

---

# The Othello AI Arena: Evaluating Intelligent Systems Through Limited-Time Adaptation to Unseen Boards

---

**Sundong Kim**

Gwangju Institute of Science and Technology  
sundong@gist.ac.kr

## Abstract

The ability to rapidly adapt to novel and unforeseen environmental changes is a cornerstone of artificial general intelligence (AGI), yet it remains a critical blind spot in most existing AI benchmarks. Traditional evaluation largely focuses on optimizing performance within fixed environments, failing to assess systems' flexibility and generalization capabilities when faced with even subtle rule or structural modifications. Addressing this gap, I introduce the Othello AI Arena, a novel benchmark framework designed to evaluate intelligent systems based on their capacity for limited-time adaptation to unseen environments. Our platform poses a meta-learning challenge: participants must develop systems that can analyze the specific configuration and rules of a novel Othello board within a strict time limit ( $\approx 60$  seconds) and generate a tailored, high-performing strategy for that unique environment. With this, evaluation of the meta-level intelligence (which performs the analysis and strategy generation) can be separated from the task-level strategy performance of the generated strategy. The Arena features a diverse set of game stages, including public stages for development and private stages with structural and rule variations (such as altered board sizes, blocked cells, non-standard capture mechanics, and modified turn dynamics) designed to test genuine adaptive and generalization capabilities. Implemented as an accessible web-based platform, the Arena provides real-time visualization, automated evaluation using multi-dimensional metrics (including adaptation speed and efficiency), and comprehensive logging for post-hoc analysis. Initial observations from pilot tests and preliminary student engagements highlight fascinating patterns in adaptation approaches, ranging from rapid parameter tuning to rudimentary environmental model learning through simulation. The Othello AI Arena offers a unique educational tool and a valuable research benchmark for fostering and evaluating the crucial skill of rapid, intelligent adaptation in AI systems.

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>The Othello AI Arena Framework</b>	<b>3</b>
2.1	System Architecture . . . . .	4
2.2	Challenge Design . . . . .	4
2.3	API and Environment Interaction . . . . .	5
<b>3</b>	<b>Meta-Learning Formulation and Adaptation Mechanisms</b>	<b>5</b>
3.1	Adaptation Strategies under Time Constraints . . . . .	6
3.2	Connection to General Intelligence . . . . .	7
<b>4</b>	<b>Evaluation Methodology and Dataset Potential</b>	<b>7</b>

4.1	Evaluation Metrics . . . . .	7
4.2	Tournament Setup . . . . .	8
4.3	Dataset Structure and Potential . . . . .	9
<b>5</b>	<b>Benchmark Insights from Preliminary Experiments</b>	<b>9</b>
5.1	Illustrative Tournament Results . . . . .	9
5.2	Preliminary Observations on Adaptation Patterns . . . . .	10
5.3	Potential for Analysis and Extension . . . . .	10
5.4	Conclusion . . . . .	10
<b>A</b>	<b>System Architecture</b>	<b>12</b>
<b>B</b>	<b>Core Modules and Implementation Details</b>	<b>12</b>
B.1	Game Core Engine . . . . .	12
B.2	Intelligent System Execution Environment . . . . .	14
B.3	Strategy Management . . . . .	14
B.4	Game Flow and UI Management . . . . .	15
B.5	Logging and Replay System . . . . .	15
B.6	Tournament System . . . . .	16
<b>C</b>	<b>Othello Arena API Reference</b>	<b>16</b>
C.1	Example Intelligent System: Advanced Template . . . . .	17
<b>D</b>	<b>Board Variation Types</b>	<b>28</b>
D.1	Structural Variations . . . . .	28
D.2	Rule Variations . . . . .	28
D.3	Initial State Variations . . . . .	29
<b>E</b>	<b>Time Management Strategies in Intelligent System Development</b>	<b>30</b>
E.1	Leveraging $T_{\text{development}}$ : The Extended Preparation Phase . . . . .	30
E.2	Optimizing $T_{\text{analysis}}$ : Rapid Environment Modeling . . . . .	30
E.3	Managing $T_{\text{game}}$ : Constrained Real-Time Execution . . . . .	31
<b>F</b>	<b>Game Logging and Data Format</b>	<b>31</b>
F.1	Human-Readable Text Log . . . . .	31
F.2	Machine-Readable JSON Log . . . . .	33
<b>G</b>	<b>Human Adaptation to Novel Board Games: Observations</b>	<b>34</b>
G.1	Adaptation in Othello Arena Stages . . . . .	34
G.2	Adaptation in YINSH Game . . . . .	35

