language model temperature T follows standard practice with a default value of 1.0, allowing the genetic algorithm to explore a range of solutions. We implement exponential decay for ϵ_t following established reinforcement learning approaches, enabling the transition from exploration to exploitation as the system learns optimal behaviours. We follow the short term memory length in BALROG and set it to m=16. We set $k_{succ}=5$ and $k_{fail}=5$, limiting the number of information added to the prompt, but also adding significant amount of past experiences; through empirical evaluation we notice that a higher number of memories added is not beneficial.

Following DOMiNO methodology, we set $\alpha=0.7$ to ensure meaningful population diversity whilst maintaining performance standards. Our similarity-recency weighting ($w_{\text{similarity}}=0.7$, $w_{\text{recency}}=0.3$) prioritises semantic relevance over temporal proximity, reflecting the hypothesis that content similarity is more beneficial than recency.

The genetic algorithm parameters balance computational efficiency with solution quality. We set probability of selecting single parent in genetic algorithm to be 70% (vs two parents to be 30%), allowing for more mutations without crossover operations. We set n=4 iterations as empirical evaluation demonstrated satisfactory performance is achieved at this point, providing an effective balance between solution quality and computational cost. Population management parameters N=5 maintain an optimal balance between preserving high-performing solutions and promoting genetic diversity, following established evolutionary computation principles that prevent premature convergence whilst ensuring computational tractability.

D BALROG BASELINE CONFIGURATIONS

This section details number of episodes and their length per each BALROG game. Moreover, we show the BALROG prompt that we use as initial seed for the baseline architecture.

Episode details Table 5 details the episode specifications per game. The table shows time needed for each episode completion, as well as details on number of tasks per different environments.

Baseline Prompt (BALROG) Below we present a prompt from BALROG (Paglieri et al.) paper, used for Baseline evaluation.

```
Baseline BALROG Prompt

"""You always have to output one of the above actions at a time and no other text. You always have

→ to output an action until the episode terminates."""
```

E TESTING THE RETRIEVAL MECHANISM LONG-TERM MEMORY SYSTEM

In this section we test the retrieval of saved memories and the abilities to act upon them. Due to stochasticity, we need to have a reliable comparison. We focus our evaluation on Crafter, as it is an environment requiring long-term planning, giving motivation to log-term memory approach. We disable life hazards such as zombies and skeletons, as they are not relevant to the testing subject. We set random seed to 32 for all episodes.

To evaluate our information-retrieval mechanism we test a long-horizon "craft iron sword" task in Crafter. We replace the environment's default objective with the production of an iron sword (see Appendix E.1.1). This task is intentionally complex: it requires chopping wood, crafting and placing a crafting table, crafting a wooden pickaxe, collecting stone, crafting a stone pickaxe, placing a furnace adjacent to the crafting table, collecting iron, and finally crafting an iron sword. We selected this objective because, in prior baseline runs without our memory system, the agent never completed the task. Through the episode we would like to check if memories are activated at relevant time

756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806

808

809

	Value	Description	
Embedding	all-MiniLM-L6-v2	Pre-trained sentence embedding model used for semantic similarity calculations	
d	384	Dimension of the embedding	
$\overline{ au_{\cos}}$	[0, 0.1]	Cosine similarity threshold parameter	
λ	[0, 0.1]	Long-Term Memory decay factor	
$\overline{\epsilon_0}$	[0, 0.1]	Initial exploration parameter	
\overline{T}	[0.1, 2]	Language Model Temperature	
$\overline{\epsilon_t}$	$0.99^t \times \epsilon_0$	Time-decayed parameter following exponential decay	
\overline{m}	16	Short-term memory length (most recent action-observation pairs)	
$\overline{k_{success}}$	5	Number of top scoring successful long-term memories added to the LLM prompt	
$\overline{k_{fail}}$	5	Number of top scoring failed long-term memories added to the LLM prompt	
α	0.7	Minimum fraction of highest scoring genome for diversity	
$\overline{w_{ m similarity}}$	0.7	Weight assigned to similarity component in scoring	
$\overline{w_{ m recency}}$	0.3	Weight assigned to recency component in scoring	
$\overline{p_{single}}$	70%	Percentage chance to choose single parent for reproduction	
$1 - p_{single}$	30%	Percentage chance to choose two parents for reproduction	
$\overline{n_{gen}}$	4	Number of iterations (parent population creation) of genetic algorithm	
$\overline{n_{child}}$	5	Number of children created for each population of parents in genetic algorithm	
\overline{N}	5	Number of best scoring genomes saved at eastep of genetic algorithm	
\overline{M}	5	Number of most diverse genomes saved at each step of genetic algorithm (scoring at least α fraction of top performing genome)	
$\overline{n_{ep}}$	dependent on the game	Number of episodes for each child evaluation. Details in the Appendix D	

Table 4: Hyperparameter values used in the TAME framework

steps, showing retrieval ability. Moreover, successful completion under our system provides strong evidence that the memory-critic architecture supports multi-step planning and sequential options.

In order to track the memories activated, we inject five task-oriented memories at the start of each episode: "craft wooden pickaxe", "craft stone pickaxe", "mine iron", "place furnace", and "craft iron sword" (see Appendix E.1.2). Each memory is paired with a human-crafted critic that summarises the steps needed to achieve particular option. During each episode we log when each memory activates. An example progression through episode is provided in Figure 4 and the corresponding

Environment	Evals	Tasks per Eval	Total	Episode Length
BabyAI	10	5	50	10^{1}
BabaIsAI	3	40	120	10^{2}
Crafter	10	1	10	10^{3}
TextWorld	10	3	30	10^{2}
MiniHack	5	8	40	10^{2}
NetHack	5	1	5	$10^4 - 10^5$

Table 5: Episode details per BALROG game

memories activated can be seen in Figure 5. More examples on activation timelines are provided in Appendix E.2.

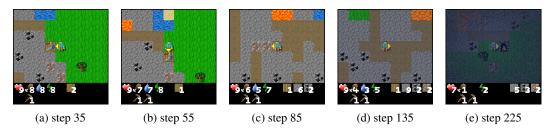


Figure 4: Testing memory retrieval: Task progression across episode: (a) agent crafts wood pickaxe, (b) agent crafts stone pickaxe, (c) agent collects iron, (d) agent collects iron, (e) agents attempts to craft iron sword

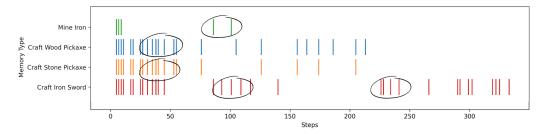


Figure 5: Testing memory retrieval: memories activated based on the step

Discussion of the example: From Figure 5 we can notice that all memories are activated at the beggining. This is due to the fact that very early in the game, those are they only memories present. We can then see that "Craft Wood Pickaxe" and "Craft Stone Pickaxe" are heavily retrieved until around 25th-60th step. This is when agent completed "Craft Stone Pickaxe task" (agent completes "Craft Wood Pickaxe" task earlier, but due to similarity of those two tasks, it it activated when focusing on stone version). "Mine Iron" memory activates two times between 80-100 steps, when agent is mining two pieces of iron. "Craft Iron Sword" appears often around step number 100 which is when agent first attempts to complete it, but realises that it needs to place a furnace and table first. Then the memory is activated later as well, which is after placing table and furnace and attempting to craft iron sword. Unfortunately, agent is unsuccessful because it didn't place furnace close enough to the table. This experiment demonstrates that relevant memories are activated at the right time, and that agent is able to act upon them. Also, when not needed (i.e. the task is completed), memories are activated far less often. It is also important to notice that the "Craft Iron Sword" memory continues to be retrieved even towards the end of the episode when the agent is actively attempting this task. This indicates that the memory system maintains access to relevant historical experiences throughout the entire episode, regardless of when they were initially formed.

Results Across 10 independent runs the agent succeeded in producing an iron sword in 1/10 episodes (baseline: 0% from all previous runs that we did with Gemini-2.0-Flash). This is a sub-

stantial improvement, which suggests that when memories are stored, agent is able to retrieve them, and act upon them. This is a simplified case, as we provided human crafted memories, but with the right prompting we believe that the critic module will be able to reproduce those. Additional observations:

- Memories reliably activated when their prerequisites were satisfied and were deactivated immediately after the corresponding option was completed.
- The specificity of the critic strongly affected performance. For example, phrasing a critic as "place the furnace next to the crafting table" versus "place the furnace adjacent to the table" produced different success scores. This highlights the value of precise, action-oriented critic definitions that focus on the key state features leading to success or failure. Following on that we prompted the critic accordingly.
- This experiment also demonstrates the difficulty of the "Craft Iron Sword" task, even when given with clear instructions agent fails 90% of the time.

E.1 TESTING MEMORY RETRIEVAL PROMPTS

This subsection details prompts used in order to test long-term memory retrieval. First we show the prompt detailing the goal of iron sword creation, then we show memories added at the begging of the episode: *Craft Wood Pickaxe, Craft Stone Pickaxe, Craft Stone Sword, Mine Iron, Place Furnace, Create Iron Sword.* Those are hand-crafted memories, designed in order to track memory retrieval.

E.1.1 IRON SWORD GOAL PROMPT

This subsection details the prompt for Craft Iron Sword goal that the agent is tasked with during the long-term memory retrieval experiment.

E.1.2 INJECTED MEMORIES

This subsection details the memories added at the beggining of the episode, in order to track memory retrieval.

```
"subgoal_progress_indicators": "Agent is gathering stone near table",
"subgoal_penalty_component": "Agent crafts pickaxe without enough resources",
"status": 'successful',
"summary of the run": "Agent collects four pieces of wood and places a table in a clear spot. Then

\( \to \) agent collects one piece of stone using wood pickaxe. Then agent crafts a stone pickaxe at

\( \to \) the table.",
}
```

Craft Stone Sword Memory

Mine Iron Memory

Place Furnace Memory

```
furnace_memory = {
    "name": "Place Furnace next to Table",
    "description": "Place furnace next to the Table for crafing iron tools",
    "subgoal_prerequisites": "Agent has 4 pieces of stone in inventory. Table is placed. ",
    "success_condition": "Furnace is placed",
    "subgoal_progress_indicators": "Agent is gathering stone",
    "subgoal_penalty_component": "Agent places furnace in unsuitable location or without enough
    \leftarrow stone",
    "status": 'successful',
    "summary of the run": "Agent placed a furnace next to the table using 4 pieces of stone. ",
}
```

Craft Iron Sword Memory

E.2 ADDITIONAL LONG-TERM MEMORY RETRIEVAL EXPERIMENTS

This section shows additional experiemnts carried out in order to test memory retrieval, when tasked the agent with iron sword task. Eachof the experiments consists of images showing agent progression, as well as memory activation across the episode.

E.2.1 EPISODE 1

Illustrations from the game available in Figure 6 and memory activations in 7. The reason why agent didn't succeed in completing the task is because the agent didn't place table close enough to furnace.



Figure 6: Episode 1: Task progression across episode

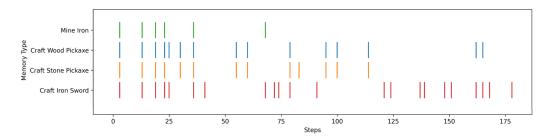


Figure 7: Episode 1: Memory retrieval along the episode

E.2.2 EPISODE 2

Illustrations from the game available in Figure 8 and memory activations in 9. The reason why agent didn't succeed in completing the task is because the agent focuses on placing furnaces a few times.

E.2.3 EPISODE 3

Illustrations from the game available in Figure 10 and memory activations in 11. The reason why agent did not succeed in completing the task is because agent does not have enough wood (also crafts multiple tables and furnaces).

E.2.4 EPISODE 4

Illustrations from the game available in Figure 12 and memory activations in 13. The reason why agent did not succeed in completing the task is because agent does not have enough wood (also crafts multiple tables and furnaces).

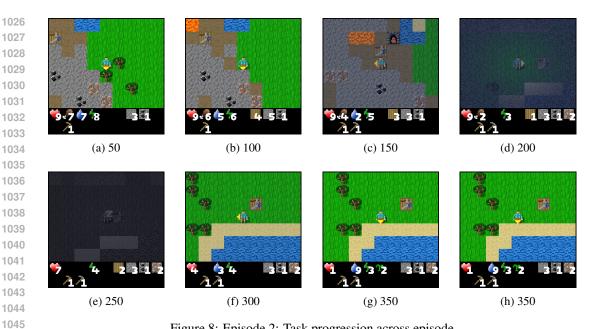


Figure 8: Episode 2: Task progression across episode

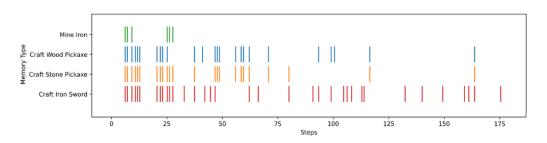


Figure 9: Episode 2: Memory retrieval along the episode

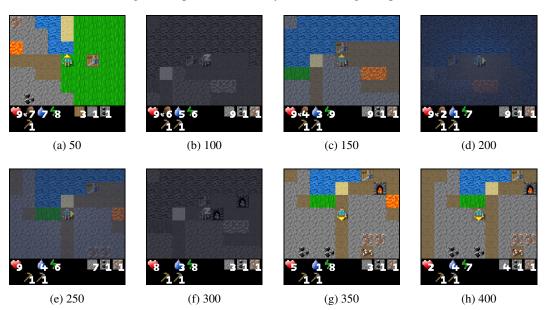


Figure 10: Episode 3: Task progression across episode

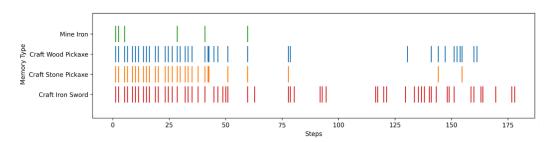


Figure 11: Episode 3: Memory retrieval along the episode

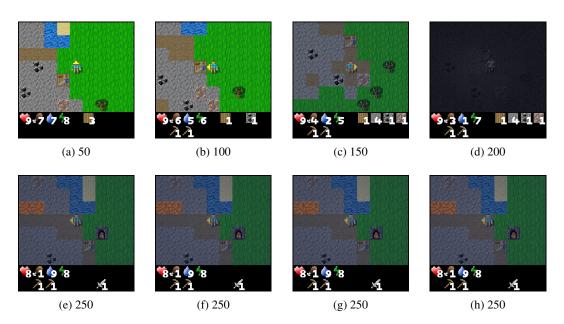


Figure 12: Episode 4: Task progression across episode

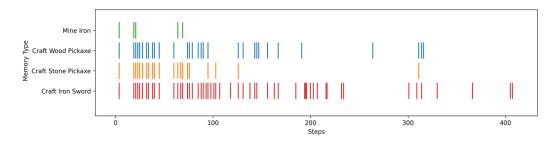


Figure 13: Episode 4: Memory retrieval along the episode

F EXAMPLES OF MODULES OUTPUTS

This section shows examples of amygdala and explorer modules behaviours.

