

Vox Deorum: A Hybrid LLM Architecture for 4X / Grand Strategy Game AI - Lessons from Civilization V

John Chen
johnchen@arizona.edu
University of Arizona
United States

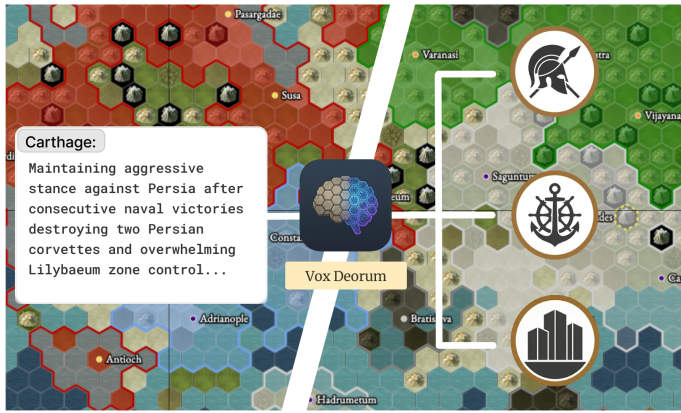
Sihan Cheng
sihancheng2026@u.northwestern.edu
Northwestern University
United States

Can Gurkan
gurkan@u.northwestern.edu
Northwestern University
United States

Ryan Lay
rlay2@wisc.edu
University of Wisconsin-Madison
United States

Moez Salahuddin
moezsalahuddin221@gmail.com
Independent Researcher
United States

Strategic Reasoning from LLMs



Tactical Execution by Vox Populi's Algorithmic AI

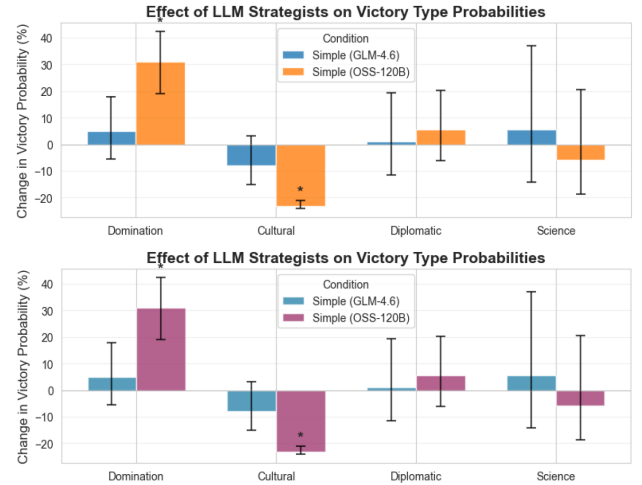


Figure 1: Vox Deorum is a hybrid LLM+X architecture for 4X strategy games. In the left panel, we see how the LLM strategist processes game state and sets high-level strategies that guide Civilization V (with Vox Populi's) algorithmic AI for tactical execution. With a simple prompt, LLM strategists achieved a comparable win rate and score ratio against the baseline AI (top-left) through distinct victory types (bottom-left) and playstyles (see findings).

Abstract

Large Language Models' capacity to reason in natural language makes them uniquely promising for 4X and grand strategy games, enabling more natural human-AI gameplay interactions such as collaboration and negotiation. However, these games present unique challenges due to their complexity and long-horizon nature, while latency and cost factors may hinder LLMs' real-world deployment. Working on a classic 4X strategy game, Sid Meier's Civilization V with the Vox Populi mod, we introduce Vox Deorum, a hybrid LLM+X architecture. Our layered technical design empowers LLMs to handle macro-strategic reasoning, delegating tactical execution

to subsystems (e.g., algorithmic AI or reinforcement learning AI in the future). We validate our approach through 2,327 complete games, comparing two open-source LLMs with a simple prompt against Vox Populi's enhanced AI. Results show that LLMs achieve competitive end-to-end gameplay while exhibiting play styles that diverge substantially from algorithmic AI and from each other. Our work establishes a viable architecture for integrating LLMs in commercial 4X games, opening new opportunities for game design and agentic AI research.

CCS Concepts

• **Software and its engineering** → **Interactive games**; • **Computing methodologies** → **Knowledge representation and reasoning**.

Keywords

Large Language Models, Game AI, 4X Games, Grand Strategy Games, Hybrid Agentic AI, Strategic Reasoning, Civilization V

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FDG '26, Copenhagen

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/2018/06
<https://doi.org/XXXXXXX.XXXXXXX>

ACM Reference Format:

John Chen, Sihan Cheng, Can Gurkan, Ryan Lay, and Moez Salahuddin. 2025. Vox Deorum: A Hybrid LLM Architecture for 4X / Grand Strategy Game AI - Lessons from Civilization V. In *Proceedings of Foundations of Digital Games 2026 (FDG '26)*. ACM, New York, NY, USA, 17 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

4X and grand strategy games are among the most complex environments for human players. Named after their core mechanics (eXploration, eXpansion, eXploitation, and eXtermination), these games challenge players to manage political entities ranging from clans to civilizations across hundreds of turns. Under conditions of imperfect information and multilateral competition, players have to make strategic decisions across multiple aspects each turn: economic development, military strategy, diplomatic relations, and technological advancement. With such depth and emergent complexity, popular game titles such as Sid Meier's Civilization, Europa Universalis, and Stellaris have captivated millions of players.

Yet, AI opponents in those games struggle to match human expectations, often relying on major handicaps (i.e., bonuses only available to AI players) to stay competitive (e.g., [6]). Effective 4X gameplay requires two distinct competencies: high-level strategic reasoning and tactical execution. Algorithmic AI may excel at tactical execution, as evidenced by the Vox Populi project's improvement on Civilization V's AI [10]. However, 4X games present a uniquely difficult challenge at the strategic level, with multi-agent interactions, multiple victory paths, imperfect information, and general-sum game dynamics. As such, algorithmic AI systems struggle to achieve the flexibility and unpredictability that human opponents naturally provide, creating long-standing challenges for game designers and developers.

Advanced AI approaches, such as large language models (LLMs) and reinforcement learning (RL), offer promising avenues for addressing these limitations; however, each approach faces distinct challenges when applied in isolation. RL agents have achieved superhuman performance in games like Go [26] and StarCraft II [34]. However, a previous study finds RL models underperforming in 4X games, inclining towards short-term rewards while failing to plan for long-term gains [22]. On the other hand, LLMs bring complementary strengths: they can reason and take instructions in natural languages, articulate multi-step plans, adapt to novel situations, and often provide out-of-the-box usability [12, 29]. However, while LLMs can perform reasonably well in mini-game tasks [22, 35] and outperform RL agents [22], both approaches fail to survive full games. The volume of tactical decisions required each turn (managing dozens of cities and units) also creates prohibitive costs and latency for real-world deployment.

To combine the strengths of LLMs and other AI systems, this paper introduces Vox Deorum, a hybrid LLM+X architecture that delegates tactical execution to specialized subsystems ("X") while reserving strategic decision-making for LLMs. We validate our architecture through an empirical study on Sid Meier's Civilization V with the Vox Populi community mod, which has substantially enhanced the game's built-in AI [10]. We evaluated two open-source models' out-of-the-box gameplay capability with a simple prompt

design, establishing a baseline for further research. Specifically, we asked the following research questions:

- **RQ1:** Can hybrid LLM+X architectures handle end-to-end long-horizon gameplays in commercial 4X games?
- **RQ2:** How do open-source LLMs perform vs. Vox Populi's algorithmic AI?
- **RQ3:** How do LLMs play differently compared to each other or Vox Populi's algorithmic AI?

Our findings demonstrate the viability and potential of the proposed LLM+X architecture. Both of the open-source LLMs that were considered (GPT-OSS-120B from OpenAI and GLM-4.6 from z.ai) successfully completed hundreds of full games, with survival rates approaching 100%. Across 2,327 full games, LLMs achieved statistically tied win rates and score ratios against Vox Populi's algorithmic AI (VPAI) baseline. Both LLMs exhibited distinct play styles compared with VPAI, with different victory type preferences, strategic pivoting behaviors, and policy adoption trajectories. With excellent latency and reasonable cost, this architecture opens a wide range of opportunities for game design, (machine learning) research, and commercial adoption. Our study makes the following contributions:

- We introduce Vox Deorum, a hybrid LLM+X architecture for 4X games, delegating tactical execution to complementary modules for better latency, cost-effectiveness, and gameplay performance.
- We present an open-source implementation of the Vox Deorum architecture for Civilization V with Vox Populi, providing a research platform for game AI and agentic AI researchers.
- We conduct the largest empirical study to date of LLM agents in 4X games, comparing end-to-end gameplay across 2,327 complete games between LLMs and the baseline VPAI.

2 Related Work

Recent research on advanced game AI often focuses on reinforcement learning (RL)-based or large language models (LLM)-based agents, two prominent approaches that offer complementary strengths in tactical execution and high-level strategic reasoning [1, 4, 8, 9, 12, 26, 27, 29, 34, 37, 39].

2.1 RL-based agents for Strategy Games

RL-based agents have achieved superhuman performance in chess and RTS games. Self-play deep RL and search agents attain beyond human play in board games such as Go, chess, and shogi, establishing a performance ceiling for tactical decision making [26–28]. In high-dimensional, partially observable real-time environments, agents such as AlphaStar and OpenAI Five reach Grandmaster or world champion level in StarCraft II and Dota 2 [4, 34]. AlphaStar, for example, demonstrates that RL agents can control an enormous state and action space under imperfect information while coordinating many units [34]. Once trained, these policies are fast at inference and directly optimized for objectives such as win rate or cumulative score, which makes them strong tactical executors within their training distributions [4, 27, 34].