## 1 Introduction

Current AI benchmarks predominantly focus on measuring peak performance in static environments. While valuable, this approach overlooks the dynamic nature of real-world problems, where environments are constantly evolving, requiring intelligent agents to be robust and adaptable. Furthermore, many benchmarks evaluate only the final output or performance, without explicitly assessing the underlying process by which an intelligent system analyzes a new task and generates a suitable strategy. This makes it difficult to understand and evaluate the meta-cognitive abilities, such as efficient knowledge transfer, rapid model building, and strategic reasoning under uncertainty, that are critical for true general intelligence.

From an educational standpoint, providing students with opportunities to design systems that can adapt to changing circumstances offers a richer learning experience than simply optimizing for fixed problems. Such challenges cultivate not just problem-solving skills but also meta-reasoning [6], time management under computational constraints, and a deeper understanding of the exploration-

exploitation trade-off. Incorporating adaptive challenges is particularly relevant in AGI education, where fostering the ability to handle novelty is a core objective.

Addressing these limitations, I introduce the Othello AI Arena, a novel benchmark framework specifically designed to evaluate the adaptive intelligence of AI systems. This is grounded in the belief that a key aspect of AGI is the ability to quickly analyze an unfamiliar environment and synthesize an effective strategy within a limited timeframe. The Othello AI Arena provides a platform where participants must develop intelligent systems capable of this rapid adaptation. The Othello AI Arena is accessible at <https://sundong.kim/courses/agi25sp/othello-leaderboard/hw2.html>

**Othello AI Arena for Intelligent System (HW2)**

Black: 46 White: 14

Black: 0.13s / 10.00s White: 0.00s / 10.00s

Game over. Black wins! (46-14)

**Leaderboard**

Begin Stop Refresh

Competition results will appear here.

Rank	Strategy	Win Rate	W	L	D	G
1	smart-lv2	89.3%	12	1	1	14
2	intelligent_is-template-adv_Standard_8x8	75.0%	10	3	1	14
3	smart-lv1	71.4%	10	4	0	14
4	Corners	50.0%	7	7	0	14
5	Positional	50.0%	7	7	0	14
6	Greedy	42.9%	6	8	0	14
7	Random	21.4%	3	11	0	14
8	smart-lv3-updated-slow	0.0%	0	14	0	14

Upload Intelligent System (with Interaction) - [How to use API](#)

Your system will have maximum 60 seconds to analyze the stage and generate a tailored strategy.

Choose File No file chosen Upload Intelligent System

Upload Strategy (without Interaction) - [Sample Strategy](#)

Choose Files No file chosen Upload Strategies

smart-lv1 Edit Delete

smart-lv2 Edit Delete

smart-lv3-updated-slow Edit Delete

intelligent\_is-template-adv\_Standard\_8x8 Edit Delete

Stage: 3: 8x8 (Partial C-Squares-cw)

Black: smart-lv2

White: Greedy

Start Game Reset Board Clear All Data

Greedy(W): f7  
 smart-lv2(B): g6  
 Greedy(W): h7  
 smart-lv2(B): f5  
 Greedy(W): h6  
 smart-lv2(B): c1  
 Greedy(W): d1  
 smart-lv2(B): d8  
 Greedy passes  
 smart-lv2(B): e1  
 Greedy(W): f1  
 Game over: Final score 46-14

Save Log

Figure 1: Screenshot of the Othello AI Arena web interface on the "8 × 8 (Partial C-Squares-cw)" stage. The leaderboard displays the performance of various strategies, including a generated intelligent system strategy.

## 2 The Othello AI Arena Framework

The Othello AI Arena is a comprehensive platform for evaluating intelligent systems' adaptive capabilities in Othello [11]. It comprises core components, a structured challenge design with diverse environmental variations, and a defined API for system-environment interaction.