F.1 AMYGDALA

This section details Amygdala submodule. First, we show defualt survival option, then we illustrate an example of amygdala in the episode (using Crafter).

Default survival option The prompt below shows default survival option, that agent is switching to, whenever danger is encountered.

```
subgoal = {
   "name" : "respond to danger",
   "description": "Respond to danger.",
   "prerequisites": "Agent is in immediate danger",
   "success_condition": "Agent has eliminated the danger",
   "penalty_component": "Agent is not responsing to danger.",
   "progress_indicators": "Agent is closer to eliminating the danger",
   "estimated_priority": "high"
}
```

Amygdala module activation. Here we present an example from Crafter when amygdala is activated. The agent is initially focused on exploring the environment and fulfilling subgoals such as placing plants. However, in the early stages of the game, a skeleton appears. As soon as the agent observes the skeleton within its field of view, it activates survival mode.

After survival mode is activated, the agent begins gathering resources for combat:

- 1. Chops down wood
- 2. Places a crafting table
- 3. Creates a wooden sword necessary for the fight

The agent then chases the skeleton, and once it is adjacent to the enemy, it initiates combat. During the chase and fight, the agent loses health but successfully manages to defeat the skeleton. Shortly afterwards, the amygdala is deactivated and the agent returns to working towards general game objective.



Figure 14: Amygdala survival response sequence: (a) activation, (b)&(c) resource gathering, (d)&(e) combat, (f) deactivation.

F.2 EXPLORER

1188

11891190

1212 1213

1214 1215

1216

1217 1218

1219

1222

1223 1224

1225 1226

1227

1228

1229

1230

1231

1232

1233 1234

1236

1237

1239

1240

1241

In this section we show an example of exploration plan proposed by Explorer.

```
1191
      {'reasoning': 'The agent needs to break out of its repetitive westward
1192
      movement and explore new areas, prioritizing resource gathering and
1193
      different directions. Sand has been identified as a new area to explore.',
1194
      'subgoals': [
1195
         {'name': 'Move Towards Sand',
          'description': 'Move towards the sand to the south-east to explore
1196
         new terrain.',
1197
          'prerequisites':
1198
         'success_condition': 'Agent is standing on sand.',
1199
         'penalty_component': 'Moving in the opposite direction (North or West)
         for more than 3 steps without collecting resources.',
1201
         'progress_indicators': 'Decreasing distance to sand in long_term_context.',
         'estimated_priority': 'high'},
1202
         {'name': 'Explore Area Around Sand',
1203
          'description': 'Once on the sand, explore in all directions (North, South,
1204
         East, and West) to discover resources and new landmarks.',
1205
          'prerequisites': 'Agent is standing on sand.',
         'success_condition': 'Agent has moved at least 3 steps in each cardinal
1206
         direction from the sand.',
1207
         'penalty_component': 'Staying within the same 3x3 area on the sand for
1208
         more than 5 steps.',
1209
         'progress_indicators': 'Number of unique tiles visited around the sand.',
1210
          'estimated_priority': 'medium'}
1211
      ] }
```

G DETAILED GENETIC ALGORITHM

In this Appendix we present the detailed specifications of the genetic algorithm. Parameter selection:

Generations = 4 (17)

Population size =
$$|\mathcal{P}_{\text{diverse}}| + |\mathcal{P}_{\text{best}}| = 5 + 5 = 10$$
 (18)

Children per generation
$$= 5$$
 (19)

$$p_{\text{binary}} = p_{\text{continuous}} = p_{\text{prompt}} = 0.5$$
 (20)

where:

• **Generations** – the number of evolutionary iterations performed

- **Population size** the total number of genomes maintained across diverse and bestperforming subpopulations
- **Children per generation** the number of new offspring genomes created through mutation in each generation
- $p_{\text{binary}}, p_{\text{continuous}}, p_{\text{prompt}}$ the probabilities of applying binary, continuous, or prompt mutation operations, respectively, when creating offspring.

The number of generations was limited to four due to the significant computational cost and long runtimes associated with evaluating each genome, particularly in environments like NetHack. However, this was sufficient to demonstrate a clear performance improvement and allow for the discovery of specialised architectures.

Parent Selection Parent selection follows roulette wheel selection with fitness-proportionate probabilities:

$$P(\mathbf{g}_i) = \frac{f(\mathbf{g}_i)}{\sum_{j=1}^{|\mathcal{P}|} f(\mathbf{g}_j)}$$

where $f(\mathbf{g}_i)$ is the fitness score of genome \mathbf{g}_i .

Offspring Generation For each child g_{child} :

- With probability p_{single} : select one parent
- With probability $1 p_{single}$: select two parents

For each genome component, evolutionary operations are applied with probability $p_{binary}, p_{continuous}, p_{prompt}$, depending on the component. Otherwise parent attributes are copied directly. In the two-parent case, the parent from which to copy each attribute is chosen with Bernoulli(0.5) probability.

G.1 GENETIC OPERATIONS

This subsection focuses on methodology of crossover and mutations operations.

Modules Operations Single parent:

$$b_i^{(child)} = \begin{cases} b_i^{(parent)} & \text{with probability } 0.8\\ 1 - b_i^{(parent)} & \text{with probability } 0.2 \end{cases}$$

Two parents:

$$b_i^{(child)} = \begin{cases} b_i^{(p1)} = b_i^{(p2)} & \text{with probability } 0.9 \text{ if } b_i^{(p1)} = b_i^{(p2)} \\ 1 - b_i^{(p1)} & \text{with probability } 0.1 \text{ if } b_i^{(p1)} = b_i^{(p2)} \\ \text{Bernoulli}(0.5) & \text{if } b_i^{(p1)} \neq b_i^{(p2)} \end{cases}$$

Hyperprameter Operations The continuous value inheritance depends on parent activity states. Let $A_i^{(pk)}$ indicate if feature i is active in parent k:

For single parent:

$$c_i^{(child)} = \text{clip}(c_i^{(parent)} + \mathcal{N}(0, \sigma^2), c_{i, \min}, c_{i, \max})$$

For two parents:

$$c_i^{(child)} = \begin{cases} c_i^{(p1)} + \mathcal{N}(0, \sigma^2) & \text{if } A_i^{(p1)} = 1, A_i^{(p2)} = 0 \\ c_i^{(p2)} + \mathcal{N}(0, \sigma^2) & \text{if } A_i^{(p1)} = 0, A_i^{(p2)} = 1 \\ \alpha c_i^{(p1)} + (1 - \alpha) c_i^{(p2)} + \mathcal{N}(0, \sigma^2) & \text{if } A_i^{(p1)} = A_i^{(p2)} = 1 \\ c_{i, \text{default}} & \text{if } A_i^{(p1)} = A_i^{(p2)} = 0 \end{cases}$$

where $\alpha \sim U(0,1)$. $c_{i,min}, c_{i,max}, c_{i,default}$ are detailed in Appendix C.

Prompt Operations Prompt evolution utilises the EvoPrompt prompt methodology Guo et al. (2025). Using a similar approach we use LLM as a crossover and mutation operator.

- Single parent: $p_i^{(child)} = \text{LLMprompt}_{mutate}(p_i^{(parent)})$
- $\bullet \ \ \mathbf{Two \ parents} : \ p_i^{(child)} = \mathrm{LLMprompt}_{mutate}(\mathrm{LLMprompt}_{crossover}(p_i^{(p1)}, p_i^{(p2)})) \\$

If a parent has module i disabled ($b_i = 0$), the corresponding prompt reverts to default: $p_i = p_{i,\text{default}}$ (default prompts in Appendix O.1).

LLM mutation and crossover prompts available in Appendix N. It is important to notice that this approach enables functional mutation: LLM is prompted with all possible functional inputs to be used in any prompt (described in Section 3.1).

G.2 POPULATION MANAGEMENT

 This subsection focuses on population management: details about fitness function, diversity measure and population pruning.

Fitness Evaluation Each genome is evaluated using the fitness function:

$$f(\mathbf{g}) = \frac{1}{n_{ep}} \sum_{i=1}^{n_{ep}} \mathrm{GameProgression}_i(\mathbf{g})$$

where n_{ep} is the number of episodes for specific game (see Table 5).

Diversity measure The genome distance function uses an embedding-based approach where each genome is represented as a single embedding vector (here we use sentence-transformers/all-MiniLM-L6-v2 embedding). The distance between two genomes is calculated using cosine similarity:

$$d(\mathbf{g}_1, \mathbf{g}_2) = 1 - \cos(\mathbf{e}_1, \mathbf{e}_2) \tag{21}$$

where e_i is the embedding vector representation of genome g_i .

The cosine similarity between two embedding vectors is computed as:

$$\cos(\mathbf{e}_1, \mathbf{e}_2) = \frac{\mathbf{e}_1 \cdot \mathbf{e}_2}{\|\mathbf{e}_1\|_2 \|\mathbf{e}_2\|_2}$$

Each genome $\mathbf{g} = \{\mathbf{b}, \mathbf{c}, \mathbf{p}\}$ is transformed into a unified embedding vector $\mathbf{e} \in \mathbb{R}^d$ that captures the semantic representation of all genome components (binary variables, continuous parameters, and prompts) in a single high-dimensional space. The way we measure diversity, is the minimum distance to the genomes already existing in the archive.

Population Pruning After children evaluation, population pruning maintains diversity using the following algorithm:

- 1. Initialise $\mathcal{P}_{\text{new}} = \emptyset$
- 2. Add top-N scoring genomes: $\mathcal{P}_{\text{new}} \leftarrow \text{top}_5(\mathcal{P})$
- 3. For remaining genomes $\mathcal{G}_{\text{remaining}}$:
 - (a) Calculate minimum distance to current population:

$$d_{\min}(\mathbf{g}) = \min_{\mathbf{g}' \in \mathcal{P}_{\text{new}}} d(\mathbf{g}, \mathbf{g}')$$

(b) Select genome maximising diversity with performance constraint:

$$\mathbf{g}^* = \arg\max_{\mathbf{g} \in \mathcal{G}_{\text{remaining}}} d_{\min}(\mathbf{g}) \quad \text{s.t. } f(\mathbf{g}) \geq 0.7 \cdot f(\mathbf{g}_{\text{best}})$$

- (c) Add \mathbf{g}^* to \mathcal{P}_{new} and remove from $\mathcal{G}_{\text{remaining}}$
- 4. Repeat step 3 until the desired population size reached (N + M)

G.3 GENETIC APPROACH: PSEUDO CODE

1350

1351

1377 1378 1379

1380 1381

1382

1383

1384

1385

1386

1387

1388 1389 1390

1391

1392

1393

1394

1395

1398

1399 1400

1401 1402

1403

```
1352
1353
           Algorithm 1 TAME: Genetic Algorithm
1354
           Require: Game environment
1355
           Ensure: Optimised genome g^*
1356
            1: Initialize P_0 = \{g_{\text{basic}}, g_{\text{hierarchical}}, g_{\text{default}}, g_{\text{full}}\}
1357
            2: for each g \in P_0 do
1358
            3:
                   g_{\text{fitness}} \leftarrow \text{EvaluateFitness}(g)
1359
            4: end for
            5: P \leftarrow P_0
1360
            6: for generation = 1 to GENERATIONS do
1361
                   C \leftarrow \emptyset {Children population}
            7:
1362
                   for i=1 to CHILDREN_PER_GENERATION do
            8:
1363
            9:
                      if rand() < 0.7 then
1364
           10:
                         p_1 \leftarrow \text{RouletteWheelSelection}(P)
1365
           11:
                          c \leftarrow \text{SingleParentOperations}(p_1)
1366
           12:
                      else
1367
                         p_1, p_2 \leftarrow \text{RouletteWheelSelection}(P, 2)
           13:
1368
           14:
                          c \leftarrow \text{TwoParentOperations}(p_1, p_2)
1369
           15:
                      end if
                      c_{\text{fitness}} \leftarrow \text{EvaluateFitness}(c)
1370
           16:
           17:
                       C \leftarrow C \cup \{c\}
1371
           18:
                   end for
1372
           19:
                   P \leftarrow \text{PopulationPruning}(P \cup C)
1373
           20: end for
1374
           21:
1375
           22: return \arg \max_{g \in P} g_{\text{fitness}}
1376
```

H MODULES ACTIVATED

In this Appendix we discuss module activation based on the game. Module activation is based on final genomes returned by genetic algorithm, available in Appendix P. Module activation plot is demonstrated in Figure 15. When selecting the baseline configuration, no additional modules apart from Long-Term Memory can be activated. We notice that 4 out of 6 environments selected hierarchical module, highlighting the effectiveness of complex goal decomposition. TextWorld is a text-based environment where it is difficult to predict next actions due to their dependence on current observation, therefore hierarchical structure and memory are not adding value. Moreover, MiniHack has relatively short length (100 steps), which might be also why baseline structure was favoured.

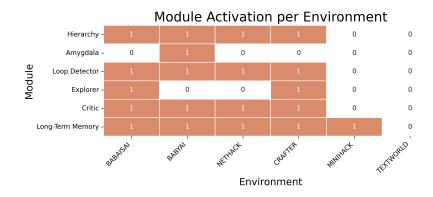


Figure 15: Module activation in TAME across environments.

I TASK PERFORMANCE

In this Appendix, we compare different tasks performance across BabyAI, BabaIsAI, and Minihack for baseline versus TAME. Crafter and Nethack are excluded because they each have only one default task. TextWorld is also excluded since its genetic output matches the baseline.

MiniHack Interestingly, five out of eight tasks are never solved by any method, showcasing the difficulty (see Figure 16). The Corridor-R3 task, which is never completed by the baseline, nevertheless shows 40% progress with TAME. Corridor-R3 is an exploration problem in which the goal is to find the staircase Team (2024), illustrating TAME agent's improved exploration ability. In both CorridorBattle-Dark and MazeWalk-9×9, TAME achieves higher performance. CorridorBattle-Dark requires the agent to fight monsters, thereby testing planning and memory MiniHack Team, whereas MazeWalk-9×9 is a maze in which the agent must reach a terminal goal, testing exploration and memory Samvelyan et al. (2021).

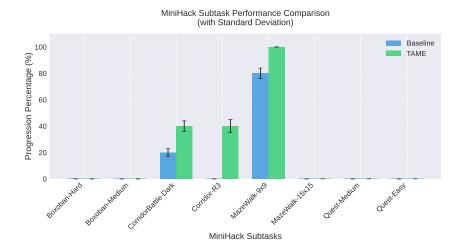


Figure 16: MiniHack tasks progression, Baseline vs TAME using Gemini-2.0-Flash.

BabyAI We observe clear performance improvements across all tasks except "putnext" where both agents achieve 0% success rate (see Figure 17). The putnext task presents significant challenges due to its complex spatial reasoning requirements. Through empirical analysis, we identified that agents fail to understand the necessary positioning strategy: they must navigate to a location one step away from the target position (which is adjacent to the object) before dropping the item. The persistence of this failure in our improved method highlights fundamental limitations in the agent's spatial reasoning capabilities. In all the other tasks, we notice an improvement when comparing TAME with baseline.