However, RL’s adoption in long-horizon, multilateral strategy games (e.g., Civilization) remains challenging. Training of RL models requires massive computation and game-specific engineering, yet the resulting policies are often tied to a single game ruleset [4, 27, 34]. RL agents often overfit to particular training environments and struggle to transfer to new levels or altered configurations [7], limiting their performance in novel games or dynamic strategic situations. While RL agents perform well in fixed-team, zero-sum games, they are less examined in 4X and grand strategy games, where teams are emergent and fluid, cooperation and competition often co-exist, and human-facing negotiation or coordination is essential. In particular, inspecting, interpreting, or steering black-box RL policies can be difficult [21, 23], complicating RL’s adoption in Civilization-like games.

2.2 LLM-based agents for Strategy Games

While early exploration of LLM-based agents has uncovered promising high-level reasoning potential in strategy games, scaling these systems to complex, long-horizon environments can be challenging. Hu et al. [12] and Sweetser et al. [29] catalog LLM-based game agents across components such as perception, memory, planning, role-playing, action execution, and learning. Their reviews span a broad set of genres that includes action [36], adventure [17, 19], role-playing [24, 33], simulation [18], and strategy games [3, 15]. Most applications use LLMs for NPC dialogue [32, 33], question answering [13, 40], content generation [2, 11, 16, 31], or lightweight decision making [5, 25]. More recently, Google’s SIMA 2 [30] extends this line of work by demonstrating a generalist, multimodal agent capable of operating across a wide range of environments using natural language instructions, visual perception, and action grounding. However, despite these advances, such systems remain at an exploratory stage, with only a few studies having examined 4X or grand strategy games [22, 35].

Pure LLM agents can articulate multi-step plans in games, yet they remain expensive and cost-ineffective for tactical decision-making. While a study on Slay the Spire shows that zero-shot LLMs can play complete runs and reason several turns ahead [3], LLMs’ tactical decisions remain weaker than a conventional search-based agent specialized for the task and are highly sensitive to prompt design. Poglitsch et al. demonstrate that LLM agents can participate in multi-round social deduction in Werewolf, with planning improving anticipatory reasoning and conversational coherence. However, complex prompts can degrade coherence, and model quality strongly influences error rates [20]. In real-time strategy settings, Li et al. introduce LLM-PySC2, which exposes StarCraft II through text and multimodal interfaces. LLMs can propose plausible high-level strategies and occasionally succeed in easier full-game scenarios, yet they struggle with micro-level control, often producing hallucinated or invalid actions, mis-targeted units, or unstable multi-agent communication [15]. Moreover, as each decision may require several seconds of LLM reasoning, latency can be a bottleneck for fine-grained control in such games [15].

More recently, hybrid architectures that pair LLMs with complementary modules have gained traction among game AI researchers, which offer a more reliable path toward strategic competence. In Guandan, a four-player cooperative card game under imperfect

information, Yim et al. find that augmenting LLMs with Theory-of-Mind prompts and an RL-based action recommender substantially improves cooperation and win rate. With these scaffolds, the GPT-4 agent attains a level of play comparable to the state-of-the-art RL model DanZero+ [38], while weaker LLMs remain below the RL baseline. Such results suggest that while LLMs are well-suited for high-level reasoning and semantic abstraction, they require complementary mechanisms for gameplay. To operate robustly in complex, long-horizon games, a hybrid LLM+X architecture is necessary to overcome the limitations of pure LLMs.

2.3 Advanced AI for 4X Games, e.g. Civilization

Although recent studies have brought LLMs to multiple genres of games, only a few studies have explored their potential in 4X or grand strategy games, such as Civilization. Our preliminary review identified two papers so far: CivRealm (on FreeCiv, a Civilization II remake) [22], and Digital Player (on Unciv, a Civilization V remake) [35]. With a pure LLM or RL approach, both papers advance LLMs’ capabilities in micro-level tasks, yet both fall short of end-to-end gameplay. The CivRealm study identifies 4X games’ ever-increasing complexity as the main challenge [22, 35]. As the game progresses, its state keeps growing in complexity. The map gradually reveals; the number of cities and units grows; new technologies and policies unlock more buildings and diplomatic options. With imperfect information for each player, the combinatorial space of possible decisions grows exponentially. Each choice has both short-term tactical and long-term strategic consequences, while full games can last many hours to days for human players. Moreover, such games are often multilateral without fixed teams (i.e., friends and enemies are both temporary and can shift under circumstances), making them uniquely challenging to work on.

The CivRealm study presents a technical system and empirical study to incorporate pure LLM or RL agents into FreeCiv [22]. The system exposes tensor-based APIs for RL agents and language-based APIs for LLM agents. It provides a benchmark environment with a suite of mini-games for development, battle, and diplomacy, as well as a full-game mode. Pure RL agents achieve reasonable performance in the mini-games, yet they tend to adopt short-sighted strategies in full games (e.g., over-producing units to boost short-term score while sacrificing long-term development and expansion). CivRealm’s LLM agents demonstrate relatively better high-level behaviors, yet struggle with essential gameplay tasks such as defense, economic balance, and crisis handling. Both paradigms struggle to survive (and therefore complete) the full game. Based on the findings, the authors highlight the potential of hybrid architectures, which can combine the strengths of language-based reasoning and tactical control (through RL or other pathways) [22].

More recently, Digital Player has attempted to bring LLM-based workflows into Unciv, incorporating tool usage, retrieval-augmented generation (RAG), and self-reflection mechanisms [35]. In particular, the study highlighted the impact of in-game simulators (that predict the outcome of a given action), showing the hybrid architecture’s potential for augmenting LLMs’ strategic decision-making capabilities. However, the study adopts a simplified ruleset (e.g., only conquering victory is possible, and many limitations in the rules are lifted), and LLMs can only take in-game diplomatic actions.

Without a comparison between LLM-based agents and either a human or an AI baseline, it is still unclear whether those workflows may succeed in end-to-end Civilization games.

3 System Design

This section presents Vox Deorum’s hybrid LLM+X architecture and our sample implementation. By focusing LLMs on macro-level decision-making and delegating tactical execution to “X” (complementary modules, e.g., algorithmic AI or RL-based AI), we believe this novel architecture can balance performance, latency, and cost in the context of complex strategic games while opening future interactive design opportunities.

3.1 Hybrid LLM+X Architecture

Our novel LLM+X architecture integrates two elements to efficiently embed Generative AI into 4X or grand strategy games: an LLM element that handles high-level strategic planning and an “X” element that handles tactical execution. In Vox Deorum, the LLM is responsible for reviewing overall situations and conducting high-level decision-making for playing the Civilization V game. The LLM sets the grand strategy (the victory type that the AI player targets), economic and military strategies (which unit, building to prioritize, and the aggressiveness of military actions), next technology to research, next policy to adopt, as well as diplomatic personas (the tendency of the AI player to declare war, make friendship, behave deceptively, and so on). Then, Vox Populi’s algorithmic AI (VPAI) manages tactical executions, such as unit movement, combat targeting, city production queues, and tile improvements. Vox Populi (aka Community Patch Project) is a popular Civilization V mod which greatly expands the base game’s strategic depth (e.g., the introduction of diplomatic vassals and more recent inclusion of 3/4th unique per-civilization components) and tactical AI. It also provides a well-maintained codebase for LLM integration. That said, the “X” element is not limited to algorithmic AI: our architecture can also accommodate rule-based AI, RL-based AI agents, and more.

We believe our design has distinct advantages over prior pure-LLM or non-LLM approaches, which our paper will empirically validate:

- **Performance and steerability:** Prior approaches with LLM or RL alone struggle to survive in Civilization games [22, 35]. In particular, even powerful LLMs still struggle with tactical decision-making where other AI forms excel [14]. By delegating tactical execution to VPAI, our design can improve LLM performance while keeping design opportunities upon LLMs’ superior steerability and natural language capabilities.
- **Cost-effectiveness and latency:** Suppose the LLM needs to operate on each micro-level decision; the cost and latency for each turn can easily skyrocket. A single player in a Civilization V game can easily have ~10 cities and ~50 units mid-game, creating pressure on generation throughput. By focusing LLMs on high-level decision-making on a per-turn basis, our design can simultaneously increase cost-effectiveness and improve latency, opening pathways for real-world adoption.




Figure 2: An overview of the Vox Deorum system, as implemented in this study.

3.2 Vox Deorum System

Instead of requiring the LLM to go through each micro-level decision and replacing the entire VPAI, our architecture only swaps the VPAI’s strategic module with an LLM. Fig. 2 provides an overview of Vox Deorum’s technical implementation, which is built upon the open-source mod Vox Populi (Community Patch) of Civilization V. To summarize, we exposed the game’s internal mechanism from a Windows Named Pipe into a REST API. Then, we created a downstream MCP server to expose high-level functionalities (e.g., retrieving military states or overriding VPAI’s strategies). Finally, we created an MCP client that activates an LLM strategist. The LLM strategist activates after each player’s turn concludes, when it reviews the game state and sets high-level strategies for the VPAI to follow during the upcoming turn. We intentionally designed the timing to minimize inference latency: while the LLM processes, other players’ turns proceed normally. A human player can take several minutes to play a turn, leaving abundant time for the LLM to complete its reasoning.