## 2.1 System Architecture

The modular architecture includes the User Interface (UI), Core Game Engine, Intelligent System Execution Environment, Evaluation Module, and Logging and Analysis System. Figure 3 illustrates their interaction. The UI (`hw2.html`, `app.js`, `game-ui.js`) provides web-based visualization and management. The Core Game Engine (`game-core.js`) implements Othello rules and stage variations. The Execution Environment (`intelligent-system-loader.js`) safely runs submitted code, enforcing  $T_{\text{analysis}}$  ( $\approx 60$  seconds) for strategy generation (`analyzeStage`) and  $T_{\text{game}}$  ( $\approx 10$  seconds total) for game execution. The Evaluation Module (`tournament.js`) manages tournaments and metrics. The Logging and Analysis System (`game-logger.js`) records detailed game data for post-hoc analysis (`Using-game-logs-for-world-model-learning.md`). The system flow involves system upload, `analyzeStage` execution per stage within  $T_{\text{analysis}}$ , returning a strategy, and game execution using the strategy under  $T_{\text{game}}$ , with logging throughout.

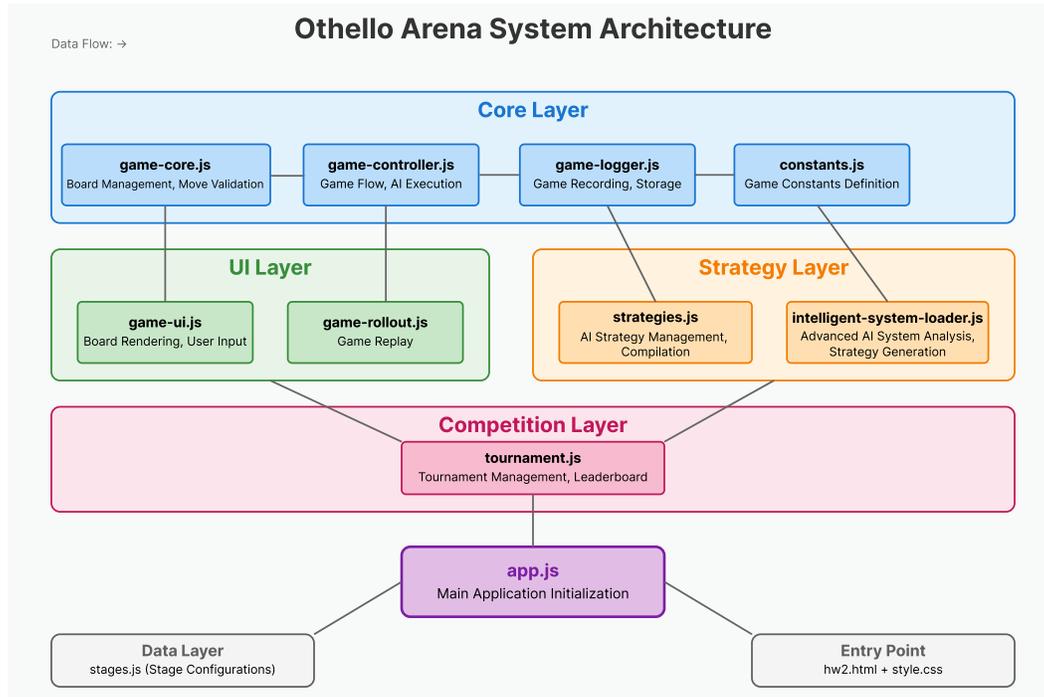


Figure 2: Overall system architecture of the Othello AI Arena, depicting modular layers and data flow. Intelligent systems interact primarily with components within the Strategy Layer.

## 2.2 Challenge Design

The challenge uses public (for development) and private (unseen for evaluation) stages to assess real-time adaptation. Stages feature variations probing different cognitive aspects. I categorize variations (`stages-extended.js`) not limited to below settings:

- **Structural Variations:** Modifications to the physical layout of the board.
  - *Board Size:* Different grid dimensions (e.g.,  $6 \times 6$ ,  $10 \times 10$  from standard  $8 \times 8$ ). This tests spatial generalization and the ability to adapt positional evaluation and search strategies to new scales.
  - *Blocked Cells:* Introducing impassable cells on the board. This challenges pathfinding, region control, and the re-evaluation of move validity and positional value in a constrained space.
- **Rule Variations:** Alterations to the fundamental game mechanics.
  - *Capture Mechanics:* Modifying how pieces are captured (e.g., ‘ignore occlusion’ where blocked cells don’t stop capture lines). This demands rapid rule induction and adaptation of move simulation and evaluation logic.