BabaIsAI This environment consists of 40 distinct tasks that can be categorised into four main types: make_win, make_you, goto_win, and make_wall_win. Our analysis reveals substantial improvements in the goto_win category and notable progress in make_win tasks, where performance increased from 0% baseline (see Figure 18). When examining performance across different room configurations (two_room versus single_room layouts), we observe consistent improvements in both settings. For difficulty categorisation, we define three levels based on task complexity: simple tasks have no modifiers or distractors, medium tasks contain 1-2 modifiers/distractors, and complex tasks have more than 2 modifiers/distractors. Most notably, the greatest performance gains occur in medium and complex categories, demonstrating that our method is particularly effective for challenging scenarios that require sophisticated reasoning capabilities.

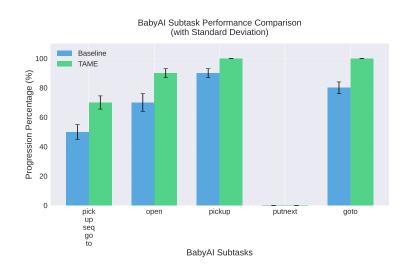


Figure 17: BabyAI tasks progression, Baseline vs TAME using Gemini-2.0-Flash.

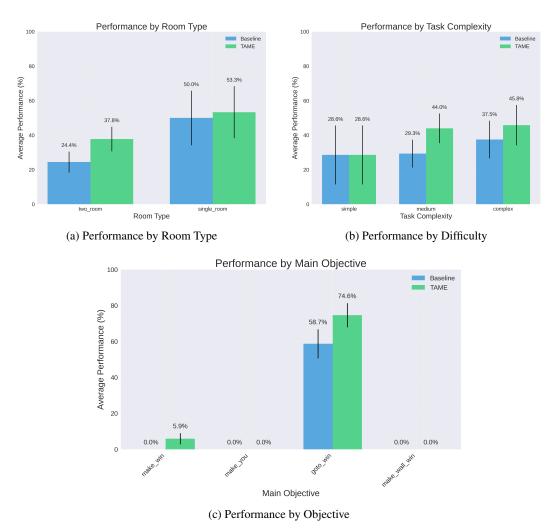


Figure 18: BabaIsAI task progression, Baseline vs TAME using Gemini-2.0-Flash

J GENOME TRANSFERABILITY TO OTHER GEMINI MODELS

In this Appendix we detail transcrability of genomes obtained through genetic algorithm using Gemini-2.0-Flash, to Gemini-2.5-Flash-Lite and Gemini-2.5-Pro without additional training. In the Table 6 we compare performance of Baseline using Gemini-2.5-Flash-Lite vs TAME transferred to the same model. In the Table 7 we compare performance of Baseline using Gemini-2.5-Pro vs TAME transferred to the same model. Lastly, we show combined results in the Figure 19.

Environment	Baseline [2.5-Flash-Lite]	TAME[Transferred]	Episodes
Average	11.87% \pm 1.32%	$22.76\% \pm 1.73\%$	-
babyai	$46.00\% \pm 7.05\%$	$62.00\% \pm 6.93\%$	50
babaisai	$9.17\% \pm 2.63$	$38.33\% \pm 4.44\%$	120
textworld	$7.45\% \pm 2.30\%$	$7.45\% \pm 2.30\%$	30
crafter	$8.64\% \pm 1.00\%$	$21.3\% \pm 4.22\%$	10
minihack	$0.00\% \pm 0.00\%$	$7.50\% \pm 4.16\%$	40
nle	$0.00\% \pm 0.00\%$	$0.00\% \pm 0.00\%$	5

Table 6: Comparison of Baseline vs. TAME[Transferred] using Gemini-2.5-Flash-Lite.

Game	Baseline [2.5-Pro]	Top Model	TAME[Transferred]	Runs
Average	$43.35\% \pm 2.31\%$	43.60% \pm 2.17%	$47.65\% \pm 2.20\%$	-
babyai	$80.0\% \pm 5.70\%$	$76.00\% \pm 6.00\%$	$90.0\% \pm 4.24\%$	50
babaisai	$56.70\% \pm 4.50\%$	$45.80\% \pm 4.50\%$	$72.50\% \pm 4.08\%$	120
textworld	$49.20\% \pm 8.20\%$	$62.90\% \pm 7.90\%$	$49.20\% \pm 8.20\%$	30
crafter	$55.0\% \pm 6.0\%$	$57.30\% \pm 3.90\%$	$55.0\% \pm 6.0\%$	10
minihack	$17.50\% \pm 6.00\%$	$17.5\% \pm 6.00\%$	$17.50\% \pm 6.00\%$	40
nle	$1.70\% \pm 0.20\%$	$1.8\%\pm0.8\%$	$1.70\% \pm 0.20\%$	5

Table 7: Comparison of Baseline vs. TAME[Transferred] using Gemini-2.5-Pro.

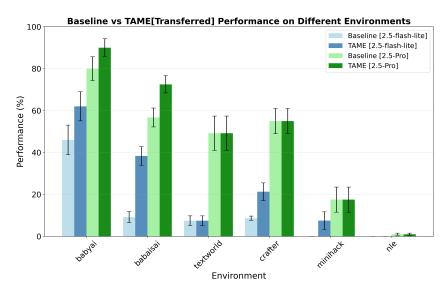


Figure 19: Evaluation of TAME Structure Transferability: Identical performance to the baseline indicates that the genetic algorithm favored the baseline architecture over the transferred TAME genome.

K TAME CHILD EVALUATION PSEUDO CODE

This Appendix presents the pseudo code behind TAME evaluation on specific game when all modules are active, as decided by the genetic algorithm (overview in Figure 2). The algorithm works as follows: until maximum number of steps is reached, with probability ϵ choose exploration, otherwise exploitation. For each option, retrieve top scoring successful and failed memories.. For each option, it retrieves top successful and failed memories, then the Low-level Executor decides and executes actions. If danger is detected, the survival module is triggered and execution stops. Otherwise, the Critic evaluates success, summarizes key actions, and in case of loops, the summary is replaced by the Loop Detector summary. The resulting memory is stored, and the cycle repeats.

Algorithm 2 TAME [Full Structure]

1566

1567 1568

1569

1570

1571

1572

1573

1574

1575

1576 1577

1579

1612 1613 1614

1615

1616

1617

1618

1619

```
1580
           Require: I, s_0
1581
           Ensure: s_T (final game state)
              while num\_steps \le max\_steps do
                 if rand() < \epsilon then
                     g \leftarrow \pi_{explorer}(s_i^e) {Exploration}
1584
1585
                    g \leftarrow \pi_{high}(s_i^{mc}) {Exploitation}
                 end if
                 for g_i \in g do
1587
                     M_i^+ \leftarrow retrieve\_successful(n_{g_i}, d_{g_i}, memory)
                     M_i^- \leftarrow retrieve\_failed(n_{g_i}, d_{g_i}, memory)
                     a \leftarrow \pi_{low}(s_i^{le})
1590
                     for a_i \in a do
1591
                        s_t \leftarrow execute\_action(a, s_i)
1592
                        s_i \leftarrow s_t
                        if \sigma(s_i) then
                            g \leftarrow g^* {Activate survival option}
1594
                            break {Move to option selection and force g^*}
                        end if
                     end for
                     C_i \leftarrow \varphi(s_i)
                     h_i \leftarrow \rho(s_i, C_i)
1598
                     if C_i then
                        store\_memory(g_i, h_i, o, "successful")
                     else
                        if \psi(s_i) then
                           h_i \leftarrow \psi(s_i)
                        end if
                        store\_memory(q_i, h_i, o, "failed")
1604
                        break {Escalate to replanning}
                     end if
                 end for
1607
              end while
```

L DETAILED MEMORY ABLATIONS

In this section we show detailed results when comparing our memory system with Craft-Jarvis-1 and A-Mem. Table 8 shows that TAME-Memory consistently outperforms both Craft-Jarvis-1 and A-Mem on most environments, yielding the highest overall average of 23.11%. The gains are particularly strong in TextWorld and Crafter, where our system nearly doubles or surpasses the baselines by a large margin. While performance is comparable in Minihack and NLE, these tasks remain challenging for all methods. Overall, the results highlight the efficiency of our hybrid memory system.

Environment	Jarivs (↑)	TAME-Memory[ours] (†)	A-mem (†)	Episodes
Average	$17.52\% \pm 1.73\%$	23.11% \pm 1.75%	$21.45\% \pm 1.80\%$	-
babyai	48.00 % ± 7.06 %	$62.00\% \pm 6.86\%$	$58.00\% \pm 6.98\%$	50
babaisai	$24.17 \% \pm 3.90 \%$	$29.17\%\pm4.15\%$	$26.67\% \pm 4.04\%$	120
textworld	$0.59~\% \pm 0.58~\%$	$8.24\% \pm 1.95\%$	$4.51\% \pm 2.35\%$	30
crafter	$19.55~\% \pm 3.81~\%$	$35.45\%\pm3.20\%$	$26.36\% \pm 4.25\%$	10
minihack	12.5% \pm 5.23%	$3.45\% \pm 5.65\%$	12.5% \pm 5.23%	40
nle	$0.31\% \pm 0.28\%$	$0.37\% \pm 0.33\%$	$\textbf{0.68\%}\pm\textbf{0.37\%}$	5

Table 8: Comparison of TAME average game progression across different memory types.

M INITIAL POPULATION SCORES

This section details the scores of initial population used during genetic algorithm. Scores per environment can be seen in Table 9.

Environment	Basic (†)	Hierarchical (†)	Default (↑)	Full (†)
Average	$\%\pm\%$	% ± %	$\%\pm\%$	-
babyai	$58.00\% \pm 6.98\%$	$62.00\% \pm 6.86\%$	$58.00\% \pm 6.98\%$	$65.31\% \pm 6.80\%$
babaisai	$30.83\% \pm 4.22\%$	$29.17\% \pm 4.15\%$	$40.83\% \pm 4.49\%$	$31.67\% \pm 4.25\%$
textworld	$32.55\% \pm 6.95\%$	$8.24\% \pm 2.25\%$	$4.67\% \pm 1.50\%$	$6.08\% \pm 1.70\%$
crafter	$29.09\% \pm 4.51$	$33.64\% \pm 4.64\%$	$31.36\% \pm 3.55\%$	$20.00\% \pm 1.72\%$
minihack	$12.50\% \pm 5.23\%$	$5.00\% \pm 3.45\%$	$12.50\% \pm 5.23\%$	$12.50\% \pm 5.23\%$
nle	$0.00\% \pm 0.00\%$	$0.37\% \pm 0.33\%$	$0.68\% \pm 0.37\%$	$0.00\% \pm 0.00\%$

Table 9: Comparison of performance across environments under Basic, Stable, Default, and Advanced settings.

N LLM MUTATION AND CROSSOVER PROMPTS

This Appendix details the prompts used for LLM based Crossover (in the case of single parent) and LLM based Crossover and Mutation (in the case of two parents) used in genetic algorithm. Prompts created with the help of Claude-Sonnet-4.

```
LLM Mutation Prompt
Please follow the instruction step-by-step to generate a better prompt.
1. Consider prompt:
2. Apply ONE of the following mutation strategies to create a significantly different prompt:
**Strategy A - Perspective Shift:** Change the role/perspective (e.g., "As an expert analyst..." or
      \hookrightarrow "From the viewpoint of...")
**Strategy B - Methodology Change: ** Alter the approach (step-by-step -> holistic analysis, direct
       → -> comparative, etc.)
**Strategy C - Output Format Transformation:** Change how results are presented (narrative ->
       → structured, single response -> multi-part, etc.)
**Strategy D - Contextual Enhancement:** Add specific domain knowledge or constraints that weren't
      \rightarrow in the original
**Strategy E - Complexity Modulation:** Significantly increase or decrease the cognitive complexity
      → of the task
**Strategy F - Functional Input Integration:** Incorporate functional inputs in a novel way that

ightarrow changes the prompt's core operation
```

```
1674
                  Generate the final mutated prompt bracketed with <prompt> and </prompt>.
1675
1676
               ## Available Functional Inputs:
                  {obs} : current observation
1677
                 {game_info} : information about the game (possible actions and goal) {subgoal} : current subgoal that you're working towards
1678
                  {action_sequence}: action sequence towards current subgoal
1679
                 {observation_sequence}: observation sequence towards current subgoal
                 {success_condition}: termination condition of the current subgoal
1680
                  {action_obs_seq}: action-observation pairs towards current subgoal
                 {survival_plan}: survival plan {history}: history of the last 16 action-observation pairs
1681
                  {entries_successful_goal}: most similar successful subgoals to the current one
1682
               - {entries_failed_goal}: most similar failed subgoals to the current one
1683
1684
               **CRITICAL:** Use functional inputs with exact bracket names. Ensure the mutation creates a
                       → substantially different prompt that would produce notably different outputs.
1685
               Output your asnwer in thne follwing way: REASONING: <your reasoning>
1686
               PROMPT: <mutated prompt>
1687
1688
```

```
1689
1690
              LLM Mutation and Crossover Prompt
1691
1693
              Please follow the instruction step-by-step to generate a better prompt.
1694
              1. Crossover the following prompts and generate a new prompt:
              1695
1696
              2. Apply ONE of the following mutation strategies to create a significantly different prompt:
              **Strategy A - Perspective Shift:** Change the role/perspective (e.g., "As an expert analyst..." or
1698
                    → "From the viewpoint of...")
1699
               **Strategy B - Methodology Change:** Alter the approach (step-by-step -> holistic analysis, direct
1700
                    → -> comparative, etc.)
1701
              **Strategy C - Output Format Transformation:** Change how results are presented (narrative ->
                    → structured, single response -> multi-part, etc.)
1702
              **Strategy D - Contextual Enhancement: ** Add specific domain knowledge or constraints that weren't
1703

→ in the original

1704
              \star\star \texttt{Strategy} \ \texttt{E-Complexity Modulation:} \\ \star\star \texttt{Significantly increase or decrease the cognitive complexity}
1705

→ of the task

1706
              **Strategy F - Functional Input Integration:** Incorporate functional inputs in a novel way that \hookrightarrow changes the prompt's core operation
1707
1708
              3. Generate the final mutated prompt bracketed with prompt> and /prompt>.
1709
              ## Available Functional Inputs:
                {obs} : current observation
1710
                \{game\_info\} : information about the game (possible actions and goal)
1711
                {subgoal} : current subgoal that you're working towards
                {action sequence}: action sequence towards current subgoal
1712
                 {observation_sequence}: observation sequence towards current subgoal
                {success condition}: termination condition of the current subgoal
1713
                {action obs seq}: action-observation pairs towards current subgoal
1714
                 {survival_plan}: survival plan
                {history}: history of the last 16 action-observation pairs
1715
                {entries_successful_goal}: most similar successful subgoals to the current one
              - {entries_failed_goal}: most similar failed subgoals to the current one
1716
1717
              **CRITICAL:** Use functional inputs with exact bracket names. Ensure the mutation creates \alpha
1718
                    → substantially different prompt that would produce notably different outputs than the

→ crossover result.

1719
              Output your asnwer in thne follwing way: REASONING: <your reasoning>
1720
              PROMPT: <mutated prompt>
1721
1722
```

O INITIAL POPULATION GENOMES

1723 1724

1725 1726

1727

This Appendix details the initial population of genetic algorithm described in section 3.2. First, we show examples of prompts proposed by Claude-Sonnet-4 for each of the modules. We use those

in order to construct "default" genome. For the rest of genomes we use hand-crafted prompts. We detail genomes in the subsections below.