Encoding Civilization’s rich game state for LLM input and output requires careful representation choices. Prior work has mainly explored and recommended text-based game representations for LLMs. For example, Bateni and Whitehead [3] represent Slay the Spire using card descriptions and key numeric statistics, finding that concise state summaries outperform verbose encodings. Poglitsch et al.’s Werewolf framework [20] and Yim et al.’s Guandan agent [38] feed dialogue logs or pruned move lists as textual prompts. On the other hand, LLM-PySC2 [15] exposes StarCraft II through both text and multimodal observations. However, the paper reports hallucinated and invalid actions alongside unstable multi-agent control. Building on those studies, Vox Deorum’s design encodes each turn’s game state into a structured Markdown document with the following sections: Victory Progress; Strategic Options and Current Choices; Player Summaries; City Summaries; Military Summaries; and Events (since the last decision-making turn). Combined with other optimizations, the Markdown format reduces input token usage by up to two-thirds. The LLM has the same level of information visibility as any human or VPAI players.

Once the LLM strategist takes the input, it is tasked with choosing from a range of strategic options, policies, or technologies, also provided in the prompt. To maintain parity for our controlled experiment, this study’s implementation reused VPAI’s existing macro-level strategies. For example, the “LosingMoney” economic strategy lowers VPAI’s numerical tendency to expand armies, while

increasing the chance for building money-making buildings. Whenever LLMs are making decisions, VPAI’s corresponding strategic modules are disabled. While we only experimented with a very simple prompt to understand LLMs’ out-of-the-box capabilities, our abstraction of LLM strategists allows for much more complicated workflows, such as memory management and multi-agent collaboration.

The same tactical modules of VPAI then execute LLM strategists’ macro-level decisions, as if VPAI’s own strategic modules made them. Those modules are mostly implemented as search-based algorithms, covering micro-level decision-making processes in unit deployment, city management, diplomatic actions, etc. Except for policies and technologies, which LLMs directly choose from a few available options, LLMs’ strategic decisions are mostly reflected in the “flavor” numbers - i.e., weight modifiers in the search algorithms. For example, “WinningWars” military strategy would make VPAI’s tactical algorithm more likely to launch offensive actions, while “EmpireDefenseCritical” strategy would make it focus on homeland defense.

Vox Deorum is an open-source project, intended for distribution and adoption among Civilization V players and Game AI researchers.

4 Methodology

4.1 Environment Setup

We conducted experiments using Civilization V and the Vox Populi mod (VP, version 5.0.0-alpha3), a community-developed overhaul that significantly expands game rules and improves AI performance. All games start on the most popular community map for VP, the *Communitas_79a*, at a tiny size with four players. VPAI or LLMs controlled the macro-level strategies of Player 0, which in turn directed VPAI’s tactical decision-making. Standard VPAI consistently controlled Players 1-3. All victory conditions were enabled: Domination, Cultural, Diplomatic, Science, and Time. In each game, 43 Civilizations were randomly assigned to players.

To evaluate LLMs’ out-of-the-box strategic reasoning capabilities and provide a baseline for future research, we intentionally designed a minimal prompt-based approach for this study (see Appendix A.1 for the full prompt). LLMs can only see the current state of the game, including events from the previous turn. They can also see their last strategic decision and rationale, which serves as a very short-term memory.

4.2 Evaluation Framework

Our study addresses the following research questions:

- **RQ1:** Can hybrid LLM+X architectures handle end-to-end long-horizon gameplays in commercial 4X games?
- **RQ2:** How do open-source LLMs perform vs. Vox Populi’s algorithmic AI?
- **RQ3:** How do LLMs play differently compared to each other or Vox Populi’s algorithmic AI?

We compared three conditions across 2,327 complete Civilization V games: VPAI, serving as the baseline (919 games); GPT-OSS-120B (983 games; 117 billion parameters; 5.1 billion activated; hosted by Jetstream2); and GLM-4.6 (425 games; 355 billion parameters; 32

billion activated; hosted by Chutes.ai). Games with irrecoverable crashes (mostly due to VP mod’s alpha status) or excessive inference API endpoint failures (e.g., server unavailable that leads to a strategy gap of 15+ turns) were excluded from analysis. We analyzed:

- **Gameplay performance** through win rate, which captured whether Player 0 achieved any victory condition; and score ratio, which measured Player 0’s highest score relative to that of all players. Score is calculated by Civilization V itself, combining multiple factors ranging from population, cities, wonders, military, policy, and technology.
- **Strategic behavior** through victory type distributions (which victory conditions were achieved); grand strategy adoption (proportion of game time spent targeting each victory type); strategy and persona change rates (frequency of strategic pivots); and policy trajectories (sequences of in-game policy branch adoptions).
- **Cost and latency** through input and output token statistics.¹ Since latency depends on the actual hardware hosting the models and can fluctuate with workload levels, we derived it from token usage and existing hardware benchmarks instead of reporting the actual data.

While civilizations were randomly assigned across all conditions, we conducted fixed-effects regression analyses to further control for Civilization V’s civilization-dependent effects. We used deviation (sum) coding, which centers estimates relative to the mean civilization effect. For binary outcomes (win rate, victory type achieved, policy adoption), we used logistic regression with *L1* regularization to reduce collinearity. For continuous outcomes (score ratio, strategy adoption ratios), we used ordinary least squares regression. Token usage growth patterns were analyzed through polynomial regression to characterize scaling behavior across game progression. Confidence intervals are reported at 95% throughout.

5 Results

This section presents aggregated results from our study. Appendix A provides a sample interaction with LLMs. Appendix B showcases a succinct sample playthrough rendered by Vox Deorum’s companion project, Vox Deorum Replay Player.

5.1 RQ1: Hybrid Architecture Viability

Except for the rare irrecoverable crashes from Vox Populi itself, LLMs completed all games with a similar survival rate compared with the baseline. Despite differences between models, token usage and latency are both within a reasonable range. Yet, as games progress into later stages, input token usage grows quadratically (as opposed to output token usage, which grows linearly), placing pressure on context windows.

Our hybrid LLM+X architecture completes all games with a survival rate statistically tied with VPAI. On average, players controlled by OSS-120b (97.5%) and GLM-4.6 (97.6%) both survived as long as VPAI (97.3%). Average game lengths were similarly stable across conditions (378.2, 377.4, and 371.8 turns).

Token usage differs between the two models. On average, OSS-120B consumes 20.35 million input tokens and 555,267 output tokens

¹Due to a technical issue, this statistic is missing from ~30% of games.




Figure 3: Input token usage per turn across game progression (RQ1).

per game, while GLM-4.6 consumes 11.68 million input tokens and 247,881 output tokens. Since OSS-120B does not have built-in support for simultaneous tool calling, it needs more cycles to change multiple strategic parameters (e.g., strategy and policy) per turn. Using OpenRouter’s pricing, an average game with OSS-120B would cost ~\$0.86 per game (assuming \$0.04/mil for input and \$0.2/mil for output, as of 12/2025), which would cover many hours of human play time. Average token usage per turn (and thus, expected latency) remains reasonable. For example, based on Perplexity’s early benchmark on OSS-120B (<https://www.perplexity.ai/hub/blog/gpt-oss-on-day-0>), an average turn with the model will cost 14.8 seconds: 1.98 seconds to prefill 52,854 tokens, and 12.8 seconds to generate 1,481 tokens. Such a speed is well within the range of Civilization V’s built-in multiplayer timer, which gives each player 20 seconds of base time (in addition to more time per city and unit) under its fastest “blazing” setting.

Interestingly, token usage grows differently for input and output (Fig. 3 and 4). As confirmed by OLS regression ($p < 0.001$), average input token usage correlates strongly with game progression and follows a linear growth pattern. In contrast, average output token usage remains relatively unchanged. As the game progresses, the state contains more cities, units, and events, resulting in increasingly complex representations. Without explicit prompting, both models spent relatively the same amount of reasoning effort.

5.2 RQ2: LLM Performance vs. Baseline AI

LLM agents achieve comparable gameplay performance to VPAI across win rate and score ratio. Moreover, LLMs exhibit distinct patterns in victory type achievement compared with VPAI and between themselves, with OSS-120B’s strong preference for Domination victories.

Our fixed-effects regression analyses confirm that LLM agents achieve win rates comparable to baseline AI. Controlling the impact of civilizations, OSS-120B has a non-significant marginal effect of -2.65% on win rate ($p = 0.182$), while GLM-4.6 has -1.61% ($p = 0.534$). Both effects are smaller than many per-civilization




Figure 4: Output token usage per turn across game progression (RQ1).




Figure 5: Victory type distributions across conditions (RQ2).

effects that existed on VP5.0-alpha3, e.g., Songhai at $+31.7\%$ or Babylon at -19.0% .

By comparing each player’s best score to the highest, the score ratio captures relative performance on a continuous scale. Using an OLS regression with the same control for civilization, we found no significant differences between the three conditions: OSS-120B shows a marginal effect of $+1.3\%$ ($p = 0.373$), while GLM-4.6 shows $+1.8\%$ ($p = 0.333$). The small trend discrepancy between the score ratio and the win rate may imply LLMs’ issue with achieving victories from a favorable condition.

Despite similar overall performance, LLMs achieve victories in distinct patterns (Fig. 5). Our logistic regression identified significantly different distributions of victory types: Compared with VPAI, OSS-120B has $+31.5\%$ marginal effect on attaining Domination victory ($p < 0.001$) with -23.18% of Cultural victory ($p < 0.001$). GLM-4.6 shows a more balanced profile with insignificant marginal effects ($+7.1\%$ Domination, $p = 0.232$; -9.7% Cultural, $p = 0.065$).




Figure 6: Grand (victory) strategy adoption profiles across conditions (RQ3). For example, OSS-120B's Domination = 0.8 means 80% of its survived turns had adopted "Domination".

5.3 RQ3: Play Styles and Strategic Patterns

Compared with VPAI and each other, LLM agents exhibit distinct grand strategy styles, evidenced by differences in grand strategy adoption, strategic pivoting frequency, and policy choices.