- *Turn Dynamics*: Changes to whose turn it is (e.g., ‘fewer pieces continue’ where the player with fewer pieces takes consecutive turns). This affects temporal planning and requires understanding and predicting the flow of the game based on piece counts.
- *Winning Conditions*: Altering the goal of the game (e.g., Reverse Othello, where the player with the *least* pieces wins). This challenges goal reorientation and counter-intuitive strategic thinking.
- **Initial State Variations**: Non-standard starting configurations.
  - *Non-standard Initial Placement*: Pieces starting in unusual positions. This invalidates standard opening books and requires dynamic early-game strategy generation.
  - *Pre-placed Special Pieces*: Introducing pieces with unique properties or interactions (potential future extension).

Each variation imposes a cognitive hurdle for the intelligent system within  $T_{\text{analysis}}$ . The generated task-level strategy function takes board state, player, and valid moves, returning a move within the cumulative  $T_{\text{game}}$  limit.

### 2.3 API and Environment Interaction

The Othello AI Arena provides a well-defined API<sup>12</sup> that serves as the sole channel through which the intelligent system can interact with and learn about a new stage during the 60-second analysis ( $T_{\text{analysis}}$ ) phase. Key functions include:

- `getValidMoves(board, player)`: Returns a list of valid moves for a player on a given board state, respecting the stage’s specific rules.
- `simulateMove(board, player, row, col)`: Simulates the outcome of making a specific move, returning the resulting board state and number of pieces captured. This is crucial for exploring game dynamics and evaluating potential move consequences without playing the game.
- `evaluateBoard(board, player)`: Provides an evaluation of a given board state using pre-defined metrics (e.g., piece count, mobility, corner control). Note that the weights or interpretation of these metrics might need to be learned or adjusted by the intelligent system based on the stage’s characteristics (as seen in `intelligent-system-template-adv.js`).

Systems effectively use these APIs within  $T_{\text{analysis}}$  (e.g., via simulation-based learning as in `intelligent-system-template-adv.js`) to model the environment and synthesize a strategy. A successful intelligent system effectively utilizes these tools within the time limit. For instance, the template in `intelligent-system-template-adv.js` uses `api.simulateMove` and `api.getValidMoves` extensively in its self-play loop to explore the stage’s game dynamics and collect data (e.g., position win rates, evidence for rule variations) from which it could learn a tailored strategy. This process exemplifies how systems must learn an environmental model and synthesize a strategy through active interaction constrained by time and the API.

Game logging (`game-logger.js`) captures detailed information about every move and board state during gameplay. While not directly part of the ‘analyzeStage’ API for the *current* stage analysis, these logs from *previous* games on public stages can potentially be leveraged by an intelligent system during its analysis phase to transfer knowledge or identify patterns (`Using-game-logs-for-world-model-learning.md`).

Move detailed API descriptions and code snippets to Appendix.

## 3 Meta-Learning Formulation and Adaptation Mechanisms

The adaptation task in the Othello AI Arena can be formalized as follows. Let  $\mathcal{S} = \{s_1, s_2, \dots, s_N\}$  be the set of possible game stages, where each stage  $s \in \mathcal{S}$  is defined by its board configuration, rule set, and initial state. Some stages are public ( $s \in \mathcal{S}_{\text{public}}$ ), while others are private and unseen during development ( $s \in \mathcal{S}_{\text{private}}$ ). An intelligent system  $I$  is evaluated on its ability to adapt to a given stage  $s$ .

<sup>1</sup><https://sundong.kim/courses/agi25sp/othello-leaderboard/Intelligent-system-API-reference/>

<sup>2</sup><https://sundong.kim/courses/agi25sp/othello-leaderboard/Using-game-logs-for-world-model-learning/>

For each stage  $s$ , the intelligent system  $I$  is granted a limited analysis time  $T_{\text{analysis}} \approx 60$  seconds. During this time,  $I$  can interact with the environment defined by  $s$  solely through a predefined API (as described in Section 2.3). The goal of the system during this phase is to output a task-specific strategy function  $f_s$ :

$$I(s, \text{API}) \xrightarrow{T_{\text{analysis}}} f_s \quad (1)$$

This function  $f_s$  takes the current board state  $b$ , player ID  $p$ , and set of valid moves  $M$  as input, and must return a chosen move. The cumulative time taken by  $f_s$  across all its calls within a single game is limited to  $T_{\text{game}} \approx 10$  seconds:

$$\sum_{\text{moves in game}} \text{Time}(f_s(b, p, M) \rightarrow m) \leq T_{\text{game}} \quad (2)$$