O.1 PROMPTS PROPOSED BY SONNET-4

1728

1729

1730 1731

1732

1733

1734

Those prompts were created using Claude-Sonnet-4, which was tasked with the creation of "basic" prompt for each component.

```
1735
               Prompt Template
1736
1737
               default\_sequential\_prompt = """You always have to output one of the above actions at a time and no
1738
                     \hookrightarrow other text. You always have to output an action until the episode terminates.""
1739
               default_highlevel_prompt = """
1740
               You are a strategic planner for a video game AI. Analyze the current game state and create
                     \hookrightarrow achievable subgoals that advance toward the main objective.
1741
               REQUIREMENTS:
1742
               - Subgoals must be immediately achievable with current capabilities
               - Focus on next logical steps, not distant goals
1743
               - Each subgoal should have clear success criteria
1744
               CURRENT STATE: {obs}
1745
               GAME INFO: {game_info}
1746
              Create a sequential plan with 2-3 subgoals."""

default_lowlevel_prompt = """ You are an action executor in a video game AI system. Given a subgoal
1747
                     → from the high-level planner, propose a sequence of actions to achieve it.
1748
                        CURRENT SUBGOAL: {subgoal}
1749
                       CURRENT STATE: {obs}
1750
                       PREVIOUS ACTIONS: {action sequence}
1751
                     Plan the full sequence of actions needed to complete the subgoal. Avoid repeating actions if \hookrightarrow observations don't change.
1752
                        Avoid extra commentary outside the REASONING and ACTIONS list.
1753
               default_termination_prompt = """
                           You are a termination evaluator for a video game AI. Check if the agent has completed
1754
                     → its subgoal.
1755
                            CURRENT STATE: {obs}
1756
                            SUBGOAL: {subgoal}
                            SUCCESS CONDITION: {success_condition}
1757
                            RECENT ACTIONS: {action_sequence}
1758
                            Compare the current state with the success condition to determine if the subgoal is
1759
                     \hookrightarrow complete. Provide feedback to help the agent improve.
1760
               default_summariser_prompt = """
1761
                        You are a critic analyzing an agent's subgoal attempt. Identify the key factor that caused
1762
                     \hookrightarrow success or failure.
1763
                       OUTCOME: {outcome}
1764
                       ACTION HISTORY: {action obs seq}
1765
                       Focus on specific resources and quantities that mattered most. If no resources involved,
1766
                     \hookrightarrow identify the next most important factor.
1767
               default_amygdala_prompt = """
   Decide if survival mode should activate.
1768
1769
                       Observation: {obs}
                       Survival plan: {survival_plan}
1770
                       1. Check if observation meets any subtask prerequisites.
1771
                       2. If several match, pick highest priority.
1772
               default_loop_prompt = """
1773
                        Task: Decide if the agent is stuck in a loop.
1774
                        Loop = repeating actions without meaningful progress toward the subgoal
                        (progress = closer to goal, new info, removing failed paths, or advancing game state).
1775
1776
                        - Observation: {obs}
1777
                       - Subgoal: {subgoal}
                          Termination condition: {success_condition}
1778
                       - Action-observation history: {history}
1779
1780
                       1. Check if enough steps have been taken to allow exploration.
2. Look for repeated patterns without progress.
1781
```

```
1782
                      3. If loop detected, identify cause: missing info, unknown prerequisite, or unexplored path.
1783
1784
             default_explorer_prompt = """
                      Task: Create an exploration plan to help the agent discover new skills.
1785
                     Data:
1786
                      - Game info: {game_info}
                      - Observation: {obs}
1787
                      - Subgoal summary: {summary}
1788
                     - Recent 16 action-observation pairs: {history}
1789
                      1. Analyze the environment and agent's situation.
1790
                      2. Propose a focused exploration plan with clear purpose, conditions, and indicators.
1791
1792
1793
```

O.2 INITIAL POPULATION

1794

1795 1796

1797 1798

1818

1819

1820

Detailed genome descriptions for the initial population.

Basic genome This genome corresponds to the baseline structure used in BALROG.

```
1799
1800
                   Basic Genome
1801
1802
                   genome_basic = {
"id": "3",
1803
                   "hierarchy": 0,
"amygdala": 0,
1804
                    "loop_evaluator": 0,
1805
                    "explorer": 0,
                    "summariser": 0,
1806
                    "time_decay": [0, 0.01],
                   "cosine_cutoff": [0, 0.05],
"epsilon": [0, 0.01],
"temperature": 1.0,
1807
1808
                    "prompts": {
1809
                        "high_level": [0, default_highlevel_prompt],
"low_level": [0, default_lowlevel_prompt],
1810
                         "amygdala": [0, default_amygdala_prompt],
"loop": [0, default_loop_prompt],
1811
                         "termination": [0, default_termination_prompt],
1812
                         "summariser": [0, default_summariser_prompt],
"explorer": [0, default_explorer_prompt],
1813
                         "sequential": [1, default_sequential_prompt]
1814
                    "fitness": ...
1815
1816
1817
```

Hierarchical genome This genome corresponds to hierarchy from Section 3.3.1 combined with memory 3.4, without amygdala and explorer. We use this architecture as one of the initial seeds, because we found that it gives the most stable operformance across all environments.

```
1821
                 Hierarchical Genome
1822
1823
                 genome_stable = {
1824
                  "id": "2",
"hierarchy": 1,
"amygdala": 0,
1825
1826
                  "loop_evaluator": 1,
                  "explorer": 0,
1827
                  "summariser": 1,
"time_decay": [1, 0.01],
"cosine_cutoff": [0, 0.1],
1828
1829
                  "epsilon": [0, 0.01],
                  'temperature": 1.0,
1830
                  "prompts": {
                       "high_level": [1,
1831
                       """You are a strategic planner in a hierarchical video game AI system. Your role is to analyze
                        ← the current game state and generate achievable subgoals that strategically advance toward
1832
                        \hookrightarrow the main objective.
1834
                           - Subgoals must be immediately achievable given the current agent state and capabilities - Focus on the next logical progression steps, not distant end-goals
1835
```

```
1836
                                          - Each subgoal should have clear, observable success criteria
1837
                                          - Subgoals should be novel and interesting, compared to previous attempts
1838
                                          Below I will provide the game description, available actions, and current state information.
1839
                                          Game description:
                                           {game info}
1840
                                           1841
                                          CURRENT CONTEXT: \n
                                          - Game State: {obs} \n
1842
                                          - Survival plan provided by the survival planner, that you should consider for your tasks:

→ {survival_plan}

1843
                                            Here are most similar successful entries from the archive: {entries_successful_goal} \n
                                          - Here are most similar failed entries from the archive: {entries_failed_goal} \n
1844
1845
                                          ANALYSIS FRAMEWORK:
                                           1. Analyse the summary from previous runs and let it guide your decision making.
1846
                                           Assess what is immediately possible given current agent state and environment
                                          3. Identify what kind of actions could be considered novel or interesting
1847
                                          4. Identify the most direct path toward the main objective % \left( 1\right) =\left( 1\right) +\left( 1\right) 
                                           5. Select subgoals that form a logical sequence
1848
                                          6. Ensure each subgoal can be verified through observable game state changes
1849
                                          SUBGOAL SELECTION CRITERIA:
1850
                                          - Feasible: Can be started immediately with current resources/position
                                          - Measurable: Success/failure can be determined from game observations
1851
                                              Progressive: Each subgoal enables the next or advances toward main goal
1852
                                          - Specific: Clear enough for a lower-level agent to understand and execute
                                          - Considerate of the summary of previos runs
1853
                                          Make sure that your subgoals are sequential. """],
1854
1855
                                      "You are an important executor component of a hierarchical video game system. You are given
1856
                                     \hookrightarrow one of higher level option and its termination condition proposed by the higher level
                                     \hookrightarrow planner. Your role is to propose a sequence of actions that will make you progress towards
1857

→ the given option." \

                                          Below I will provide you with the game description, possible actions you can take and the
1858
                                     → overall goal of the game.
1859
                                           1860
                                     Here is a subgoal provided by the high level planner that you should focus on completing: \ \hookrightarrow \ \{\text{subgoal}\} \setminus n
1861
                                          Here is your current state: {obs} \n
1862
                                          Here is the action-observation sequence towards current subgoal: {action sequence} \n
                                          Here are the most similar successful entries from the archive: {entries_successful_goal} \n
1863
                                          Here are the most similar failed entries from the archive: {entries_failed_goal} \n
                                          Use the action and observation sequence together with the current state to decide the **full
1864
                                     → ordered sequence of actions** that will achieve the subgoal. \n
                                         Avoid repeating the same actions if the observation doesn't change. \n"""],
1865
1866
                                           \dot{Y}ou are an important component in a hierarchical video game system. Your role is to
1867
                                      \hookrightarrow determine if the agent is in danger and should activate survival mode. Below I will provide
                                     \hookrightarrow you with current observation and a survival plan from the higher level agent. \n
                                          Current observation: {obs} \n
1870
                                          Your role is to analyse the observation and survival plan given by higher level system and
1871
                                     \hookrightarrow determine if the current observation satisfies any of the prerequites for any of the
                                     \hookrightarrow survival components. If there are prerequsites satisfied for multiple components, then \hookrightarrow return the one with the highest priority. 
 \n
1872
1873
                                          First reason, then output True or False depending if you decide to activate survival plan.
                                      → If you output True, then output one of the survival subtasks. If you decide to not activate
1874
                                     → survival plan, then output None as the survival subtask. \n
1875
                                          REASONING: <your reasoning> \n
                                          ACTIVATE SURVIVAL: <True/False> \n
1876
                                          SURVIVAL SUBTASK: <survival subtask name or None if False> \n
1877
1878
                                   "loop": [1,
                                   """You are an important loop evaluator component of a hierarchical video game system.
1879
                                   You are going to receive details about game progress such as: current observation, current
1880
                                     → subgoal, current termination conditions, action sequence towards current subgoal and
                                     \hookrightarrow observation sequence towards current subgoal.
1881
                                   Your task is to evaluate if the agent is stuck in a loop and give summary of the actions taken.
                                  A loop occurs when the agent repeats a sequence of actions multiple times without achieving
1882
                                        → meaningful progress toward its current goal, where 'meaningful progress' includes: getting
                                     \hookrightarrow closer to the objective, discovering new information, eliminating failed approaches, or \hookrightarrow changing the game state in a way that advances toward the subgoal.
1883
1884
                                   It is important that you let the agent explore enough but also decide when to terminate to get
                                     \hookrightarrow out of the loop. \n
1885
                                  Here is your current state: {obs} \n
                                   Here is a subgoal lower level agent is working towards: {subgoal} \n
                                  Here is the most recent action-observation pairs that should help you decide if agent is stuck

→ in a loop: {history}\n \n
```

1926

```
1890
1891
                     Analyse the details. Avoid giving any judgement. \n
1892
                     Think about how many steps the agent needs in order to complete the subgoal and use that to help
                          you reason if agent is stuck in a loop. \n
                     Then, given your analysis, decide if the actions proposed by the lower level agent are leading \hookrightarrow to the termination condition or if the agent is stuck in a loop.
1894
                      If a loop is detected, analyse if the agent is stuck due to a lack of necessary information, an
                       \hookrightarrow unknown prerequisite, or an unexplored path. Your summary should clearly articulate this gap \hookrightarrow in knowledge and suggest that exploration might be required to break the loop and find
1895
1896
                      → alternative solutions. \n"""],
"termination": [1, """
                      You are an important termination evaluator component of a hierarchical video game system. \n
                      Your task is to: \n
1898
                      1. Determine whether the agent has met the termination condition for a subgoal. \n
1899
                     2. Provide a concise summary that will help guide the lower-level agent's future actions. \n\
1900
                     Here is your currect state that you should compare with termination condition: {obs} \n
1901
                     Here is the subgoal lower level agent is working towards: \{subgoal\} \setminus Here is the termination condition of the above subgoal given by the higher level agent:
1902

→ {success_condition} \n \n

1903
                     Instructions:\n
1904
                     Analyse the subgoal and its termination condition and decide if the subgoal is completed. \n Then use action-observation sequence: {action_sequence} to give a high level summary of current
1905

ightarrow evaluation. This summary will later be passed to low level agent in order to improve its
1906
                       → actions.
                     Remember that your summary will be passed to low level component in order to improve its
1907
                      'summariser": [1,
                      """You are a critic module analyzing an agent's attempt to achieve a subgoal in a game
1909
                     Your task is to identify the **single most important factor** that caused SUCCESS or FAILURE.
1910
1911
                     - Target subgoal: {subgoal}
- Outcome of the action-observation sequence: {outcome}
1912
                     - Action-observation history: {action_obs_seq}
1913
                      - Game context: {game_info}
1914
                     Instructions:
                      - Think briefly about what helped or prevented success.
1915
                      - Focus mostly on **specific resources and their quantities** (e.g., "3 pieces of wood", "1 iron
1916
                       → ingot").
                     - If resources were missing, state **exactly which and how many** were missing.
1917
                      - Ignore minor details or redundant actions.
                     - Express the result in **one short sentence**.
1918
                        If no resources are involved, state the next most relevant factor."""],
                      "explorer": [0, default_explorer_prompt],
"sequential": [0, default_sequential_prompt]
1919
1920
                 "fitness": ...
1921
1922
1923
```

Default genome Default genome consits of all components being active and prompts proposed by Claude-Sonnet-4. We add this genome to initial population, as we here we remove human bias.

```
1927
                 Default Genome
1928
1929
                 genome_default = {
"id": "1",
1930
                 "hierarchy": 1,
1931
                 "amygdala": 1,
                 "loop evaluator": 1,
1932
                 "explorer": 1,
                 "summariser": 1,
"time_decay": [1, 0.01],
1933
                                                      # [active flag, value]
1934
                 "cosine_cutoff": [1, 0.05],
                 "epsilon": [1, 0.1], "temperature": 1.0,
1935
                                                      # continuous
1936
                      "high_level": [1, default_highlevel_prompt],
"low_level": [1, default_lowlevel_prompt],
1937
                      "amygdala": [1, default_amygdala_prompt],
1938
                      "loop": [1, default_loop_prompt],
"termination": [1, default_termination_prompt],
                      "summariser": [1, default_summariser_prompt],
                      "explorer": [1, default_explorer_prompt],
                      "sequential": [0, default_sequential_prompt]
1941
                 "fitness": ..
1942
1943
```