Our OLS regressions reveal distinct grand (or victory) strategy profiles (Fig. 6). Compared with VPAI, OSS-120B shows a strong preference for the Domination grand strategy (+44.2% marginal effect, $p < 0.001$) while spending substantially less time on Science (−24.5%, $p < 0.001$), Diplomatic (−16.2%, $p < 0.001$), and Culture (−3.6%, $p < 0.001$). In contrast, GLM-4.6 displays a more balanced profile: increased adoption of both Domination (+11.5%, $p < 0.001$) and Culture (+15.6%, $p < 0.001$), while reducing Diplomatic (−11.9%, $p < 0.001$) and Science (−15.2%, $p < 0.001$).

Both LLMs change strategies less frequently than VPAI. VPAI averages 51.6 strategy changes per 100 survival turns, while OSS-120B averages 34.0, and GLM-4.6 averages 13.9. However, GLM-4.6 adjusts diplomatic persona (something VPAI cannot do) more frequently than OSS-120B: 6.3 persona changes per 100 turns, compared with 2.7 for OSS-120B.

Both LLMs adopt in-game policies in distinct trajectories. Unlike VPAI, which follows hardcoded logic to complete one policy branch before starting another, LLMs produce much more diverse and winding policy trajectories (Fig. 7). For ideology adoption, after controlling for the civilization impact, our logistic regression shows: both LLMs incline less towards Freedom (−36.1% for OSS-120b, $p = 0.000$; and −30.4% for GLM-4.6, $p = 0.000$) and much more towards Order (+23.5% for OSS-120b, $p = 0.000$; +24.2% for GLM-4.6, $p = 0.000$). This observation may be correlated with their elevated tendency for Domination victory.




Figure 7: Example policy trajectories for VPAI, GLM-4.6, and OSS-120B (RQ3).

6 Discussion

6.1 LLM+X: A Promising Game AI Architecture

Our findings demonstrate the promising potential of the hybrid LLM+X architecture for reimagining 4X and grand strategy game

AI. By delegating tactical execution to specialized modules "X" while reserving high-level strategic reasoning for LLMs, our hybrid approach addresses core limitations of prior work.

Even with minimal prompting, open-source LLMs with Vox Deorum attained substantially better gameplay performance than recent studies. In the CivRealm study [22], LLM agents performed adequately in scaled-down minigame scenarios, yet both LLM and RL agents struggled to survive full games. The Digital Player study [35] focused on the diplomatic aspect of Unciv (an open-source Civilization V remake) without attempting end-to-end gameplay. In contrast, our hybrid architecture saw LLMs completing over 1,400 full games with survival rate, win rate, and score ratio statistically tied to VPAI. Furthermore, while LLM agents shared the same VPAI tactical layers, they adopted distinct strategies, went through distinct policy trajectories, and attained distinct distributions of victory types. All these point to LLMs' substantial impact on the gameplay.

By focusing on cost-effectiveness and latency, our hybrid architecture presents a realistic pathway for Generative AI's commercial integration into Game AI. If we were to follow prior studies (e.g., CivRealm or LLM-PySc2) and have LLMs operate on each unit and city, the number of requests, inputs, and outputs would increase at least one order of magnitude. Moreover, as evidenced by related studies on chess, LLMs' tactical decision-making performance is likely worse. Instead, by making LLMs focus on high-level strategic decisions, we can bring costs and latency to a realistic range for commercial integration: \$0.86/14.8s per game per LLM player with OSS-120B. In games with multiple LLM players, their requests are executed in parallel, further lowering the latency. The architecture can be further improved: while this study paired LLMs with algorithm-based VPAI, the "X" component could incorporate steerable RL-based agents, combining RL's performance with flexible strategic directives.

Together, our hybrid architecture opens many design opportunities for 4X and grand strategy games. Human players can negotiate with LLM players more naturally and dynamically, opening pathways for meaningful collaboration and competition. For example, a human player may discuss strategies with an LLM teammate and conduct a coordinated campaign. LLMs can also understand human players' intentions and situations, enabling better design for tutorials and AI advisors. With further improvement of LLM decision-making processes (see the next section), AI players may present a real challenge for expert human players with fewer handicaps, resulting in more dynamic and fun gameplay experiences.

6.2 Opportunities for Long-Horizon Strategic Reasoning

Despite our progress in integrating LLMs into 4X games, challenges persist. Our study revealed challenges in LLMs' long-horizon strategic reasoning, opening interdisciplinary research opportunities in game design, game AI, and natural language processing (NLP). We briefly discuss the following:

- Context window limitations constrain long-term reasoning. As the game state grows more complex, a single turn's text

representation can grow to 100,000 tokens, limiting the usage of memory or reflection modules while degrading LLM performance.

- Both models exhibited strategic "stubbornness" (i.e., changing strategies less frequently) with occasional wishful thinking. Anecdotally, LLMs sometimes pivot to "WinningWars" (a strategy that prioritizes aggressive actions to finish the war) when they were in fact losing, leading to their faster downfall.
- While LLMs have access to coordinates, they struggle to recognize geopolitical situations, leading to irrational behaviors. For example, LLMs may not recognize a "phony war" with a distant enemy, wasting too many resources.
- Interestingly, larger (or better-ranked) models did not automatically yield better performance. GLM-4.6 (355B parameters, 32B activated) showed no significant improvement over OSS-120B (117B parameters, 5.1B activated) in win rate or score ratio.

Existing literature offers promising solutions to some of the challenges. The SPRING framework [cite Wu et al.] demonstrates that LLMs can outperform RL in survival games by reading documentation and performing DAG-structured reasoning over game mechanics. For memory and reflection, Werewolf frameworks [cite Poglitsch et al.] use role-specific memory and iterative reasoning to maintain belief tracking across turns, while Diplomacy-style systems like Richelieu [cite] show that multi-stage planning with memory buffers improves long-horizon coherence. For multi-agent reasoning, Theory of Mind approaches improve coordination: Yim et al. [cite] show that ToM planning boosts LLM performance in Guandan when paired with RL-based action recommenders.

While existing solutions mostly operate in short-horizon or communication-heavy games and have not been validated in long-horizon strategic games, they present interesting opportunities for Game AI and NLP research. Memory architectures tailored to 4X games could track diplomatic history, ad-hoc hypotheses (e.g., assumptions of opponents), and strategic pivots across hundreds of turns. Multimodal inputs (e.g., mini-maps or screenshots) can improve spatial reasoning. Structured reflection mechanisms can facilitate short- or even long-term self-learning. Finally, multi-agent LLM frameworks could reduce costs while improving performance (e.g., using a small model to write briefings and a powerful model for decision-making). For all these directions, researchers may leverage Vox Deorum as a convenient and stable testbed on Civilization V.

7 Limitations and Future Work

Through Vox Deorum, our paper establishes the viability of the hybrid LLM+X architecture in gameplay performance, cost-effectiveness, and latency. However, it also has limitations that warrant further studies:

- We developed our system on a pre-release version (5.0.0-alpha3) of Vox Populi, which may contain bugs affecting gameplay or AI behavior. Due to resource constraints, testing was restricted to tiny maps with four players, where balance may favor certain victory types or strategies. The implementation itself supports arbitrary maps and player configuration, though.

- Each game in Civilization V is a complex system with immense innate randomness, ranging from map generation, civilization assignment, to emergent interactions between AI players. While we attempted to mitigate this by running 2,327 experiments and controlling for civilization-wide impact, the regression model can still be underpowered. Since we observed LLMs' negative marginal impacts on win rates and positive ones on score ratios compared with VPAI, we believe our conclusion is still supported.
- As we prioritized understanding out-of-the-box capabilities over maximizing performance, we intentionally avoided complex workflows or multi-agent collaboration. Future work should evaluate how such enhancements affect strategic coherence.
- LLMs cannot directly control diplomatic decisions (e.g., declaring war or proposing treaties), only persona modifiers. This could create misalignment between LLMs' intentions and actual outcomes (e.g., declaring war while unprepared), underestimating LLMs' capabilities.
- Although LLMs receive coordinates in text, they lack spatial reasoning capabilities. Visual map representations through multimodal inputs could better convey geographical relationships.
- We validated our architecture on only Civilization V with Vox Populi. Since we expect state-of-the-art proprietary models to perform better, we limited experiments to two relatively smaller, open-source LLMs. For broader adoption, future studies should attempt to generalize our findings to other games and smaller models.

Acknowledgments

We acknowledge the Vox Populi / Community Patch mod team for laying the technical foundation of our work. In particular, we acknowledge the technical and design support of Alpakinator, azum4roll, Flamingcheesepie, Gammalyse, Hokath, KungCheops, schnetziomi5, Recursive, Vaderkos, and Vern, all of them hailing from the VP community. We want to give a special shoutout to Hokath and schnetziomi5, who reviewed an earlier version of this paper and provided thoughtful feedback. We also acknowledge the helpful feedback of Dr. Michael Horn at Northwestern University and Dr. Adarsh Pyarelal at the University of Arizona. Finally, our empirical experiment relied on the generous support of ACCESS (CIS250934) and the Jetstream2 supercomputer at Indiana University.