The performance of the intelligent system  $I$  on stage  $s$  is evaluated based on the effectiveness of the generated strategy  $f_s$  in playing games on stage  $s$ . The overall evaluation score for system  $I$  is a composite metric considering various factors across a set of evaluation stages  $\mathcal{S}_{\text{eval}} \subseteq \mathcal{S}_{\text{private}}$ . Key metrics (Section 4.1) are Task Performance ( $P$ ), Adaptation Speed ( $A$ ), Efficiency ( $E$ ), Generalization ( $G$ ), and Robustness ( $R$ ). This aligns with few-shot adaptation [1] under tight constraints.

### 3.1 Adaptation Strategies under Time Constraints

Effectively addressing the Othello AI Arena challenge requires intelligent systems to devise strategies for analyzing a new environment and generating a tailored strategy within the  $T_{\text{analysis}}$  time limit. This involves an exploration-exploitation trade-off: how much time should be spent exploring the new environment’s dynamics versus optimizing the strategy based on the acquired knowledge?

Successful adaptation mechanisms in this context often involve a combination of techniques as:

- **Environmental Modeling through Interaction:** Since the intelligent system does not have explicit knowledge of the specific rule or structural variations of a novel stage beyond what is revealed through interaction via the API, it must learn an implicit or explicit model of that environment’s dynamics. A common approach, demonstrated in `intelligent-system-template-adv.js`, is to perform rapid self-play simulations using the provided API functions (`getValidMoves`, `simulateMove`, `evaluateBoard`). By playing a large number of simulated games ( $\approx 3000$  games in the template) within the  $T_{\text{analysis}}$  window, the system can:<sup>3</sup>
  - Infer the rules of piece movement and capture, including any variations like ‘ignore occlusion’ or ‘fewer pieces continue’ (as shown by the rule detection logic in the template).
  - Gather statistics on the value of different board positions or move sequences in the context of the current stage (e.g., calculating position win rates in the template).
  - Understand the typical game flow and identify effective opening sequences for this specific environment.
- **Learning Environmental Features:** Based on the data gathered through simulation or initial board inspection, the system can identify key features of the environment, such as board size (`stageConfig.boardSize`), presence of blocked cells (`hasBlockedCells` in the code snippet), or detected rule variations. These features inform the subsequent strategy generation.
- **Adaptive Strategy Synthesis:** Using the learned environmental model and identified features, the intelligent system synthesizes a strategy function  $f_s$ . This can involve:
  - Adjusting parameters of a general Othello strategy template (e.g., modifying positional weights or search depth based on board size or blocked cells).
  - Selecting from a portfolio of pre-existing strategy components or heuristics based on the detected environmental features.
  - Constructing a new strategy using learned elements, such as a data-driven positional value matrix (`positionValueMatrix`) or an opening book (`openingBook`) derived from the self-play simulations in the current stage.

<sup>3</sup>Notably, this level of simulation ( $\approx 3000$  games) is significantly higher than what a human typically requires (e.g., 3-4 games) to grasp the rules and achieve intermediate play, highlighting a current difference in learning efficiency between such systems and human adaptability [7].

- **Efficient Time Allocation:** The  $T_{\text{analysis}}$  constraint necessitates a trade-off between the depth of environmental exploration (e.g., number of simulations) and the complexity of the strategy synthesis process. Successful systems are expected to adaptively manage the total  $T_{\text{game}}$  budget, spending more time exploring highly novel environments and perhaps relying on very fast heuristics or shallow search for familiar moves to stay within the cumulative 10-second limit.

This process requires a meta-cognitive ability to understand the task at hand (the specific stage), assess what information is needed for effective play, and efficiently acquire and utilize that information within the given time budget.

### 3.2 Connection to General Intelligence

The challenges posed by the Othello AI Arena are deeply connected to the core principles of general intelligence. True AGI is not merely about excelling at a single task but demonstrating robust performance across a wide range of tasks and the capacity to adapt to novel ones.