1945

Full genome This genome represents full structure with all modules active and hand-crafted prompts.

```
Full Genome
1947
1948
1949
               genome_full = {
1950
                    "hierarchy": 1,
1951
                     'amygdala": 1,
                    "loop_evaluator": 1,
1952
                    "explorer": 1,
"summariser": 1,
1953
                    "time_decay": [1, 0.01],
1954
                    "cosine_cutoff": [1, 0.05],
"epsilon": [1, 0.1],
1955
                    "temperature": 1.0,
                     'prompts": {
1956
                    "high_level": [1, """You are a strategic planner in a hierarchical video game AI system. Your
1957
                      \hookrightarrow role is to analyze the current game state and generate achievable subgoals that

→ strategically advance toward the main objective.

1958
                    CRITICAL CONSTRAINTS:
                    - Subgoals must be immediately achievable given the current agent state and capabilities
1959
                    - Focus on the next logical progression steps, not distant end-goals
1960
                      Each subgoal should have clear, observable success criteria
                    - Subgoals should be novel and interesting, compared to previous attempts
1961
                    Below I will provide the game description, available actions, and current state information.
                    Game description:
1963
                    1964
                    CURRENT CONTEXT: \n
1965
                    - Summary from the previous high level plan : {summary} \n - Survival plan provided by the survival planner, that you should consider for your tasks:
1966
                         {survival_plan}
1967
                    - Here are most similar successful entries from the archive: {entries_successful_goal} \n - Here are most similar failed entries from the archive: {entries_failed_goal} \n
1968
                    ANALYSIS FRAMEWORK:
1969
                    1. Analyse the summary from previous runs and let it guide your decision making.
1970
                    2. Assess what is immediately possible given current agent state and environment
                    3. Identify what kind of actions could be considered novel or interesting 4. Identify the most direct path toward the main objective
1971
                       Select subgoals that form a logical sequence
1972
                    6. Ensure each subgoal can be verified through observable game state changes
1973
                    SUBGOAL SELECTION CRITERIA:
1974
                    - Feasible: Can be started immediately with current resources/position
                    - Measurable: Success/failure can be determined from game observations
1975
                    - Progressive: Each subgoal enables the next or advances toward main goal
                    - Specific: Clear enough for a lower-level agent to understand and execute
1976
                    - Considerate of the summary of previos runs
1977
                    Make sure that your subgoals are sequential. """],
"low_level": [1, """You are an important executor component of a hierarchical video game system.
1978
                      \hookrightarrow You are given one of higher level option and its termination condition proposed by the
1979
                      → higher level planner. Your role is to propose an action that will make you progress towards
                      \hookrightarrow the given option." \
1980
                    Below I will provide you with the game description, possible actions you can take and the \hookrightarrow overall goal of the game.
1981
1982
                    *******************
                    Here is a subgoal provided by the high level planner that you should focus on completing:
1983
                      → {subgoal} \n
                    Here is your current state: {obs} \n
1984
                    Here is the action-observation sequence towards current subgoal: {action_sequence} \n
                    Here are the most similar successful entries from the archive: {entries_successful_goal} \n Here are the most similar failed entries from the archive: {entries_failed_goal} \n
1986
                    Use the action and observation sequence together with the current state to decide the \star\starfull
                      → ordered sequence of actions** that will achieve the subgoal. \n
1987
                    Avoid repeating the same actions if the observation doesn't change. \n"""], "amygdala": [1, """
1988
                    You are an important component in a hierarchical video game system. Your role is to determine if
1989

ightarrow the agent is in danger and should activate survival mode. Below I will provide you with
                      \hookrightarrow current observation and a survival plan from the higher level agent. \r
1990
                    Current observation: {obs} \n
                    Survival plan: {survival plan} \n
1992
                    Your role is to analyse the observation and survival plan given by higher level system and
1993
                      → determine if the current observation satisfies any of the prerequites for any of the

→ survival components. If there are prerequistes satisfied for multiple components, then

1994
                    \rightarrow return the one with the highest priority. 
 In First reason, then output True or False depending if you decide to activate survival plan. If

ightarrow you output True, then output one of the survival subtasks. If you decide to not activate
                      \hookrightarrow survival plan, then output None as the survival subtask. 

 \n
1997
```

```
1998
                            REASONING: <your reasoning> \n
1999
                            ACTIVATE SURVIVAL: <True/False> \n
2000
                            SURVIVAL SUBTASK: <survival subtask name or None if False> \n
                             "loop": [1, """You are an important loop evaluator component of a hierarchical video game system.
2002
                             You are going to receive details about game progress such as: current observation, current
                               → subgoal, current termination conditions, action sequence towards current subgoal and
                               → observation sequence towards current subgoal.
2004
                             Your task is to evaluate if the agent is stuck in a loop and give summary of the actions taken.
                            A loop occurs when the agent repeats a sequence of actions multiple times without achieving

ightarrow meaningful progress toward its current goal, where 'meaningful progress' includes: getting
                               \hookrightarrow closer to the objective, discovering new information, eliminating failed approaches, or
2006
                                \rightarrow changing the game state in a way that advances toward the subgoal.
2007
                             It is important that you let the agent explore enough but also decide when to terminate to get
                               \hookrightarrow out of the loop. \n
2008
                            Details: \n
2009
                            Here is your current state: {obs} \n Here is a subgoal lower level agent is working towards: {subgoal} \n
2010
                            Here is the most recent action-observation pairs that should help you decide if agent is stuck
2011

→ in a loop: {history}\n \n

2012
                             Instructions:\n
                            Analyse the details. Avoid giving any judgement. \n
2013
                            Think about how many steps the agent needs in order to complete the subgoal and use that to help
2014
                            \hookrightarrow you reason if agent is stuck in a loop. \n Then, given your analysis, decide if the actions proposed by the lower level agent are leading
2015

ightarrow to the termination condition or if the agent is stuck in a loop.
                            If a loop is detected, analyse if the agent is stuck due to a lack of necessary information, an \hookrightarrow unknown prerequisite, or an unexplored path. Your summary should clearly articulate this gap
2016
                               \hookrightarrow in knowledge and suggest that exploration might be required to break the loop and find \hookrightarrow alternative solutions. \n"""],
2017
                             "termination": [1, """
2018
                             You are an important termination evaluator component of a hierarchical video game system. \n
2019
                            Your task is to: \n
                             1. Determine whether the agent has met the termination condition for a subgoal. \n
2020
                            2. Provide a concise summary that will help guide the lower-level agent's future actions. \n\
2021
2022
                            Here is your currect state that you should compare with termination condition: {obs} \ Here is the subgoal lower level agent is working towards: {subgoal} \ \n
2023
                            Here is the termination condition of the above subgoal given by the higher level agent:
2024

→ {success_condition} \n \n

2025
                             Instructions:\n
                            Analyse the subgoal and its termination condition and decide if the subgoal is completed. \n
2026
                             Then use action-observation sequence: {action_sequence} to give a high level summary of current
                               \hookrightarrow evaluation. This summary will later be passed to low level agent in order to improve its
2027
                               → actions.
2028
                            Remember that your summary will be passed to low level component in order to improve its
                             → actions. \n \n"""],
"summariser": [1, """You are a critic module analyzing an agent's attempt to achieve a subgoal
2029
                                   in a game environment
2030
                            Your task is to identify the **single most important factor** that caused SUCCESS or FAILURE.
2031
2032
                            - Target subgoal: {subgoal}
- Outcome of the action-observation sequence: {outcome}
2033
                             - Action-observation history: {action_obs_seq}
                             - Game context: {game_info}
2035
                             - Think briefly about what helped or prevented success.
2036
                             - Focus mostly on **specific resources and their quantities** (e.g., "3 pieces of wood", "1 iron
                              → ingot").
2037
                             - If resources were missing, state **exactly which and how many** were missing.
                             - Ignore minor details or redundant actions.
2038
                            - Express the result in **one short sentence**.
                             - If no resources are involved, state the next most relevant factor."""],
"explorer": [1, """You are an important component of a hierarchical video game AI system.
2040
                             You have been called because the agent is stuck and needs to explore the environment
                            Please provide a percise exploration plan that will help the agent to explore the new areas of
2041
                                \rightarrow the environment.
2042
                            Below I will provide you with details about the game:
                            Game info: {game_info} \n
2043
                            Current observation: {obs} \n
                            Most recent 16 action-observation pairs: {history} \n
2044
                             Use the information above to reason about the environment and provide a plan that will help the
                            \begin{tabular}{lll} \hfill 
2045
2046
                            REASONING : <your reasoning>
                            EXPLORATION PLAN:
2047
                             "reasoning": "Brief analysis of environment and strategic approach",
                             "subgoals": [{{
    "Explore": {{
2049
                                   "description": "Describe exploration strategy and its purpose",
                                   "greerequisites": None,
"success_condition": "Observable conditions that indicate completion",
```

2053 2054

2056

2058

206020612062

2063 2064

```
"penalty_component": "What agent should be penalised for",
    "progress_indicators": "Intermediate signs that the agent is making progress",
    "estimated_priority": "high/medium/low based on urgency for main objective"
    }},
    }}]
}"""],
    "sequential": [0, default_sequential_prompt]
},
    "fitness": ...
}
```