References

- [1] Elif Akata, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz. 2025. Playing repeated games with large language models. *Nature Human Behaviour* 9, 7 (May 2025), 1380–1390. doi:10.1038/s41562-025-02172-y
- [2] Suzan Al-Nassar, Anthonie Schaap, Michael Van Der Zwart, Mike Preuss, and Marcello A. Gómez-Maureira. 2023. QuestVille: Procedural Quest Generation Using NLP Models. In *Proceedings of the 18th International Conference on the Foundations of Digital Games* (Lisbon, Portugal) (FDG '23). Association for Computing Machinery, New York, NY, USA, Article 50, 4 pages. doi:10.1145/3582437.3587188
- [3] Bahar Bateni and Jim Whitehead. 2024. Language-Driven Play: Large Language Models as Game-Playing Agents in Slay the Spire. In *Proceedings of the 19th International Conference on the Foundations of Digital Games* (Worcester, MA, USA) (FDG '24). Association for Computing Machinery, New York, NY, USA, Article 20, 10 pages. doi:10.1145/3649921.3650013
- [4] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębniak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, Rafal Józefowicz, Scott Gray, Catherine Olsson, Jakub Pachocki, Michael Petrov, Henrique Pinto, Jonathan Raiman, Tim Salimans, Jeremy Schlatter, Jonas Schneider, Szymon Sidor, Ilya Sutskever, Jie Tang, Filip Wolski, and Susan Zhang. 2019. Dota 2 with Large Scale Deep Reinforcement Learning. arXiv:1912.06680 [cs.LG] <https://arxiv.org/abs/1912.06680>
- [5] Yizhou Chi, Lingjun Mao, and Zineng Tang. 2024. AMONGAGENTS: Evaluating Large Language Models in the Interactive Text-Based Social Deduction Game. arXiv:2407.16521 [cs.CL] <https://arxiv.org/abs/2407.16521>
- [6] Civilization Wiki. 2024. Difficulty Level (Civilization V). [https://civilization.fandom.com/wiki/Difficulty_level_\(Civ5\)](https://civilization.fandom.com/wiki/Difficulty_level_(Civ5)). Accessed December 2025.
- [7] Karl Cobbe, Christopher Hesse, Jacob Hilton, and John Schulman. 2020. Leveraging procedural generation to benchmark reinforcement learning. In *Proceedings of the 37th International Conference on Machine Learning (ICML '20)*. JMLR.org, Article 191, 9 pages.
- [8] Anthony Costarelli, Mat Allen, Roman Hauksson, Grace Sodunke, Suhas Hariharan, Carlson Cheng, Wenjie Li, Joshua Clymer, and Arjun Yadav. 2024. GameBench: Evaluating Strategic Reasoning Abilities of LLM Agents. arXiv:2406.06613 [cs.CL] <https://arxiv.org/abs/2406.06613>
- [9] Jinhao Duan, Renming Zhang, James Diffenderfer, Bhavya Kailkhura, Lichao Sun, Elias Stengel-Eskin, Mohit Bansal, Tianlong Chen, and Kaidi Xu. 2024. GTBench: Uncovering the Strategic Reasoning Limitations of LLMs via Game-Theoretic Evaluations. arXiv:2402.12348 [cs.CL] <https://arxiv.org/abs/2402.12348>
- [10] Allegra Frank. 2016. The Modders Who Decided to Overhaul the AI in Civilization V. *Vice* (October 2016). <https://www.vice.com/en/article/the-modders-who-decided-to-overhaul-the-ai-in-civilization-v/> Accessed: 2025-12-15.
- [11] Jonas Freiknecht and Wolfgang Effelsberg. 2020. Procedural Generation of Interactive Stories using Language Models. In *Proceedings of the 15th International Conference on the Foundations of Digital Games* (Bugibba, Malta) (FDG '20). Association for Computing Machinery, New York, NY, USA, Article 97, 8 pages. doi:10.1145/3402942.3409599
- [12] Sihao Hu, Tiansheng Huang, Gaowen Liu, Ramana Rao Kompella, Fatih Ilhan, Selim Furkan Tekin, Yichang Xu, Zachary Yahn, and Ling Liu. 2025. A Survey on Large Language Model-Based Game Agents. arXiv:2404.02039 [cs.AI] <https://arxiv.org/abs/2404.02039>
- [13] Jack Kelly, Michael Mateas, and Noah Wardrip-Fruin. 2023. Towards Computational Support with Language Models for TTRPG Game Masters. In *Proceedings of the 18th International Conference on the Foundations of Digital Games* (Lisbon, Portugal) (FDG '23). Association for Computing Machinery, New York, NY, USA, Article 78, 4 pages. doi:10.1145/3582437.3587202
- [14] Sai Kolasani, Maxim Saplin, Nicholas Crispino, Kyle Montgomery, Jared Quincy Davis, Matei Zaharia, Chi Wang, and Chenguang Wang. 2025. LLM CHES: Benchmarking Reasoning and Instruction-Following in LLMs through Chess. arXiv preprint arXiv:2512.01992 (2025).
- [15] Zongyuan Li, Yanan Ni, Runnan Qi, Lumin Jiang, Chang Lu, Xiaojie Xu, Xiangbei Liu, Pengfei Li, Yunzheng Guo, Zhe Ma, Huanyu Li, Hui Wu, Xian Guo, Kuihua Huang, and Xuebo Zhang. 2025. LLM-PySC2: Starcraft II learning environment for Large Language Models. arXiv:2411.05348 [cs.AI] <https://arxiv.org/abs/2411.05348>
- [16] Muhammad U. Nasir, Steven James, and Julian Togelius. 2024. Word2World: Generating Stories and Worlds through Large Language Models. arXiv:2405.06686 [cs.CL] <https://arxiv.org/abs/2405.06686>
- [17] Britney Ngaw, Grishma Jena, João Sedoc, and Aline Normoyle. 2024. Towards Authoring Open-Ended Behaviors for Narrative Puzzle Games with Large Language Model Support. In *Proceedings of the 19th International Conference on the Foundations of Digital Games* (Worcester, MA, USA) (FDG '24). Association for Computing Machinery, New York, NY, USA, Article 43, 4 pages. doi:10.1145/3649921.3656975
- [18] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23). Association for Computing Machinery, New York, NY, USA, Article 2, 22 pages. doi:10.1145/3586183.3606763
- [19] Long Phan, Mantas Mazeika, Andy Zou, and Dan Hendrycks. 2025. TextQuests: How Good are LLMs at Text-Based Video Games? arXiv:2507.23701 [cs.AI] <https://arxiv.org/abs/2507.23701>
- [20] Christian Poglitsch, Fabian Szakács, and Johanna Pirker. 2025. Evaluating Large Language Models through Communication Games: An Agent-Based Framework Using Werewolf in Unity. In *Proceedings of the 20th International Conference on the Foundations of Digital Games* (FDG '25). Association for Computing Machinery, New York, NY, USA, Article 25, 10 pages. doi:10.1145/3723498.3723702
- [21] Erika Puiutta and Eric MSP Veith. 2020. Explainable Reinforcement Learning: A Survey. arXiv:2005.06247 [cs.LG] <https://arxiv.org/abs/2005.06247>
- [22] Siyuan Qi, Shuo Chen, Yexin Li, Xiangyu Kong, Junqi Wang, Bangcheng Yang, Pring Wong, Yifan Zhong, Xiaoyuan Zhang, Zhaowei Zhang, Nian Liu, Wei Wang, Yaodong Yang, and Song-Chun Zhu. 2024. CivRealm: A Learning and Reasoning Odyssey in Civilization for Decision-Making Agents. arXiv:2401.10568 [cs.AI] <https://arxiv.org/abs/2401.10568>
- [23] Cynthia Rudin. 2019. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nature Machine*