- **Rapid Task Acquisition:** The limited-time adaptation challenge directly evaluates a system’s ability to quickly understand and become proficient in a new task (stage), a fundamental aspect of intelligent behavior.
- **Environmental Modeling:** Adapting to unseen rule and structural variations requires building or updating an internal model of the environment’s dynamics. The simulation-based learning approach in `intelligent-system-template-adv.js` is one method for achieving this.
- **Transfer Learning and Generalization:** Effective intelligent systems should leverage knowledge gained from public stages or prior experience to accelerate adaptation on private stages [10]. This requires abstracting principles that generalize across different Othello variants.
- **Meta-Cognition and Resource Management:** The time constraints compel systems to reason about their own cognitive processes – deciding how to allocate the  $T_{\text{analysis}}$  budget, when to stop exploring, and which learning methods are most efficient for the given stage.
- **Flexible Strategy Synthesis:** Instead of relying on fixed algorithms, intelligent systems must be able to compose or modify strategies based on learned environmental properties, demonstrating strategic flexibility.

By providing a controlled yet challenging environment for evaluating these capabilities, the Othello AI Arena serves as a valuable microcosm for studying and advancing the development of general intelligent agents.

## 4 Evaluation Methodology and Dataset Potential

Evaluating adaptive intelligence requires metrics and methodologies that go beyond traditional performance measures in static environments. The Othello AI Arena employs a multi-faceted approach to assess the capabilities of submitted intelligent systems and offers the potential for a valuable dataset for future research.

### 4.1 Evaluation Metrics

Performance of intelligent systems is evaluated based on several key dimensions, designed to capture different facets of adaptive intelligence:

- **Task Performance ( $P$ ):** This is the most straightforward metric, measuring the success of the generated strategy function ( $f_s$ ) in playing games on the target stage  $s$ . It is typically quantified by win rate, average score difference against various opponents (including built-in strategies like Random, Greedy, Corners, Positional, and potentially other submitted intelligent systems), and the number of pieces controlled at the end of the game.
- **Adaptation Speed ( $A$ ):** This measures how quickly an intelligent system is able to generate an effective strategy for a *new* stage. In the context of the 60-second analysis limit, adaptation speed can be implicitly assessed by the performance of the generated strategy in the initial games played on a private stage. A system that rapidly identifies key environmental features and generates a

reasonable strategy will likely perform better early on than one that is slow to understand the new environment. Future metrics could involve explicitly measuring performance within a fixed number of initial moves or games on the new stage.

- **Efficiency ( $E$ ):** This dimension evaluates the computational cost of the intelligent system within the imposed time limits. It includes:
  - *Analysis Time Utilization:* How effectively the system uses the allocated  $T_{\text{analysis}}$  ( $\approx 60$  seconds). This is measured by the actual time taken by the `analyzeStage` function. Systems that exceed this limit fail to produce a strategy.
  - *Game Time Utilization:* How effectively the generated strategy function  $f_s$  uses its *total game time budget of  $\approx 10$  seconds* across all moves in a single game. Exceeding this cumulative limit results in a time forfeit. This constraint necessitates fast individual move computations, typically on the order of tens or a couple of hundred milliseconds per move.
- **Generalization ( $G$ ):** A critical measure of adaptive intelligence is the ability to perform well on unseen tasks. Generalization is assessed by comparing the performance of intelligent systems on the public training stages versus the private evaluation stages. A small performance drop on private stages indicates strong generalization capabilities.
- **Adaptation Robustness ( $R$ ):** This is evaluated by testing the system’s consistent performance across a *variety* of different, challenging stages, including those with combinations of structural and rule variations. A system that maintains strong performance across diverse unseen environments demonstrates high adaptation robustness.

A final evaluation score for an intelligent system is computed as a weighted combination of these dimensions, allowing the benchmark to highlight systems that are not just strong players but are also fast, efficient, and robust adaptors to novel environments.

Central to the philosophy of the Othello AI Arena, and aligned with definitions of intelligence emphasizing skill-acquisition efficiency (e.g., [4]), is the concept that intelligence is demonstrated not merely by high performance (skill), but by the efficiency with which that skill is acquired and generalized to novel situations, controlling for prior knowledge and experience. From this perspective, a system that relies heavily on a large amount of pre-encoded prior knowledge tailored to potential variations might be considered less intelligent than one that can achieve comparable performance and generalization with less prior knowledge, by more effectively learning and adapting from limited exposure to the new environment.