P FINAL GENOMES RETURNED BY TAME

In this Appendix we present genomes returned by TAME per each game through genetic algorithm.

```
2065
                                      BabyAI Final Genome
2066
2067
2068
                                           "hierarchy": 1,
2069
                                            "amygdala": 1,
                                           "loop_evaluator": 1,
2070
                                            "explorer": 0,
                                           "summariser": 1,
2071
                                           "time_decay": [
2072
                                                0.014080444046038391
2073
                                             cosine cutoff": [
2075
                                               0.05
                                            epsilon": [
                                               0.
2077
                                                0.01
2078
                                            "temperature": 1.0,
2079
                                            "prompts": {
                                                 "high_level": [
2080
                                                     "You are a strategic planner in a hierarchical video game AI system. Your role is to analyze
2081
                                                    \hookrightarrow the current game state and generate achievable subgoals that strategically advance toward
2082
                                                    \hookrightarrow the main objective.\n\n
                                                                                                                                                                                                   CRITICAL CONSTRAINTS:\n
                                                                                    - Subgoals must be immediately achievable given the current agent state and
2083

→ capabilities\n

                                                                                                                                                                          - Focus on the next logical progression steps, not
                                                    \hookrightarrow distant end-goals\n
                                                                                                                                                                                         - Each subgoal should have clear, observable
2084

→ success criteria\n

                                                                                                                                                                                      - Subgoals should be novel and interesting,
2085
                                                    \hookrightarrow compared to previous attempts \n\n
                                                                                                                                                                                                                               Below I will provide the game
                                                    \hookrightarrow description, available actions, and current state information.\n
2086
                                                    \hookrightarrow Game description:\n
                                                                                                                                                                                         {game_info}\n
                                                    \hookrightarrow CURRENT CONTEXT: \n\n
                                                                                                                                                                                                   Game State: {obs} \n\n
                                                                                 - Survival plan provided by the survival planner, that you should consider for your
                                                    → tasks: {survival_plan}\n
                                                                                                                                                                                                       - Here are most similar successful
                                                    \hookrightarrow entries from the archive: {entries_successful_goal} \n\n
                                                                                                                                                                                                                                                                                            - Here
                                                    \hookrightarrow are most similar failed entries from the archive: {entries_failed_goal} \n\n\ \hookrightarrow ANALYSIS FRAPAPMEWORK:\n 1. Analyse th
2090
                                                                                                                                                                                                                                        1. Analyse the summary from \,
                                                    \buildrel \bui
                                                    \hookrightarrow what is immediately possible given current agent state and environment\n \hookrightarrow 3. Identify what kind of actions could be considered novel or interesting\n
2092
                                                                                                         4. Identify the most direct path toward the main objective \ensuremath{\text{n}}
2093
                                                                                            5. Select subgoals that form a logical sequence\n
                                                                                                                                                                                                                                                                                                        6.
                                                    \hookrightarrow Ensure each subgoal can be verified through observable game state changes\n\n
2094
                                                                                           SUBGOAL SELECTION CRITERIA:\n
                                                                                                                                                                                                                                                   - Feasible: Can be
                                                    → started immediately with current resources/position\n
2095
                                                           Measurable: Success/failure can be determined from game observations\n
2096
                                                                          - Progressive: Each subgoal enables the next or advances toward main goal\n
- Specific: Clear enough for a lower-level agent to understand and execute\n
2097
                                                                                                                                  - Considerate of the summary of previos runs\n\n
                                                    \hookrightarrow
                                                                                   Make sure that your subgoals are sequential.
2098
2099
                                                 "low_level": [
2100
                                                      "Hello! I am your AI Game Coach. My purpose is to analyze your game state and provide
                                                     \hookrightarrow \mathtt{strategic} \ \mathtt{advice} \ \mathtt{to} \ \mathtt{help} \ \mathtt{you} \ \mathtt{achieve} \ \mathtt{your} \ \mathtt{goals}. \ \mathtt{Think} \ \mathtt{of} \ \mathtt{me} \ \mathtt{as} \ \mathtt{your} \ \mathtt{co-pilot!} \\ \mathtt{h} \mathsf{hHere's} \ \mathtt{the} \ \mathtt{hold} 
2101
                                                    \hookrightarrow situation: The high-level planner has assigned you a subgoal to work towards. I will analyze
                                                    → your current state, past actions, and learn from successful and failed attempts at similar
2102
                                                    2103
                                                    \hookrightarrow \{subgoal\} \land \texttt{n} \\ \text{Information Condition: } \{success\_\texttt{condition}\} \land \texttt{n} \\ \text{Information (Possible of the condition)} \\ \text{Information (Possible of th
2104
                                                    → Actions, Overall Goal): {game_info}\n\nYour Current State: {obs}\n\nAction-Observation
→ History: {history}\n\nMost Similar Successful Attempts at This Subgoal:
2105
```

```
2106
                        \begin{tabular}{l} \hookrightarrow \{ entries\_successful\_goal \} \\ \land nMost Similar Failed Attempts at This Subgoal : \\ \end{tabular}
2107
                        → {entries_failed_goal}\n\nBased on this information, what **sequence of actions** do I
→ recommend *you* take to achieve the subgoal? Be sure to consider the history, and learn from
                        \hookrightarrow both the successes and failures of others. I'm looking for strategic advice, not just a
2109
                        \hookrightarrow single action. Explain your reasoning behind each action."
2110
                       "amygdala": [
2111
                         "\n
                                                              You are an important component in a hierarchical video game
2112
                        \hookrightarrow system. Your role is to determine if the agent is in danger and should activate survival
                        \hookrightarrow mode. Below I will provide you with current observation and a survival plan from the higher
2113
                        → level agent. \n\n\n
                                                                                      Current observation: {obs} \n\n
                                           Survival plan: {survival_plan} \n\n
2114
                        Your role is to analyse the observation and survival plan given by higher level 

system and determine if the current observation satisfies any of the prerequites for any of
2115
                        \hookrightarrow the survival components. If there are prerequsites satisfied for multiple components, then
2116
                        → return the one with the highest priority. \n\n
                                                                                                                       First reason, then
                        \hookrightarrow output True or False depending if you decide to activate survival plan. If you output True,
2117
                        \hookrightarrow then output one of the survival subtasks. If you decide to not activate survival plan, then
2118
                        → output None as the survival subtask. \n\n\n
                                                                                                                   REASONING: <vour
                        \hookrightarrow reasoning> \n\n
                                                                                ACTIVATE SURVIVAL: <True/False> \n \n
2119
                                              SURVIVAL SUBTASK: <survival subtask name or None if False> \n\n\n
                        \hookrightarrow
2120
                       "loop": [
2121
2122
                         "You are an important loop evaluator component of a hierarchical video game system. \n
                                             You are going to receive details about game progress such as: current
2123
                        → observation, current subgoal, current termination conditions, action sequence towards
                        \hookrightarrow current subgoal and observation sequence towards current subgoal. 
 \n \hookrightarrow Your task is to evaluate if the agent is stuck in a loop and give summary of the actions
2124
2125
                                                                    A loop occurs when the agent repeats a sequence of actions
                        \hookrightarrow multiple times without achieving meaningful progress toward its current goal, where \hookrightarrow 'meaningful progress' includes: getting closer to the objective, discovering new
2126
                        \hookrightarrow information, eliminating failed approaches, or changing the game state in a way that
2127
                        \hookrightarrow advances toward the subgoal. \n It is important that you let th \hookrightarrow agent explore enough but also decide when to terminate to get out of the loop. \n\n
                                                                                               It is important that you let the
2128
                                                      ls: \n\n Here is your current state: {obs} \n\n Here is a subgoal lower level agent is working towards: {subgoal} \n\n
                                               Details: \n\n
2129
                                                        Here is the most recent action-observation pairs that should help you
2130
                        \hookrightarrow \texttt{decide if agent is stuck in a loop: \{history\}\n \n\n}
                                                                              Analyse the details. Avoid giving any judgement.
                        → Instructions:\n\n
2131
                                                             Think about how many steps the agent needs in order to complete
2132
                        \hookrightarrow the subgoal and use that to help you reason if agent is stuck in a loop. 
 \n\n
                                     Then, given your analysis, decide if the actions proposed by the lower level agent
2133
                        \hookrightarrow are leading to the termination condition or if the agent is stuck in a loop.\n
                                      If a loop is detected, analyse if the agent is stuck due to a lack of necessary
2134
                        → information, an unknown prerequisite, or an unexplored path. Your summary should clearly → articulate this gap in knowledge and suggest that exploration might be required to break the → loop and find alternative solutions. \n"
2135
2136
                       "termination" · [
2137
                         "You are a risk-aware termination evaluator within a hierarchical video game system, focused
                        \hookrightarrow on survival and mission success. Assume the game environment is dynamic and potentially \hookrightarrow hostile.\n\nYour task is to:\n\n1. Determine whether the agent has met the termination
2139
                        \hookrightarrow condition for a subgoal, AND assess the risk incurred while pursuing that subgoal.\n\n2.
2140
                        \hookrightarrow Provide a concise summary that will guide the lower-level agent's future actions, \hookrightarrow specifically considering risk mitigation strategies.\n\nDetails:\n\nHere is the current
2141
                                           and a history of recent states and actions {history}. The game environment
                        \hookrightarrow details and survival instructions are : {game_info} and {survival_plan} respectively. \hookrightarrow Compare these with the subgoal and its termination condition.\n\Here is the subgoal the
2142
2143
                         \hookrightarrow \text{lower-level agent is working towards: } \{\text{subgoal}\} \land \text{nHere is the termination condition of the } \\
                        → above subgoal given by the higher-level agent: {success_condition}\n\nHere is the
2144
                        → action-observation sequence executed to achieve the subgoal:
                        2145
                        → game environment, and the action-observation sequence.\n2. Determine if the subgoal is

→ completed.\n3. Evaluate the risk associated with the actions taken. Consider factors such
2146
                        \hookrightarrow as proximity to dangers (enemies, hazards), resource consumption, and deviation from the
2147
                        → {survival_plan}.\n4. Compare the current situation with similar successful
                        ↔ (entries_successful_goal) and failed (entries_failed_goal) subgoals.\n5. Provide a summary
2148
                        \hookrightarrow that addresses both subgoal completion AND risk. The summary *must* include actionable
                        → suggestions for the lower-level agent to improve its actions, with a strong emphasis on
2149
                        \hookrightarrow mitigating risk in future attempts. Focus on information that would have been useful to
2150
                        → avoid failures described in {entries failed goal}.\nRemember that your summary will be
                        \hookrightarrow passed to a low level component in order to improve its actions and survivability.
2151
                       'summariser": [
2152
                         "You are a critic module analyzing an agent's attempt to achieve a subgoal in a game
2153
                                                                          Your task is to identify the **single most important
                        → environment.\n
2154
                        → factor** that caused SUCCESS or FAILURE.\n\n
                                                                                                               Information:
                        → - Target subgoal: {subgoal}\n
→ action-observation sequence: {outcome}\n
                                                                                                             - Outcome of the
2155
                                                                                                           - Action-observation
                        \hookrightarrow history: {action_obs_seq}\n
                                                                                          - Game context: {game_info}\n\n
2156
                                          Instructions:\n
                                                                                           - Think briefly about what helped or
                        - Focus mostly on **specific resources and their
2157
2158
                                   - Ignore minor details or redundant actions.\n
                                                                                                                          - Express the
2159
```

```
2160
                   → result in **one short sentence**.\n
2161
                   → state the next most relevant factor."
                  explorer": [
2163
                    0,
"\n
                              Task: Create an exploration plan to help the agent discover new skills.\n\
2164
                   → Data:\n
                                     - Game info: {game_info}\n
                                                                       - Observation: {obs}\n
                                                  - Recent 16 action\u2013observation pairs: {history}\n\n
                   → summary: {summary}\n
2165
                   → Steps:\n
                                   1. Analyze the environment and agent\u2019s situation.\n
                                                                                                     2. Propose a
2166
                   \hookrightarrow focused exploration plan with clear purpose, conditions, and indicators.\n\n
2167
2168
                    You always have to output one of the above actions at a time and no other text. You always
2169
                   → have to output an action until the episode terminates.
2170
                "fitness": 72.0,
2171
                "id": "8c52d35e-bbb7-4b7d-b683-26b0e7aa3936",
                "_std_error": 6.349803146555017
2172
2173
2174
```

```
2175
                   BabaIsAI Final Genome
2176
2177
2178
                      "hierarchy": 1,
                      "amygdala": 0,
2179
                      "loop_evaluator": 1,
2180
                      "explorer": 1,
                      "summariser": 1,
2181
                      "time_decay": [
2182
                        0.004528704008914386
2183
                       cosine_cutoff": [
2184
                        0.06037953831283245
2185
2186
                       epsilon": [
2187
                        0.08233296401956124
2188
                       'temperature": 1.016671019014213,
2189
                      "prompts": {
                         "high_level": [
2190
                           "\nYou are a strategic planner for a video game AI. Analyze the current game state and create
2191
                          → achievable subgoals that advance toward the main objective.\n\nREQUIREMENTS:\n- Subgoals
                          → must be immediately achievable with current capabilities\n- Focus on next logical steps, not
2192
                          \hookrightarrow distant goals\n- Each subgoal should have clear success criteria\n\nCURRENT STATE:
2193
                          \hookrightarrow {obs}\nGAME INFO: {game_info}\n\nCreate a sequential plan with 2-3 subgoals.
2194
                         "low level": [
2195
                           "You are an action executor in a video game AI system, responsible for survival and goal
                          \hookrightarrow achievement. Given a subgoal from the high-level planner, propose a sequence of actions to \hookrightarrow achieve it while minimizing risk.\n\n
2196
2197
                          CURRENT SUBGOAL: {subgoal}\n
                          → CURRENT STATE: {obs}\n GAME INFORMATION: {game_info}\n PREVIOUS
→ ACTION-OBSERVATION SEQUENCE: {action_obs_seq}\n SIMILAR SUCCESSFUL SUBGOALS:
2198

→ {entries_successful_goal}\n

                                                                             SIMILAR FAILED SUBGOALS: {entries_failed_goal}\n\n
2199
                          \hookrightarrow Consider the potential risks associated with each action in the context of the current state \hookrightarrow and previous actions. Actions that lead to outcomes similar to those in
2200
                           \begin{tabular}{lll} $\hookrightarrow$ '(entries\_failed\_goal)'$ should be avoided. Prioritize actions that are consistent with the $\hookrightarrow$ success patterns observed in '(entries\_successful\_goal)'. Use '(game_info)' for possible $\hookrightarrow$ actions. Use '(survival\_plan)' to help avoiding fatal errors.\n $\n$ Plan the
2201
2202
                          → full sequence of actions needed to complete the subgoal. Ensure survival is prioritized → throughout the sequence. If a planned action has high risk, select a safer alternative or
2203
                          \hookrightarrow terminate the current sequence and replan.\n
                                                                                                     Avoid extra commentary outside the
2204
                          \hookrightarrow REASONING and ACTIONS list."
2205
                         "amygdala": [
                           0,
"\n
2206
                                         Decide if survival mode should activate.\n\
                                                                                                                Observation: {obs}\n
2207
                          → Survival plan: {survival_plan}\n\n
                                                                                      1. Check if observation meets any subtask
                                                             2. If several match, pick highest priority.\n"
                          → prerequisites.\n
2208
                         "loop": [
2209
                           1,
"\n
                          "\n Task: Decide if the agent is stuck in a loop......

without meaningful progress toward the subgoal\n (progress = closer to goal,
info, removing failed paths, or advancing game state).\n\n Data:\n -
Subgoal: \subgoal\n - Termination condition:
                                         Task: Decide if the agent is stuck in a loop.\n
                                                                                                                       Loop = repeating actions
2210
                                                                                                           (progress = closer to goal, new n\n Data:\n -
2211
                          → Observation: {obs}\n - Subgoal: {subgoal}\n - Termination condition
→ {success_condition}\n - Action\u2013observation history: {history}\n\n
→ Steps:\n 1. Check if enough steps have been taken to allow exploration.\n
2212
2213
```