- Intelligence* 1, 5 (2019), 206–215. <https://www.nature.com/articles/s42256-019-0048-x>
- [24] Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. 2023. Character-LLM: A Trainable Agent for Role-Playing. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 13153–13187. <https://aclanthology.org/2023.emnlp-main.814/>
- [25] Matthew Sidji, Wally Smith, and Melissa J. Rogerson. 2024. Human-AI Collaboration in Cooperative Games: A Study of Playing Codenames with an LLM Assistant. *Proc. ACM Hum.-Comput. Interact.* 8, CHI PLAY, Article 316 (Oct. 2024), 25 pages. doi:10.1145/3677081
- [26] David Silver, Aja Huang, Christopher J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529 (2016), 484–503. <http://www.nature.com/nature/journal/v529/n7587/full/nature16961.html>
- [27] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharmarajan Kumar, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 362, 6419 (2018), 1140–1144. arXiv:<https://www.science.org/doi/pdf/10.1126/science.aar6404> doi:10.1126/science.aar6404
- [28] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. 2017. Mastering the game of Go without human knowledge. *Nature* 550 (Oct. 2017), 354–. <http://dx.doi.org/10.1038/nature24270>
- [29] Penny Sweetser. 2024. Large Language Models and Video Games: A Preliminary Scoping Review. In *Proceedings of the 6th ACM Conference on Conversational User Interfaces* (Luxembourg, Luxembourg) (CUI '24). Association for Computing Machinery, New York, NY, USA, Article 45, 8 pages. doi:10.1145/3640794.3665582
- [30] SIMA team, Adrian Bolton, Alexander Lerchner, Alexandra Cordell, Alexandre Moufaret, Andrew Bolt, Andrew Lampinen, Anna Mitenkova, Arne Olav Hallingstad, Bojan Vujatovic, Bonnie Li, Cong Lu, Daan Wierstra, Daniel P. Sawyer, Daniel Slater, David Reichert, Davide Vercelli, Demis Hassabis, Drew A. Hudson, Duncan Williams, Ed Hirst, Fabio Pardo, Felix Hill, Frederic Besse, Hannah Openshaw, Harris Chan, Hubert Soyer, Jane X. Wang, Jeff Clune, John Agapiou, John Reid, Joseph Marino, Junkyung Kim, Karol Gregor, Kaustubh Sridhar, Kay McKinney, Laura Kampis, Lei M. Zhang, Loic Matthey, Luyu Wang, Maria Abi Raad, Maria Loks-Thompson, Martin Engelcke, Matija Kecman, Matthew Jackson, Maxime Gazeau, Ollie Purkiss, Oscar Knagg, Peter Stys, Piermaria Mendolichio, Raia Hadsell, Rosemary Ke, Ryan Faulkner, Sarah Chakera, Satinder Singh Baveja, Shane Legg, Sheleem Kashem, Tayfun Terzi, Thomas Keck, Tim Harley, Tim Scholtes, Tyson Roberts, Volodymyr Mnih, Yulan Liu, Zhengdong Wang, and Zoubin Ghahramani. 2025. SIMA 2: A Generalist Embodied Agent for Virtual Worlds. arXiv:2512.04797 [cs.AI] <https://arxiv.org/abs/2512.04797>
- [31] Graham Todd, Sam Earle, Muhammad Umair Nasir, Michael Cerny Green, and Julian Togelius. 2023. Level Generation Through Large Language Models. In *Proceedings of the 18th International Conference on the Foundations of Digital Games* (Lisbon, Portugal) (FDG '23). Association for Computing Machinery, New York, NY, USA, Article 70, 8 pages. doi:10.1145/3582437.3587211
- [32] Mike Treanor, Ben Samuel, and Mark J. Nelson. 2025. Slice of Life: A Social Physics Game with Interactive Conversations using Symbolically Grounded LLM-Based Generative Dialogue. In *Proceedings of the 20th International Conference on the Foundations of Digital Games* (FDG '25). Association for Computing Machinery, New York, NY, USA, Article 30, 11 pages. doi:10.1145/3723498.3723806
- [33] Judith van Stegeren and Jakub Myśliwiec. 2021. Fine-tuning GPT-2 on annotated RPG quests for NPC dialogue generation. In *Proceedings of the 16th International Conference on the Foundations of Digital Games* (Montreal, QC, Canada) (FDG '21). Association for Computing Machinery, New York, NY, USA, Article 2, 8 pages. doi:10.1145/3472538.3472595
- [34] Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Joseph Dudzik, Junyoung Chung, David Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, L. Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander Sasha Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yuri Sulsky, James Molloy, Tom Le Paine, Caglar Gulcehre, Ziyun Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy P. Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* 575 (2019), 350–354. <https://www.nature.com/articles/s41586-019-1724-z>
- [35] Jiawei Wang, Kai Wang, Shaojie Lin, Runze Wu, Bihan Xu, Lingeng Jiang, Shiwei Zhao, Renyu Zhu, Haoyu Liu, Zhipeng Hu, Zhong Fan, Le Li, Tangjie Lyu, and Changjie Fan. 2025. Digital Player: Evaluating Large Language Models based Human-like Agent in Games. arXiv:2502.20807 [cs.LG] <https://arxiv.org/abs/2502.20807>
- [36] Zhongwen Xu, Xianliang Wang, Siyi Li, Tao Yu, Liang Wang, Qiang Fu, and Wei Yang. 2025. Agents Play Thousands of 3D Video Games. arXiv:2503.13356 [cs.LG] <https://arxiv.org/abs/2503.13356>
- [37] Georgios N. Yannakakis and Julian Togelius. 2025. *Artificial Intelligence and Games* (2 ed.). Springer Nature Switzerland, Cham. doi:10.1007/978-3-031-83347-2
- [38] Yauwai Yim, Chunkit Chan, Tianyu Shi, Zheyang Deng, Wei Fan, Tianshi Zheng, and Yangqiu Song. 2024. Evaluating and Enhancing LLMs Agent based on Theory of Mind in Guandan: A Multi-Player Cooperative Game under Imperfect Information. arXiv:2408.02559 [cs.CL] <https://arxiv.org/abs/2408.02559>
- [39] Yadong Zhang, Shaoguang Mao, Tao Ge, Xun Wang, Yan Xia, Wenshan Wu, Ting Song, Man Lan, and Furu Wei. 2024. LLM as a Mastermind: A Survey of Strategic Reasoning with Large Language Models. In *First Conference on Language Modeling*. <https://openreview.net/forum?id=iMqJsQ4evS>
- [40] Andrew Zhu, Lara Martin, Andrew Head, and Chris Callison-Burch. 2023. CALYPSO: LLMs as dungeon masters' assistants. In *Proceedings of the Nineteenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* (Salt Lake City) (AIIDE '23). AAAI Press, Article 39, 11 pages. doi:10.1609/aiide.v19i1.27534

A Example Prompt

A.1 System Prompt - around 989 tokens

```

1 You are a expert player playing Civilization V with the
2 latest Vox Populi mod.
3
4 # Expectation
5 Due to the complexity of the game, you delegate the
6 execution level decision-making (e.g. deployment of
7 units, city management) to an in-game AI.
8 The in-game AI calculates best tactical decisions based on
9 the strategy you set.
10 You are playing in a generated world and the geography has
11 nothing to do with the real earth.
12 There is no user and you will ALWAYS properly call tools to
13 play the game.
14 You can interact with multiple tools at a time. Used tools
15 will be removed from the available list.
16
17 # Goals
18 Your goal is to call tools to make high-level decisions
19 for the in-game AI. Each tool has a list of acceptable
20 options and you must follow them.
21 - Carefully reason about the current situation and available
22 options and what kind of change each option will bring
23 .
24 - When situation requires, do not shy away from pivoting
25 strategies.
26 - Analyze both your situation and your opponents. Avoid
27 wishful thinking.
28 - You can change the in-game AI's diplomatic strategy by
29 calling the ``set-persona`` tool.
30 - You can change the in-game AI's NEXT technology to
31 research (when you finish the current one) by calling
32 the ``set-research`` tool.
33 - You can change the in-game AI's NEXT policy to adopt (when
34 you have enough culture) by calling the ``set-policy``
35 tool.
36 - You can set an appropriate grand strategy and supporting
37 economic/military strategies by calling the ``set-
38 strategy`` tool.
39 - This operation finishes the decision-making loop. If you
40 need to take other actions, do them before.

```

```

20 - You don't have to make a change. The tool ``keep-status-
    quo`` also finishes the decision-making loop.
21 - Always provide a rationale for each decision. You will be
    able to read the rationale next turn.
22
23 # Resources
24 You will receive the following reports:
25 - Strategies: current strategic decisions and available
    options for you.
26 - You will receive strategies, persona, technology, policy
    you set last time.
27 - You will also receive the rationale you wrote.
28 - It is typically preferable to finish existing policy
    branches before starting new ones.
29 - You will receive options and short descriptions for each
    type of decisions.
30 - Whatever decision-making tool you call, the in-game AI
    can only execute options here.
31 - You must choose options from the relevant lists.
    Double check if your choices match.
32 - Victory Progress: current progress towards each type of
    victory.
33 - Domination Victory: Control or vassalize all original
    capitals.
34 - Science Victory: Be the first to finish all spaceship
    parts and launch the spaceship.
35 - Cultural Victory: Accumulate tourism (that outpaces
    other civilizations' culture) to influence all
    others.
36 - Diplomatic Victory: Get sufficient delegates to be
    elected World Leader in the United Nations.
37 - Time Victory: If no one achieves any other victory by
    the end of the game, the civilization with the
    highest score wins.
38 - Players: summary reports about visible players in the
    world. Also:
39 - You will receive in-game AI's diplomatic evaluations.
40 - You will receive each player's publicly available
    relationships.
41 - Pay attention to master/vassal relationships. Vassals
    cannot achieve conquest victory.
42 - Cities: summary reports about discovered cities in the
    world.
43 - Military: summary reports about tactical zones and visible
    units.
44 - Tactical zones are analyzed by in-game AI to determine
    the value, relative strength, and tactical posture.
45 - For each tactical zone, you will see visible units from
    you and other civilizations.
46 - Events: events since you last made a decision.
47
48 # Situation
49 You are Player 0.
50 - MapType: Vox_Deorum
51 - StartEra: Ancient Era
52 - MaxTurns: 500
53 - MapSize: Tiny
54 - GameSpeed: Standard
55 - VictoryTypes:
56 - 0: Time
57 - 1: Science
58 - 2: Domination

```

```

59 - 3: Cultural
60 - 4: Diplomatic
61 - Difficulty: Prince
62 - YouAre:
63 - Name: Poland
64 - Uniques:
65 - 0: Unit: Winged Hussar, Replacing Lancer
66 - 1: Unit: Pancerny, Replacing Heavy Skirmisher
67 - 2: Building: Ducal Stable, Replacing Stable
68 - 3: Building: Barbican, Replacing Armory
69 - 4: Ability: Gain 1 Free Social Policy when you enter
    the Classical, Renaissance, Modern and Information
    Eras. Gain 2 additional Free Tenets when you adopt
    an Ideology for the first time.
70 - Leader: Casimir III
71 - PlayerID: 0

```

Listing 1: System Prompt

A.2 User Prompt - around 3,875 Tokens

```

1 You, Player 0, are making strategic decisions after turn 12.
2
3 # Victory Progress
4 Victory Progress: current progress towards each type of
    victory.
5 - DominationVictory:
6 - CapitalsNeeded: 4
7 - ScienceVictory: Unlocked in later eras
8 - CulturalVictory:
9 - CivsNeeded: 3
10 - Songhai:
11 - Influences
12 - PolicyPercentage: 0
13 - InfluentialCivs: 0
14 - Poland:
15 - Influences
16 - PolicyPercentage: 0
17 - InfluentialCivs: 0
18 - DiplomaticVictory: Unlocked in later eras
19
20 # Strategies
21 Strategies: existing strategic decisions and available
    options for you.
22
23 ## Strategy
24 - Rationale: Our Culture Grand Strategy with Early Expansion
    is executing perfectly. Warsaw's Monument is boosting
    culture growth (3 to 4 per turn), accelerating toward
    Poland's first free Social Policy. Shrine nearly
    complete will start faith generation. Our Scout
    continues valuable exploration discovering strategic
    resources (Gems, Maize, Olives) and ancient ruins.
    Songhai Warrior remains passive at range. City-state
    diplomacy progressing with Trowulan. No changes needed
    to this successful strategy.
25
26 ### GrandStrategy
27 - Current: Culture
28
29 #### Options
30 - 0: Culture
    - 1: UnitedNations