This understanding motivates potential benchmark design choices, such as imposing limits on the size of the intelligent system (as a proxy for the amount of prior knowledge it contains), or restricting access to external information sources or powerful pre-trained models (like large language models) during the time-constrained analysis phase. Similar to approaches seen in benchmarks like the ARC-prize, such constraints would further emphasize the system’s ability to generalize and adapt efficiently based on its intrinsic capabilities and the limited information available within the benchmark’s environment, rather than relying on extensive external resources.

## 4.2 Tournament Setup

The evaluation of intelligent systems and their generated strategies is conducted using an automated tournament system (`tournament.js`). While referred to as a tournament, the evaluation on each stage follows a round-robin format. Submitted intelligent systems analyze each private stage sequentially within the  $T_{\text{analysis}}$  limit, producing a tailored strategy function ( $f_s$ ) for that specific stage. These generated strategies then compete against each other and against a set of baseline built-in strategies (e.g., Random, Greedy, Positional from `strategies.js`), as well as various advanced strategies (including those inspired by Logistello [3]), on the corresponding stage.

For each stage in the evaluation set:

1. Each intelligent system’s `analyzeStage` function is executed with the stage configuration and API, strictly limited to  $T_{\text{analysis}}$ .
2. If a strategy function  $f_s$  is successfully generated within the time limit, it is entered into the round-robin competition for stage  $s$ .

3. In the round-robin, each generated strategy plays against every other strategy (including baselines and advanced strategies) as both Black and White.
4. Strategies compete in games on stage  $s$ , with each strategy function's total time for the entire game limited to  $T_{\text{game}} \approx 10$  seconds.
5. Match results (win/loss/draw, final scores) are recorded by the logging system (`game-logger.js`).
6. Performance metrics ( $P, A, E$ ) are calculated for each generated strategy on stage  $s$ .

After evaluating systems across all private stages, the Generalization ( $G$ ) and Robustness ( $R$ ) metrics are computed. The final leaderboard, prominently displayed on the web interface (Top right corner of Fig 1 illustrates a sample leaderboard display), is then generated based on the aggregated weighted scores, with the Task Performance ( $P$ ) metric serving as the primary visual representation of evaluation results for each strategy on a given stage.

### 4.3 Dataset Structure and Potential

The Othello AI Arena platform is designed to generate a rich and unique dataset for research into adaptive AI. A key feature enabling this is the "Save Log" function, which allows users to export comprehensive data from individual games or full tournaments for offline analysis. This data serves as a valuable record of intelligent system behavior and environmental interactions under various conditions. For every evaluation run of an intelligent system on a stage, the system captures detailed information, including:

- **Stage Configuration:** Full details of the environment presented.
- **Intelligent System Code:** Source code capturing the adaptive logic.
- **Generated Strategy Code:** Source code of the task-level strategy function.
- **Analysis Process Logs & Time:** Outputs and time usage during the  $T_{\text{analysis}}$  phase, offering insight into the analysis and learning process.
- **Game Logs & Results:** Detailed records of gameplay, including move sequences, board states, actions, time usage per move, and final outcomes (`game-logger.js`).

The detailed game data can be exported in two formats: a human-readable text log and a structured JSON format for automated analysis. See Appendix for examples and further details.

## 5 Benchmark Insights from Preliminary Experiments

Full-scale experimental results involving a cohort of diverse intelligent systems across the complete set of private stages are currently being collected as part of an ongoing educational course. However, preliminary observations from pilot tests with prototype intelligent systems (e.g., variants of `intelligent-system-template-adv.js`) and received questions provide valuable insights into the types of results the Othello AI Arena yields and the patterns of adaptation observed.

### 5.1 Illustrative Tournament Results

Figure 1 shows results from a pilot tournament on "8x8 (Partial C-Squares-cw)", including built-in, advanced, and a prototype intelligent system's generated strategy. Strategies are ranked by Task Performance ( $P$ )—Win Rate.