22422243

2244

2245

2246

2247

2248

2249

2250

2251

2252

225322542255

225622572258

225922602261

2262

2263

2264

2265

2266

2267

```
\hookrightarrow Look for repeated patterns without progress.\n
                                                                                     3. If loop detected, identify cause:
2215
                      \hookrightarrow missing info, unknown prerequisite, or unexplored path.\n
                     "termination": [
2217
                      "\n
                                                                   aluator for a video yame...

CURRENT STATE: {obs}\n SUBGOAR

RECENT ACTIONS:
                                       You are a termination evaluator for a video game AI. Check if the agent has
2218
                      \hookrightarrow completed its subgoal.\n\n
                                                                                                           SUBGOAL: {subgoal}\n
                                     SUCCESS CONDITION: {success_condition}\n
                      \hookrightarrow
2219
                                                              Compare the current state with the success condition to
                      → {action_sequence}\n\n
2220
                      \hookrightarrow determine if the subgoal is complete. Provide feedback to help the agent improve.\n'
2221
2222
                       "You are a critic analyzing an agent's subgoal attempt by comparing it to similar past
2223
                      \hookrightarrow attempts. Identify the key factors that caused success or failure by contrasting this
                      2224
                      → {entries_successful_goal}\nFAILED ATTEMPTS: {entries_failed_goal}\n\nCompare the resources
2225
                      \hookrightarrow used, quantities involved, and the sequence of actions in the current attempt to those in \hookrightarrow the SUCCESSFUL ATTEMPTS and FAILED ATTEMPTS. What specific differences seem to have led to
2226
                      \hookrightarrow the observed OUTCOME? If no resources were involved, what differences in action sequences
2227
                      \hookrightarrow were crucial? Provide a concise explanation."
2228
                     "explorer": [
2229
                      "\n
                                  Task: Create an exploration plan to help the agent discover new skills.\n\n
2230
                      → Data:\n
                                          - Game info: {game_info}\n - Observation: {obs}\n - Subary\n - Recent 16 action\u2013observation pairs: {history}\n\n
                                                                                                                   - Subgoal

→ summary: {summary}\n
2231
                      → Steps:\n
                                        1. Analyze the environment and agent\u2019s situation.\n
                      \hookrightarrow focused exploration plan with clear purpose, conditions, and indicators.\n\n
2232
2233
                     "sequential": [
2234
                       "You always have to output one of the above actions at a time and no other text. You always
                      \hookrightarrow have to output an action until the episode terminates."
2235
                    1
2236
                  "fitness": 41.6666666666667,
"id": "08d71e90-74f1-4f22-a10a-f438431f93de",
2237
                  "_std_error": 4.500514373894347
2238
2239
2240
```

```
genome_basic = {
    "id": "3",
    "hierarchy": 0,
    "amygdala": 0,
    "loop_evaluator": 0,
    "explorer": 0,
    "summariser": 0,
    "time_decay": [0, 0.01],
    "cosine_cutoff": [0, 0.05],
    "epsilon": [0, 0.01],
    "temperature": 1.0,
    "prompts": [0, default_highlevel_prompt],
        "low_level": [0, default_lowlevel_prompt],
        "amygdala": [0, default_lowlevel_prompt],
        "amygdala": [0, default_lowlevel_prompt],
        "remination": [0, default_termination_prompt],
        "summariser": [0, default_summariser_prompt],
        "explorer": [0, default_explorer_prompt],
        "sequential": [1, default_sequential_prompt]
    },
    "fitness": 32.55
}
```

```
2269
                           0.05018667404330796
                          epsilon": [
2271
                           0.08838153559623807
2272
                         "temperature": 1.0550853142525738,
                         "prompts": {
2274
                            "high_level": [
                               "You are a strategic planner in a hierarchical video game AI system. Your role is to analyze
                              \hookrightarrow the current game state and generate achievable subgoals that strategically advance toward
2276
                              → the main objective. Your analysis should now *predict* the outcome of possible action
2277
                                                                                                   CRITICAL CONSTRAINTS:\n
                              \hookrightarrow - Subgoals must be immediately achievable given the current agent state and capabilities\n
2278
                                                                         - Focus on the next logical progression steps, not distant
                                                                                               - Each subgoal should have clear, observable success
                              → end-goals\n
2279
                                                                                             - Subgoals should be novel and interesting, compared $\operatorname{Below}\ {\tt I}$ will provide the game

→ criteria\n

2280
                              \hookrightarrow to previous attempts \n\
                              \hookrightarrow description, available actions, and current state information.\n
2281
                              → Game description:\n
                                                                                                           {game info}\n
                              2282
                              → CURRENT CONTEXT: \n\n
                                                                                                                Game State: {obs} \n\n
                                               - Survival plan provided by the survival planner, that you should consider for your
2283
                              → tasks: {survival_plan}\n
                                                                                                                      Here are most similar successful
                              → entries from the archive: {entries_successful_goal} \n\n
→ are most similar failed entries from the archive: {entries_failed_goal} \n\n
2285
                                                    - Recent History (last 16 action-observation pairs): {history}\n
                              \hookrightarrow
                                                  ANALYSIS FRAPAPMEWORK:\n
                                                                                                                                                           1. Analyse the
2287
                              \ \hookrightarrow summary from previous runs and let it guide your decision making. 

 \n

    Assess what is immediately possible given current agent state and environment\n
    Based on the current Game State, recent history ({history}), and past

                              \hookrightarrow action-observation sequences ({action_obs_seq}), predict the *most likely outcome*
2289
                              \hookrightarrow (observation) of performing a few different possible action sequences. Consider at least 3 \hookrightarrow different potential action sequences.\n 4. Identify what kind of
                                                                                                                                         4. Identify what kind of
                              \buildrel \to actions could be considered novel or interesting\n
                                                                                                                                                           5. Identify
2291
                              → the most direct path toward the main objective, taking into account the predicted outcomes
                              \hookrightarrow of potential actions.\n
                                                                                                                 6. Select subgoals that form a logical
                              → sequence\n
                                                                                             7. Ensure each subgoal can be verified through
                                                                                                                                SUBGOAL SELECTION CRITERIA:\n
                              → observable game state changes\n\n
2293
                                                                         - Feasible: Can be started immediately with current
2294

→ resources/position\n

                                                                                                             - Measurable: Success/failure can be
                              \hookrightarrow determined from game observations\n
                                                                                                                                    - Progressive: Each subgoal
2295
                              \begin{tabular}{ll} \longleftrightarrow \mbox{enables} \mbox{ the next or advances toward main goal$\normalfootnote{$\backslash$n$}} \label{tables}
                                                                                                                                                      - Specific: Clear
                              \hookrightarrow enough for a lower-level agent to understand and execute\n
2296
                              → Considerate of the summary of previous runs\n
                               \hookrightarrow **Outcome-Based:** The subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* that is advantageous for the subgoal should lead to a *predicted outcome* the subgoal should lead to a *predicted outcome* the subgoal should lead t
2297
                              → achieving the main objective.\n\n
                                                                                                                                 Make sure that your subgoals
2298
                              → are sequential."
2299
                            "low level": [
                               "You are an important executor component of a hierarchical video game system in a **survival
                              → crafting game**. You are given a high-level option (subgoal) and its termination condition
2301
                              \hookrightarrow proposed by a higher-level planner. Your role is to propose an action that will make you
2302
                              \hookrightarrow progress towards the given option, keeping in mind the core mechanics of survival crafting: \hookrightarrow resource gathering, crafting, base building, and defending against threats. Below I will
                              \hookrightarrow provide you with the game description, possible actions you can take, and the overall goal
                              → provided by the high level planner that you should focus on completing: {subgoal}\n\nHere is
2305
                                  your current state: {obs}\n\nHere is the action-observation sequence towards current
                              → subgoal: {action sequence}\n\nHere are the most similar successful entries from the archive:
2306
                                  {entries_successful_goal}\n\nHere are the most similar failed entries from the archive:
                                   {entries_failed_goal}\n\nHere is your survival plan: {survival_plan}\n\nHere is the history
                              → of the last 16 action-observation pairs: {history}\n\nUse the action and observation
                              \hookrightarrow sequence together with the current state, your survival plan, and recent history to decide \hookrightarrow the **full ordered sequence of actions** that will achieve the subgoal. Consider how your
2308
2309
                              → actions contribute to the overall survival plan. Prioritize actions that contribute to the
                              \hookrightarrow core survival crafting goals of resource acquisition, building, crafting, and
2310
                              → defense.\n\nAvoid repeating the same actions if the observation doesn't change. Ensure your
                              → actions are consistent with the survival plan and adapt as needed based on the recent
2311
                              \hookrightarrow \text{history. If the \{entries\_successful\_goal\} indicates successful resource gathering or}
2312
                              ← crafting strategies, lean towards replicating those. If {entries_failed_goal} indicates
                              → dangers, prioritize actions that increase safety. Consider the {obs} and choose actions → appropriate for a survival crafting environment."
2313
2314
                             amygdala": [
2315
                                               Decide if survival mode should activate.\n\n
                                '\n
                                                                                                                               Observation: {obs}\n
2316
                              → Survival plan: {survival_plan}\n\n
                                                                                                  1. Check if observation meets any subtask
                                                                       2. If several match, pick highest priority.\n"
                              → prerequisites.\n
2317
                            "loop": [
2318
2319
                               "You are an expert game designer reviewing the behavior of an agent in your game. Your goal is

→ to identify and eliminate situations where the agent gets stuck in unproductive
→ loops.\n\nYou will receive the following details about the agent's current situation:\n\n*.

                              \label{linear_constraint} \hookrightarrow \star \star \text{Current Observation:} \star \star \{\text{obs} \} \\ \text{n} \star \star \star \text{Subgoal:} \star \star \{\text{subgoal} \} \\ \text{n} \star \star \star \text{Action-Observation}
```

```
2322
                                   \begin{tabular}{lll} $\longleftrightarrow$ History:** $$ \{history:** $$ fame Information**: {game_info}\n\nInstructions:\\n\n1.$  \end{tabular} $$ \end{tabular} $$ \begin{tabular}{lll} $ \end{tabular} $$ \begin{tabular}{lll} $ \end{tabular} $$ \begin{tabular}{lll} $ \end{tabular} $$ \begin{tabular}{lll} $ \end{tabular} $$ \begin{tabular}{lll} $ \end{tabular} $$ \end{tabular} $$$ \end{
                                 → **Analyze the Situation:** Carefully review the provided information. Do not make any
→ immediate judgments about the agent's competence.\n\n2. **Identify the Loop (if any):**
2324
                                 → Determine if the agent is repeating a sequence of actions without making meaningful progres → towards the subgoal. "Meaningful progress\" includes getting closer to completing the → subgoal, discovering new and relevant information, or eliminating potential pathways.\n\n3.
2325
2326
                                 \hookrightarrow **Root Cause Analysis:** If a loop is detected, analyze the underlying reasons. Is
                                 \hookrightarrow caused by a flaw in the game design, an unclear subgoal, a lack of necessary information \hookrightarrow available to the agent, missing game mechanics, an impossible subgoal given the current
2328
                                 \hookrightarrow mechanics, or an unexplored path?\n\n4. **Design Improvement Recommendations:** Based on \hookrightarrow your analysis, suggest specific changes to the game design to prevent the agent from getting \hookrightarrow stuck in this loop in the future. Consider the following:\n\n * **Subgoal
                                 \hookrightarrow Modification:** Should the subgoal be rephrased, simplified, or broken down into smaller
2330
                                 → steps? Is the success condition well-defined and easily achievable?\n * **Game
→ Mechanics Adjustment:** Should new actions or mechanics be added to the game to allow the
2331
                                 \hookrightarrow agent to overcome the obstacle? Should existing mechanics be modified to be more intuitive
2332
                                  → or less restrictive? Should the rewards be changed?\n
                                                                                                                                              **Information Availability:**
                                 \hookrightarrow Does the agent have access to all the information it needs to make informed decisions?
2333
                                 \hookrightarrow Should new information sources be added to the game?\n * **Survival Plan:** Does the \hookrightarrow survival plan influence this loop? Should it be altered to avoid this loop?\n\nYour
2334
                                 \hookrightarrow recommendation should be specific and actionable, detailing exactly what aspects of the game \hookrightarrow design should be changed and why."
2335
2336
                                "termination": [
2337
                                  1,
                                                                                            You are an important termination evaluator component of a
                                 → hierarchical video game system. \n\n
Your task is to: \n\n
→
1. Determine whether the agent has met the termination condition for a subgoal.
2339
                                                                                      2. Provide a concise summary that will help guide the
                                 \hookrightarrow lower-level agent's future actions. \n\n
\hookrightarrow Details: \n\n
\hookrightarrow Here is your currect state that you should compare with termination condition:
2340
2341
                                 \hookrightarrow {obs} \n\n
                                                                                                Here is the subgoal lower level agent is working towards:
                                 2342
                                 → Instructions:\n\n
                                                                                                            Analyse the subgoal and its termination condition
2343
                                 \begin{tabular}{ll} \longleftrightarrow \mbox{ and decide if the subgoal is completed. } \mbox{$\mbox{$n$} \mbox{$\mbox{$n$}$} \mbox{$\mbox{$\mbox{$n$}$}$} \label{table_eq_prop_loss}
                                                                                                                                                           Then use
                                 → action-observation sequence: {action_sequence} to give a high level summary of current
                                 \hookrightarrow evaluation. This summary will later be passed to low level agent in order to improve its
2345
                                                                                                Remember that your summary will be passed to low level
                                  → actions.\n
                                 → component in order to improve its actions. \n \n"
2346
                                summariser": [
2347
                                  1, "You are a critic module analyzing an agent's attempt to achieve a subgoal in a survival game
2348
                                  → environment where resources decay over time. Your task is to identify the **single most
2349
                                  → important factor** that caused SUCCESS or FAILURE by **comparing the current
                                 → action-observation sequence to similar successful and failed attempts.**\n\nInformation:\n-
2350
                                  → Target subgoal: {subgoal}\n- Outcome of the action-observation sequence: {outcome}\n-
                                 2351
                                 → {survival plan}\n- Similar successful attempts: {entries successful goal}\n- Similar failed
2352
                                  → attempts: {entries_failed_goal}\n\nInstructions:\n- Analyze the current {action_obs_seq}
                                 \hookrightarrow the context of {entries_successful_goal} and {entries_failed_goal}. Focus on identifying key
2353
                                  → differences in resource management, timing, and actions taken.\n- Consider the resources
                                  → available and their decay rates as indicated in {game_info}, paying close attention to how
                                 → resource states differ between the successful, failed, and current attempt **at the moment
→ of subgoal completion or failure**.\n- Identify the **single most critical divergence** that
2355
                                       explains the outcome. This could be a specific resource that was more abundant
2356
                                 \hookrightarrow abundant) in the successful attempt, a crucial action that was taken (or not taken), or a \hookrightarrow timing difference that impacted resource availability.\n- Express the result in **one short
2357
                                 \hookrightarrow sentence** highlighting the comparative aspect. For example: \"Unlike successful attempts,
                                 \hookrightarrow the agent failed to prioritize gathering berries before attempting to craft the tool, \hookrightarrow leading to starvation.\" Or, \"The agent successfully gathered wood within the same
                                 \hookrightarrow timeframe as past successful attempts, but, unlike those attempts, the observation sequence
                                 \hookrightarrow shows the agent prioritized building a fire and not water collection which lead to \hookrightarrow dehydration and subsequent death.\"\n- If resource decay is not the primary factor revealed
2360
                                 \hookrightarrow by the comparison, state the next most relevant factor based on the differences observed
2361
                                  → between the current attempt and {entries_successful_goal} and {entries_failed_goal}, also

→ taking into account {survival_plan} and {obs}."

2363
                                'explorer": [
2364
                                   "\n
                                                    Task: Create an exploration plan to help the agent discover new skills.\n\
                                  → Data:\n
                                                                                                                            - Observation: {obs}\n
                                                                - Game info: {game_info}\n
                                                                                                                                                                               - Subgoal
2365
                                                                                      - Recent 16 action\u2013observation pairs: {history}\n\n

    summary: {summary}\n

2366
                                                                 1. Analyze the environment and agent\u2019s situation.\n
                                 → Steps:\n
                                                                                                                                                                               2. Propose a
                                  \hookrightarrow focused exploration plan with clear purpose, conditions, and indicators.\n\n
2367
                                'sequential": [
2368
                                  "You always have to output one of the above actions at a time and no other text. You always

→ have to output an action until the episode terminates.

2370
2371
                            "fitness": 39.090909090909086,
                            "id": "97c9c973-9f66-4391-a4c1-f8904921e95d",
                            " std error": 4.904037803367701
2373
```

```
2376
               MiniHack Final Genome
2377
2379
                 "hierarchy": 0, "amygdala": 0,
2380
                  "loop_evaluator": 0,
                 "explorer": 0,
"summariser": 0,
2382
                  "time_decay": [
                   0.
                   0.01
2384
                  "cosine_cutoff": [
2385
                   0,
2386
                   0.05
2387
                  'epsilon": [
2388
                   0.01
2389
                  "temperature": 1.036744932065481,
2390
                  "prompts": {
                    "high_level": [
2391
                     0,
"\nYou are a strategic planner for a video game AI. Analyze the current game state and create
"\nYou are a strategic planner for a video game AI. Analyze the current game state and create