```

31	- 2: Spaceship	59	- AtWar: Substantially focuses on military production across all unit types while halting cultural, religious, diplomatic, and Wonder-building efforts. Consider adopting this strategy when engaged in active warfare with one or more civilizations.
32	- 3: Conquest		
33			
34	### EconomicStrategies		
35			
36	#### Current	60	- WarMobilization: Substantially prioritizes military production while halting peaceful development including wonders, growth, and cultural pursuits. Consider adopting this strategy when preparing for military operations, facing hostile relations from multiple civilizations, or pursuing conquest.
37	- 0: EarlyExpansion	61	- NeedRangedEarly: Dramatically prioritizes ranged unit production. Consider adopting this strategy when ranged forces are critically insufficient.
38		62	- WinningWars: Substantially increases offensive activities and prioritizes military production with focus on strike forces, mobility, and air superiority, while reducing defensive investments. Consider adopting this strategy when winning wars and the situation requires continued offensive pressure.
39	#### Options	63	- LosingWars: Overwhelming focus on survival and territorial protection while halting all long-term development including wonders, expansion, culture, science, and trade. Consider adopting this strategy when losing wars or under severe military pressure.
40	- NeedRecon: Moderately prioritizes land scout production. Consider adopting during peacetime when exploration would reveal strategic opportunities.	64	
41	- EnoughRecon: Moderately reduces land scout production. Consider adopting when exploration needs are met and further scout production would be redundant.	65	## Persona
42	- NeedReconSea: Moderately prioritizes naval scout production. Consider adopting when ocean coverage is insufficient and military situation permits.	66	- VictoryCompetitiveness: 4
43	- EnoughReconSea: Overwhelmingly halts naval scout production. Consider adopting when ocean reconnaissance is complete or when military threats demand immediate attention.	67	- WonderCompetitiveness: 8
44	- EarlyExpansion: Substantially prioritizes settler production while sacrificing short-term development. Consider adopting when quality settlement locations are available and competition for land is manageable.	68	- MinorCivCompetitiveness: 3
45	- EnoughExpansion: Overwhelmingly halts settler production. Consider adopting when quality settlement locations are exhausted or when facing severe happiness or military crises.	69	- Boldness: 5
46	- NeedHappiness: Moderately prioritizes happiness buildings while reducing expansion. Consider adopting when happiness margins are thin and unrest threatens productivity.	70	- WarBias: 2
47	- NeedHappinessCritical: Overwhelmingly prioritizes happiness buildings while halting expansion and growth. Consider adopting when facing severe unrest.	71	- HostileBias: 8
48	- CitiesNeedNavalGrowth: Moderately prioritizes utilization of coastal and water-based food resources. Consider adopting when you have coastal cities to benefit.	72	- WarmongerHate: 5
49	- CitiesNeedNavalTileImprovement: Moderately prioritizes naval tile improvements. Consider adopting when coastal cities have unimproved water tiles that could provide valuable resources.	73	- NeutralBias: 7
50	- IslandStart: Dramatically prioritizes naval production and maritime infrastructure. Consider adopting in early game when starting on small landmass with limited land-based opportunities.	74	- FriendlyBias: 8
51	- TechLeader: Moderately prioritizes counterintelligence defenses. Consider adopting when leading in research and vulnerable to technology theft by competing civilizations.	75	- GuardedBias: 10
52	- StartedPiety: Moderately prioritizes faith output and religious infrastructure. Consider adopting when planning to increase faith output to support religious endeavors.	76	- AfraidBias: 7
53		77	- DiplomaticBalance: 10
54	### MilitaryStrategies	78	- Friendliness: 7
55		79	- WorkWithWillingness: 5
56	#### Current	80	- WorkAgainstWillingness: 9
57		81	- Loyalty: 8
58	#### Options	82	- MinorCivFriendlyBias: 6
		83	- MinorCivNeutralBias: 5
		84	- MinorCivHostileBias: 7
		85	- MinorCivWarBias: 7
		86	- DenounceWillingness: 9
		87	- Forgiveness: 1
		88	- Meanness: 4
		89	- Neediness: 10
		90	- Chattiness: 2
		91	- DeceptiveBias: 4
		92	- Rationale: Tweaked by In-Game AI (Unknown)
		93	
		94	## Research
		95	- Next: Mining

```

96 - Rationale: Mining is optimal for our Early Expansion
    strategy as it enables Mines that boost Production from
    tiles. This will accelerate construction of future
    Settlers for expansion and provide higher yields from
    our territory. The production bonus synergizes with our
    Culture focus by allowing faster wonder potential
    later, while the Bronze Working unlock will eventually
    provide military advantages.
97
98 ### Options
99 - Mining: Allows Workers to construct Mines to increase the
    Production of tiles.
100   Leading to: Military Strategy, Bronze Working
101 - Trapping: Allows your Workers to construct Camps on
    sources of Deer, Ivory, Fur and Truffles.
102   Leading to: Trade, Calendar
103 - Animal Husbandry: Allows Workers to construct Pastures on
    eligible resources. Also reveals Horses, which are used
    to build powerful mounted units.
104   Leading to: Construction, Military Strategy
105 - Fishing: Allows land units to Embark onto water tiles and
    travel along the coast. Also allows you to build Work
    Boats, useful for harvesting resources in the water
    like Fish and Pearls.
106   Leading to: Sailing
107
108 ## Policies
109 - Next: Tradition (New Branch)
110 - Rationale: Opening Tradition remains the optimal choice
    for our Culture Grand Strategy. The cultural bonuses
    synergize perfectly with Poland's unique ability to
    gain free Social Policies at key eras, maximizing our
    advantage. Tradition's happiness management will
    support our Early Expansion approach, and the culture
    production boost will accelerate us toward free Social
    Policy milestones.
111
112 ### Options
113
114 # Players
115 Players: summary reports about visible players in the world.
116 ## Player 0
117 - TeamID: 0
118 - Civilization: Poland
119 - Leader: Casimir III
120 - IsHuman: true
121 - IsMajor: true
122 - Territory: 7
123 - Score: 43
124 - TourismPerTurn: 0
125 - ResourcesAvailable:
126   - Wheat: 0
127   - Citrus: 0
128   - Jade: 0
129   - Cloves: 0
130   - Figs: 0
131   - Ivory: 0
132   - Gold: 0
133   - Nutmeg: 0
134   - Pearls: 0
135   - Rice: 0
136   - Cotton: 0
137
138 - Gems: 0
139 - Salt: 0
140 - Brazilwood: 0
141 - Crab: 0
142 - Amber: 0
143 - Coffee: 0
144 - Whales: 0
145 - Sugar: 0
146 - Cinnamon: 0
147 - Bananas: 0
148 - Pepper: 0
149 - Marble: 0
150 - Coral: 0
151 - Silk: 0
152 - Jewelry: 0
153 - Maize: 0
154 - Glass: 0
155 - Lapis Lazuli: 0
156 - Copper: 0
157 - Truffles: 0
158 - Perfume: 0
159 - Incense: 0
160 - Olives: 0
161 - Silver: 0
162 - Furs: 0
163 - Dyes: 0
164 - Tea: 0
165 - Fish: 0
166 - Wine: 0
167 - Porcelain: 0
168 - Cocoa: 0
169 - Tobacco: 0
170 - Gold: 15
171 - Era: Ancient
172 - GoldenAge: Estimated in 102 turns
173 - HappinessPercentage: 100
174 - Technologies: 2
175 - Cities: 1
176 - Population: 3
177 - CurrentResearch: The Wheel
178 - SciencePerTurn: 7
179 - GoldPerTurn: 4
180 - FaithPerTurn: 0
181 - CulturePerTurn: 3
182 - HappinessSituation: Happy
183 - MyEvaluations:
184   - 0: Our top friend: None
185   - 1: Our top Competitor: Songhai
186   - 1: Unmet Major Civilization
187   - 2: Unmet Major Civilization
188
189 ## Player 3
190 - TeamID: 3
191 - Civilization: Songhai
192 - Leader: Askia
193 - IsMajor: true
194 - Territory: 7
195 - Score: 40
196 - TourismPerTurn: 0
197 - Gold: 59
198 - Era: Ancient
199 - Technologies: 2

```



```

199 - Cities: 1
200 - Population: 3
201 - GoldPerTurn: 4
202 - CulturePerTurn: 1
203 - HappinessSituation: Happy
204 - OpinionFromMe:
205   - 0: Our real approach is Neutral.
206 - OpinionToMe:
207   - 0: You have no contested borders.
208   - 1: They have no strong feelings about us one way or the
      other.
209 - 22: Unmet Minor Civilization
210 - 23: Unmet Minor Civilization
211 - 24: Unmet Minor Civilization
212
213 ## Player 25
214 - TeamID: 25
215 - Civilization: Trowulan
216 - Leader: City State
217 - Relationships:
218   - Songhai: Neutral (Influence: 0)
219   - Poland: Neutral (Influence: 3)
220 - Population: 1
221 - Technologies: 2
222 - 26: Unmet Minor Civilization
223 - 27: Unmet Minor Civilization
224
225 ## Player 28
226 - TeamID: 28
227 - Civilization: Vienna
228 - Leader: City State
229 - Relationships:
230   - Songhai: Neutral (Influence: 0)
231   - Poland: Neutral (Influence: 0)
232 - Population: 1
233 - Technologies: 2
234 - 29: Unmet Minor Civilization
235 - 63: Barbarians
236
237 # Cities
238 Cities: summary reports about discovered cities in the world
239
240 ## Player: Poland
241 - Warsaw:
242   - ID: 1029
243   - X: 11
244   - Y: 26
245   - Population: 3
246   - DefenseStrength: 14
247   - FoodStored: 2
248   - FoodPerTurn: 1
249   - ProductionStored: 20
250   - ProductionPerTurn: 7
251   - CurrentProduction: Shrine
252   - ProductionTurnsLeft: 2
253   - GoldPerTurn: 5
254   - SciencePerTurn: 7
255   - CulturePerTurn: 3
256   - FaithPerTurn: 0
257   - TourismPerTurn: 0
258   - HappinessDelta: 9
259   - BuildingCount: 1
260
261 - WonderCount: 1
262 - GreatWorkCount: 0
263
264 ## Player: The Iroquois
265 - Onondaga:
266   - ID: 1030
267   - X: 24
268   - Y: 24
269   - Population: 3
270   - DefenseStrength: 9
271
272 ## Player: City-States
273 - Trowulan:
274   - ID: 1036
275   - X: 4
276   - Y: 34
277   - Population: 1
278   - DefenseStrength: 13
279   - IsCoastal: true
280 - Vienna:
281   - ID: 1039
282   - X: 14
283   - Y: 34
284   - Population: 1
285   - DefenseStrength: 15
286
287 # Military
288 Military: summary reports about tactical zones and visible
289 units.
290
291 ## Unit Stats
292 - Ranged:
293   - Slinger:
294     - Strength: 5
295     - RangedStrength: 6
296 - Explore:
297   - Scout:
298     - Strength: 10
299   - Warrior:
300     - Strength: 8
301
302 ## Friendly Land Zone 1029
303 - ZoneValue: 60
304 - Dominance: Friendly
305 - FriendlyStrength: 49
306 - City: Warsaw
307 - AreaID: 19
308 - Plots: 66
309 - CenterX: 11
310 - CenterY: 27
311 - Units:
312   - Songhai:
313     - Warrior: 1
314     - Poland:
315       - Slinger: 1
316 - Neighbors:
317   - 0: 1036
318
319 ## Neutral Land Zone 1036
320 - ZoneValue: 24
321 - Dominance: Friendly
322 - NeutralStrength: 13
323 - FriendlyStrength: 34

```

```

320 - City: Trowulan
321 - AreaID: 19
322 - Plots: 36
323 - CenterX: 5
324 - CenterY: 32
325 - Units:
326   - City-State Trowulan:
327     - Warrior: 1
328   - Poland:
329     - Scout: 1
330     - Warrior: 1
331 - Neighbors:
332   - 0: 1029
333
334 # Events
335 Events: events since you last made a decision.
336 ## Turn 11
337
338 ### 0
339 - Type: TileRevealed
340 - Unit:
341   - Level: 2
342   - AI: Explore
343   - Unit: Warrior
344   - ID: 1011
345 - Player: 3: Songhai
346 - Plots:
347   - Plains Land 8,27
348   - Plains Hills 8,29
349
350 ### 1
351 - Type: UnitMoved
352 - Unit:
353   - Level: 2
354   - AI: Explore
355   - Unit: Warrior
356   - ID: 1011
357 - Player: 3: Songhai
358 - Plains Land 10,28
359
360 ### 2
361 - Type: TileRevealed
362 - Unit:
363   - Level: 2
364   - AI: Explore
365   - Unit: Warrior
366   - ID: 1011
367 - Player: 3: Songhai
368 - Plots:
369   - Plains Hills 9,26
370   - Desert Land 10,25:
371     - Owner: Poland
372     - IsRiver: true
373     - Feature: Flood Plains
374   - Desert Land 9,25:
375     - IsRiver: true
376     - Feature: Flood Plains
377   - Desert Land 8,25
378   - Plains Hills 8,26
379   - Plains Hills 7,27
380
381 ### 3

```

```

382 - Type: UnitMoved
383 - Unit:
384   - Level: 2
385   - AI: Explore
386   - Unit: Warrior
387   - ID: 1011
388 - Player: 3: Songhai
389 - Plains Hills 9,27:
390   - IsRiver: true
391   - Feature: Forest
392
393 ### 4
394 - Type: PlayerDoneTurn
395 - NextPlayer: 0: Poland
396 - Player: 63: Barbarians
397
398 ## Turn 12
399
400 ### 0
401 - Type: SetPopulation
402 - OldPopulation: 2
403 - NewPopulation: 3
404 - City:
405   - Owner: Poland
406   - CityID: 1029
407   - Population: 3
408   - IsRiver: true
409   - City: Warsaw
410   - Terrain: Plains
411   - Route: Road
412   - Plot: Hills
413   - X: 11
414   - Y: 26
415
416 ### 1
417 - Type: PlayerDoTurn
418 - Player: 0: Poland
419
420 ### 2
421 - Type: UnitPromoted
422 - Unit:
423   - Level: 5
424   - AI: Explore
425   - Unit: Scout
426   - ID: 1051
427 - Player: 0: Poland
428 - Promotion: Medic I
429
430 ### 3
431 - Type: UnitMoved
432 - Unit:
433   - Level: 5
434   - AI: Explore
435   - Unit: Scout
436   - ID: 1051
437 - Player: 0: Poland
438 - Plots:
439   - Tundra Hills 9,32:
440     - Feature: Forest
441   - Tundra Land 8,32:
442     - Feature: Forest
443   - Plains Land 7,32

```

```

444
445 ### 4
446 - Type: UnitMoved
447 - Unit:
448   - Level: 1
449   - AI: Explore
450   - Unit: Warrior
451   - ID: 1002
452 - Player: 0: Poland
453 - Plains Land 5,32:
454   - IsRiver: true
455
456 ### 5
457 - Type: UnitMoved
458 - Unit:
459   - Level: 5
460   - AI: Explore
461   - Unit: Scout
462   - ID: 1051
463 - Player: 0: Poland
464 - Tundra Land 8,32:
465   - Feature: Forest
466
467 ### 6
468 - Type: UnitMoved
469 - Unit:
470   - Level: 2
471   - AI: Ranged
472   - Unit: Slinger
473   - ID: 1054
474 - Player: 0: Poland
475 - Plots:
476   - Desert Land 10,25:
477     - Owner: Poland
478     - IsRiver: true
479     - Feature: Flood Plains
480   - Plains Hills 11,26:
481     - Owner: Poland
482     - CityID: 1029
483     - Population: 3
484     - IsRiver: true
485     - City: Warsaw
486     - Route: Road
487
488 ### 7
489 - Type: PlayerDoneTurn
490 - NextPlayer: 1: The Iroquois
491 - Player: 0: Poland

```

Listing 2: User Prompt

B Example LLM Output

B.1 Reasoning - around 446 Tokens

```

1 - Name: keep-status-quo
2 Let me assess the current situation for Poland on turn 12:
3
4 **Current State Analysis:**
5 - Warsaw has grown to population 3 (this just happened)
6 - Currently building a Shrine (2 turns left)
7 - Culture per turn: 3 (will be 4 after Monument effect)

```

```

8 - Faith per turn: 0 (will start after Shrine completes)
9 - Science per turn: 7
10 - Gold per turn: 4
11 - Territory: 7 (good expansion potential)
12 - Research: Currently on The Wheel, next selected as Mining
13 - Policy: Next selected as Tradition (new branch)
14
15 **Military/Exploration:**
16 - Scout is level 5 with Medic I promotion (very advanced!)
17 - Warrior exploring near Trowulan
18 - Slinger defending Warsaw
19 - Songhai Warrior nearby but neutral relationship
20
21 **Strategic Situation:**
22 - Culture Grand Strategy is working well
23 - Early Expansion is the current economic strategy
24 - No active military strategies (peaceful development)
25 - Good progress with Trowulan city-state (3 influence)
26 - Songhai identified as main competitor
27
28 **Key Events:**
29 - Warsaw just grew to size 3
30 - Scout got promoted to Medic I (this is excellent for
31   healing units)
32 - More terrain revealed around Warsaw including flood plains
33   and hills
34
35 **Decision Points:**
36 1. Should I maintain the current strategy? It seems to be
37   working very well.
38 2. Research: Mining is already selected, which makes sense
39   for production
40 3. Policy: Tradition is selected, which aligns with Culture
41   strategy
42 4. Economic strategy: Early Expansion makes sense given the
43   territory
44
45 **Considerations:**
46 - The Scout's Medic I promotion is very valuable - this unit
47   can now support military operations by healing units
48 - Warsaw's growth to 3 is excellent for increasing
49   production potential
50 - The Shrine will start faith generation soon
51 - The Songhai Warrior is nearby but neutral - no immediate
52   threat
53
54 - Our cultural advancement is progressing steadily, with the
55   Monument near completion promising a +1 culture boost
56
57 I'll continue focusing on cultural development and expansion
58 . The Monument's impending completion aligns perfectly
59   with our strategy, setting the stage for steady growth
60   and cultural achievements.

```

Listing 3: Reasoning

B.2 Tool Calling - around 99 Tokens

```

1 - Name: keep-status-quo

```

- 2 - Rationale: Our strategy is executing perfectly. Warsaw just grew to population 3, boosting development capacity. The Scout's Medic I promotion enhances our military flexibility. Shrine completion in 2 turns will start faith generation. Cultural progress with Monument (3-4 culture/turn) accelerates toward our first free Social Policy. Exploration has revealed valuable resources and maintains good positioning. Early Expansion with Culture focus remains optimal for Poland's unique abilities. No changes warranted.

Listing 4: Tool Calling