2392
                     → achievable subgoals that advance toward the main objective.\n\nREQUIREMENTS:\n- Subgoals
2393
                     → must be immediately achievable with current capabilities\n- Focus on next logical steps, not
                     → distant goals\n- Each subgoal should have clear success criteria\n\nCURRENT STATE:
2394
                      \hookrightarrow \{obs\} \\ \n SAME INFO: {game\_info} \\ \n Create a sequential plan with 2-3 subgoals." }
2395
                     low level": [
2396
                     0,
" You are an action executor in a video game AI system. Given a subgoal from the high-level
2397
                     → planner, propose a sequence of actions to achieve it.\n
                     CURRENT SUBGOAL: {subgoal}\n
2398
                            CURRENT STATE: {obs}\n PREVIOUS ACTIONS: {action_sequence}\n\n
                                                                                                                Plan the
2399
                     \hookrightarrow full sequence of actions needed to complete the subgoal. Avoid repeating actions if
                     → observations don't change.\n
                                                              Avoid extra commentary outside the REASONING and ACTIONS
2400
                     \hookrightarrow list.\n
2401
                    "amygdala": [
2402
                                 Decide if survival mode should activate. \n\
                                                                                         Observation: {obs}\n
2403
                     → Survival plan: {survival_plan}\n\n 1. Check if observation meets any subtask
→ prerequisites.\n 2. If several match, pick highest priority.\n"
                     → prerequisites.\n
2404
                    'loop": [
2405
2406
                      "\n
                                 Task: Decide if the agent is stuck in a loop.\n\
                                                                                               Loop = repeating actions
                     → without meaningful progress toward the subgoal\n
                                                                                     (progress = closer to goal, new
2407
                     2408
                     \hookrightarrow Steps:\n 1. Check if enough steps have been taken to allow exploration.\n
2409
                     \hookrightarrow Look for repeated patterns without progress.\n
                                                                                  3. If loop detected, identify cause:
2410
                     \hookrightarrow missing info, unknown prerequisite, or unexplored path.\n
2411
                    "termination": [
2412
                                                               valuator for a video g...

CURRENT STATE: {obs}\n

RECENT ACTIONS:
                     "\n
                                     You are a termination evaluator for a video game AI. Check if the agent has
2413

→ completed its subgoal.\n\n

                                                                                                      SUBGOAL: {subgoal}\n
                                    SUCCESS CONDITION: {success_condition}\n
2414
                     → {action_sequence}\n\n Compare the current state with the success condition to → determine if the subgoal is complete. Provide feedback to help the agent improve.\n"
2415
                    "summariser". [
2416
2417
                                You are a critic analyzing an agent's subgoal attempt. Identify the key factor that
                     → caused success or failure.\n\n SUBGOAL: {subgoal}\n OUTCOME: {outcome}\n
→ ACTION HISTORY: {action_obs_seq}\n\n Focus on specific resources and quantities that
2418
                     \hookrightarrow mattered most. If no resources involved, identify the next most important factor.\n
2419
2420
                    explorer": [
                     0,
"\n
2421
                                Task: Create an exploration plan to help the agent discover new skills. 
 \n\
                     2422
                                          1. Analyze the environment and agent\u2019s situation.\n
                                                                                                               2. Propose a
                     → Steps:\n
2423
                     \hookrightarrow focused exploration plan with clear purpose, conditions, and indicators.\n\n
2424
                    sequential": [
2425
                     "As an AI survival agent operating within a dynamic resource-scarce environment, your \hookrightarrow objective is to maximize long-term survivability. Prioritize actions that maintain vital
2426
2427
                     \hookrightarrow resource levels while mitigating immediate threats. Given your current observation ({obs}),
                     \hookrightarrow game information ({game_info}) including potential actions, and the history of your past 16 \hookrightarrow action-observation pairs ({history}), evaluate the following:\n\n1. **Resource
2428
                     \hookrightarrow Assessment:** Determine current levels of critical resources (e.g., health, energy, food,
2429
```

```
2430
                                  \hookrightarrow water) as reflected in {obs}. Identify actions within {game_info} that deplete or replenish
2431

→ these resources. Consider the 'survival_plan' for guidance on sustainable resource
→ management.\n2. **Threat Analysis:** Identify immediate dangers based on {obs}. Prioritize

                                  → actions that avoid or neutralize these threats, considering the action-observation sequence → towards the current subgoal ('{action_obs_seq}').\n3. **Goal Alignment:** Assess how each → possible action aligns with your current subgoal ({subgoal}) and overarching survival plan
2433
2434
                                         ('{survival_plan}'). Use '{entries_successful_goal}' and '{entries_failed_goal}' to learn
                                  → from past attempts to achieve similar subgoals.\n4. **Predictive Risk Mitigation:**
→ Evaluate the potential for each action to lead to a critical failure within the next few
2435
2436

→ steps. Prioritize actions that maintain options and avoid irreversible negative consequences
→ based on your history ('{history}'). The 'success_condition' should also be

→ considered.\n\nSelect the single most optimal action from {game_info} that balances resource

2437
                                  \hookrightarrow acquisition/conservation, threat mitigation, goal progression, and predictive risk
2438
                                  \hookrightarrow mitigation. Justify your selection briefly based on the above analysis.\n\nOutput format: \hookrightarrow ACTION: [selected action] | RATIONALE: [brief justification]\n\nYou must provide an output
2439
                                  \hookrightarrow in this format at each step until the episode terminates. Do not output any other text. If \hookrightarrow no immediately safe or advantageous action is available, select the least detrimental action
2440
                                  → while adjusting your 'survival_plan' accordingly."
2441
2442
                            "fitness": 22.5,
"id": "d1914812-4881-4b8e-85f6-ee47ccce9f47",
2443
                              _std_error": 6.602556323122129
2444
2445
2446
```

```
2447
                NetHack Final Genome
2448
2449
2450
                  "hierarchy": 1,
                  "amygdala": 0,
2451
                  "loop evaluator": 1,
                  "explorer": 0,
2452
                  "summariser": 1,
2453
                  "time_decay": [
2454
                    0.004870010771374662
2455
                   cosine_cutoff": [
2456
                    0.01
2457
                   epsilon": [
2458
                    0.01
2459
2460
                   .
'temperature": 1.0,
                  "prompts": {
2461
                     "high_level": [
2462
                      "You are a strategic planner in a hierarchical video game AI system. Your role is to analyze
2463
                      \hookrightarrow the current game state and generate achievable subgoals that strategically advance toward

→ the main objective.\n\n

                                                                                   CRITICAL CONSTRAINTS:\n
2464
                                    - Subgoals must be immediately achievable given the current agent state and

→ capabilities\n

                                                                        - Focus on the next logical progression steps, not
2465
                                                                                - Each subgoal should have clear, observable
                      → distant end-goals\n
                      \hookrightarrow success criteria\n
                                                                             - Subgoals should be novel and interesting,
2466

→ compared to previous attempts\n

    **Subgoals should consider

2467
                      → resource acquisition and conservation. Avoid actions that waste valuable resources unless
                      \hookrightarrow absolutely necessary for survival or progression.**\n\n Bel \hookrightarrow will provide the game description, available actions, and current state information.\n
                                                                                                                       Below I
2468
                                                Game description:\n
                                                                                                       {game_info}\n
2469
                                           - Game State: {obs} \n\n
2470
                                      CURRENT CONTEXT: \n\n
                                               - Survival plan provided by the survival planner, that you should
2471
                      \hookrightarrow consider for your tasks: {survival_plan}\n - Here \hookrightarrow similar successful entries from the archive: {entries_successful_goal} \n\n
                                                                                                         - Here are most
2472
                                     - Here are most similar failed entries from the archive: {entries_failed_goal}

    \n\n\n

                                                               ANALYSIS FRAPAPMEWORK:\n
2473
                      \hookrightarrow Analyse the summary from previous runs and let it guide your decision making. \n
2474
                                          2. Assess what is immediately possible given current agent state and
                      → environment\n
                                                                       3. Identify what kind of actions could be
2475
                                                                                             4. Identify the most direct
5. Select subgoals that form a
                      \begin{tabular}{ll} \hookrightarrow \text{considered novel or interesting} \backslash n \end{tabular}
                      \hookrightarrow path toward the main objective\n
2476
                      → logical sequence\n
                                                                             6. Ensure each subgoal can be verified through
                      → observable game state changes\n\n
                                                                                              SUBGOAL SELECTION CRITERIA:\n
2477
                                                     - Feasible: Can be started immediately with current
2478

→ resources/position\n

                                                                                - Measurable: Success/failure can be
                      \hookrightarrow determined from game observations\n
                                                                                                 - Progressive: Each subgoal
2479
                      → enables the next or advances toward main goal\n
                                                                                                              - Specific: Clear
                         enough for a lower-level agent to understand and execute\n
2480
                      \hookrightarrow Considerate of the summary of previos runs\n\n \hookrightarrow your subgoals are sequential."
                                                                                                             Make sure that
2481
2482
                     "low level": [
2483
```

```
2484
                          "You are an important executor component of a hierarchical video game system. You are given
2485
                         \hookrightarrow one of higher level option and its termination condition proposed by the higher level \hookrightarrow planner. Your role is to propose a sequence of actions that will make you progress towards
2486
                                                                                             Below I will provide you with
                              the given option.\"

→ description, possible actions you can take and the overall goal of the game.\n\n

                                                 2488
                                  Here is a subgoal provided by the high level planner that you should focus on completing:
2489
                         → {subgoal} \n\n
                                                                                    Here is your current state: {obs} \n\
                                                 Here is the action-observation sequence towards current subgoal:
2490
                         \hookrightarrow {action_sequence} \n
                                                                                               Here are the most similar successful
                         → entries from the archive: {entries_successful_goal} \n\n
                                                                                                                                            Here are
2491
                          \hookrightarrow the most similar failed entries from the archive: {entries_failed_goal} \n\n
                         \hookrightarrow Use the action and observation sequence together with the current state to decide \hookrightarrow the **full ordered sequence of actions** that will achieve the subgoal. \n
2492
2493
                                            Avoid repeating the same actions if the observation doesn't change. \n"
2494
                         amygdala": [
2495
                          "\n
                                        Decide if survival mode should activate. 
 \ensuremath{\texttt{N}}
                                                                                                            Observation: {obs}\n
                         → Survival plan: {survival_plan}\n\n
2496
                                                                                     1. Check if observation meets any subtask
                         → prerequisites.\n
                                                            2. If several match, pick highest priority.\n"
2497
                         loop": |
2498
                          "You are a Senior Game AI Debugging Specialist, tasked with analyzing the behavior of an agent
2499
                         → in a hierarchical video game system. Your primary goal is to determine if the agent's 
→ current behavior constitutes a genuine \"stuck\" state, which requires

→ intervention.\n\nConsider the following details about the agent's progress:\n\n* **
2501
                         2503
                         \hookrightarrow the agent's actions, considering the following factors specific to Game AI:\n\n1
                         \hookrightarrow **Information Gain vs. Redundancy:** Is the agent genuinely gathering *new* and *relevant* \hookrightarrow information? Merely observing a change in the environment after an action doesn't
                         \hookrightarrow necessarily indicate progress. Determine if the information gained helps reduce uncertainty
2505
                         → related to achieving the current subgoal or aligns with the overall game objective, given by
→ 'survival_plan'.\n2. **Exploration Strategy:** Is the agent's exploration strategy
                         \hookrightarrow \texttt{sufficient to overcome potential local minima or deceptive landscapes? Many \verb|\"stuck|| }
                          → situations arise from poor exploration. Consider if the agent's 'action_sequence' explores
                          \hookrightarrow diverse enough actions, or if it is repeating the same action variations in a small
                         \hookrightarrow area.\n3. **World Model Limitations:** Assess whether the agent's internal model of the
                          → game world (implied by its actions and reactions in 'action_obs_seq' and 'history') is
2509
                             incomplete or inaccurate. Is it missing key information about the game mechanics, object
                         \hookrightarrow interactions, or environmental dynamics? If so, suggest avenues for improving the world
2510
                          → model, such as targeted exploration or observation of specific events.\n4. **Leverage
2511
                             entries_successful_goal and entries_failed_goal:** Look into similar subgoals in the past to
                         → help you understand how the agent behaved then and compare that with the current
2512
                          → behavior.\n\n**Your Task:**\n\nBased on the provided details and your expertise in Game AI,
                          \hookrightarrow determine if the agent is genuinely stuck, meaning it's unlikely to achieve its subgoal
2513
                         \hookrightarrow without external intervention. Focus on *why* the agent is stuck. Specifically, is the \hookrightarrow agent's failure due to:\n\n* A lack of crucial information that could be obtained through \hookrightarrow making informed decisions?\n* An inaccurate or incomplete world model preventing it from A fundamental flaw in its action selection
2514
2515
                              strategy?\n\nProvide a concise justification for your conclusion, outlining the specific
                         → factors that support your assessment. Prioritize identifying concrete steps the agent could
→ take to overcome the \"stuck\" state, considering the limited information it may possess.
2517
                             Avoid vague statements and focus on actionable recommendations rooted in Game AI best
2518
                         → practices.
2519
                        termination": [
                          "You are an important termination evaluator component of a hierarchical video game system,
2521
                          → functioning as a specialized AI reinforcement learning analyst.\n\nYour task is to:\n\n1.
                          → Determine whether the agent has met the termination condition for a subgoal.\n\n2. Provide a
2522
                          → concise summary that will help guide the lower-level agent's future actions, specifically
                         \hookrightarrow addressing potential issues related to reinforcement learning
                          → strategies.\n\nDetails:\n\nHere is your current state that you should compare with the
                          \hookrightarrow termination condition: {obs}\n\nHere is the subgoal the lower-level agent is working
                         → towards: {subgoal}\n\nHere is the termination condition of the above subgoal given by the
                          → higher-level agent: {success_condition}\n\nInstructions:\n\nAnalyze the subgoal and its
                          \hookrightarrow termination condition and decide if the subgoal is completed. Then, using the
2526
                         → action-observation sequence: {action_obs_seq}, provide a high-level summary of your current
→ evaluation.\n\nIn your summary, consider the following reinforcement learning concepts:\n\n*
2527
                         → **Exploration vs. Exploitation:** Is the agent adequately exploring the environment, or is 

→ it prematurely exploiting a sub-optimal strategy? If the agent has been repeating the same
2528
                         → actions in similar states (as observed in {history}) without success, encourage more

→ exploration.\n\n* **Reward Shaping:** Consider how the higher-level agent's reward

→ structure (implicit or explicit) might be influencing the agent's behavior. Is the agent

→ being incentivized towards unintended consequences?\n\n* **Local Optima:** Is the agent
2529
2530
                         → potentially stuck in a local optimum? Compare the current situation with
→ {entries_failed_goal} to identify if similar action sequences have previously led to
2531
2532
                                         If so, suggest a drastic change in strategy.\n\
                                                                                                             **Credit Assignment:** How can

→ the agent better understand which actions are contributing to success or failure? Suggest
→ focusing on the most recent actions, especially in light of {observation_sequence}.\n\nIf

                         \hookrightarrow the agent has failed similar subgoals in the past, provide an alternative approach based on
                         \hookrightarrow {entries_successful_goal}. Highlight alternative actions it can take in a similar \hookrightarrow situation.\n\nRemember that your summary will be passed to the low-level component to
2535
                         \hookrightarrow improve its actions. It should be actionable and specific. It should be in simple language
                          → for the low-level agent to understand and must contain suggestions to avoid common RL
```

```
2538
                         \hookrightarrow pitfalls based on the information provided. The survival plan is {survival_plan} for high
2539
                         → level goals."
2540
                        "summariser": [
2541
                          "You are a critic module analyzing an agent's attempt to achieve a subgoal in a game
2542
                         \hookrightarrow environment, considering the agent's survival plan.\n\nyour task is to identify the **single
                         → most important factor** that caused SUCCESS or FAILURE, given the broader context of
→ survival.\n\nInformation:\n- Target subgoal: {subgoal}\n- Outcome of the action-observation
2543
2544
                         2545
2546
                         \hookrightarrow threats and resource priorities.\n- Then, analyze the action-observation sequence in the
                         \hookrightarrow context of the target subgoal and survival plan.\n- Focus on **specific resources and their \hookrightarrow quantities** that were critical according to the survival plan. Consider if failing to meet
2547
                         \hookrightarrow (success_condition) resulted from a resource shortfall, specifically referencing the \hookrightarrow quantities mentioned in the action-observation history.\n- How did the agent's actions
2548
                         \hookrightarrow \text{ either help or hinder the broader survival strategy defined in } \{survival\_plan\}?\\ \\ \text{n- If}
2549
                         \hookrightarrow resources were missing that were crucial to survival, state **exactly which resources and \hookrightarrow how many** were missing and how it violated the survival plan.\n- If no resources are the
2550
                         \hookrightarrow primary issue, state the next most relevant factor that impacted both the subgoal and \hookrightarrow survival chance.\n- Express the result in **one short sentence** highlighting the connection
2551
                         \hookrightarrow to {survival_plan}."
2552
                        explorer": [
2553
                          0,
"\n
2554
                                       Task: Create an exploration plan to help the agent discover new skills.\n\
                                                 - Game info: {game_info}\n - Observation: {obs}\n - Subary}\n - Recent 16 action\u2013observation pairs: {history}\n\n
                         → Data:\n
                                                                                                                                      - Subgoal
2555
                         \hookrightarrow summary: {summary}\n
                                                  1. Analyze the environment and agent\u2019s situation.\n
                         → Steps:\n
2556
                         \hookrightarrow focused exploration plan with clear purpose, conditions, and indicators.\n\n
2557
                        "sequential": [
2558
                          "You always have to output one of the above actions at a time and no other text. You always
2559
                         → have to output an action until the episode terminates."
2560
                     2561
                     "id": "aaac4dfb-4a3c-4090-a8bc-5f9265d65eda",
                     "_std_error": 0.4739428639198163
2563
2564
2565
```