- Simple strategies like Random and Greedy show relatively lower performance on the leaderboard.
- More sophisticated fixed strategies like `smart-lv2` (from `smart-lv2.js`, having alpha-beta pruning with reasonable MCTS depth, and some opening books with robust heuristics) perform better.
- The strategy generated by the prototype intelligent system is competitive, outperforming some fixed strategies but not necessarily the best-performing ones. This highlights the challenge: generating a near-optimal strategy for a novel environment within a strict time limit is difficult, but achievable to a degree that surpasses non-adaptive or poorly generalizing fixed strategies.

This illustrative example demonstrates how the Othello AI Arena provides quantitative comparisons of strategy performance on specific, potentially unseen, environmental variations.

## 5.2 Preliminary Observations on Adaptation Patterns

Beyond win/loss records, the Othello AI Arena enables deeper analysis of adaptation. Preliminary observations from pilot tests suggest several patterns:

- **Time Allocation:** Systems varied in using the  $T_{\text{analysis}}$  budget. A balanced approach between environmental modeling (simulation) and strategy synthesis likely correlates with better adaptation.
- **Game Time Limit Impact:** The strict  $T_{\text{game}}$  limit ( $\approx 10\text{s}$  total) constrains computationally intensive strategies (like deep MCTS), favoring efficient, adaptive heuristics.
- **Simulation-Based Learning:** Self-play simulations within  $T_{\text{analysis}}$  improved robustness to structural variations by learning stage-specific values.
- **Rule Adaptation Challenges:** Adapting to subtle rule variations appeared more challenging than structural changes, often requiring explicit detection logic or highly flexible simulation.
- **Generalization Gap:** Fixed strategies showed larger performance drops on private vs. public stages than adaptive systems, illustrating the specialization/generalization trade-off.

These preliminary findings demonstrate the benchmark’s ability to reveal distinct adaptation patterns under time constraints.

## 5.3 Potential for Analysis and Extension

The detailed game logs and analysis outputs collected by the Othello AI Arena (as described in Section 4.3) enable extensive post-hoc analysis of adaptation processes. Researchers can analyze correlations between analysis behaviors and strategy quality, system strategy adjustments, the relationship between strategy complexity and performance, and successes/failures in adapting to variations through log inspection. This transparency is a key strength supporting fine-grained analysis.

Building upon this framework, the Othello AI Arena is highly extensible and offers numerous avenues for future work and more complex challenges. These potential expansions aim to explore broader concepts of generalization and more sophisticated adaptive mechanisms. This could involve introducing challenges with incomplete information, such as limited observation or reasoning under uncertainty, or exploring multi-agent interaction scenarios, like N-Player or cooperative games. Further complexity could be added through dynamic environments where rules or configurations change during gameplay, or by requiring cross-domain adaptation to different game mechanics or environments. Additionally, the evaluation itself could evolve to include adaptive elements, such as variable time limits or automated curriculum generation based on system performance. These directions provide a roadmap for increasing benchmark complexity and pushing the boundaries of intelligent adaptation research.

## 5.4 Conclusion

In conclusion, this paper introduces the Othello AI Arena, a novel benchmark designed to address the critical gap in evaluating AI systems’ capacity for limited-time adaptation to unseen environments. I frame this challenge as a meta-learning problem, where participants develop intelligent systems capable of analyzing a novel Othello stage configuration and rules within a strict time limit ( $T_{\text{analysis}} \approx 60\text{s}$ ) to generate a tailored, high-performing strategy. The benchmark explicitly evaluates the meta-level intelligence of the system, separating it from the task-level performance of the generated strategy. The Othello AI Arena platform supports this higher-level competition, enabling the submission and evaluation of adaptive intelligent systems through tournaments with multi-dimensional metrics. While preliminary results offer initial insights into adaptation patterns, the comprehensive data generated also provides a valuable resource for analyzing the adaptation process itself. By focusing on rapid, intelligent adaptation under realistic constraints, the Othello AI Arena serves as a valuable tool for fostering and evaluating progress towards artificial general intelligence.

## References

- [1] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language Models Are Few-Shot Learners. In *NeurIPS*, 2020.
- [2] Guga Burduli and Jie Wu. Time Management in a Chess Game through Machine Learning. *International Journal of Parallel, Emergent and Distributed Systems*, pages 14–34, 2022.
- [3] Michael Buro. LOGISTELLO — A Strong Learning Othello Program. *19th Annual Conference Gesellschaft für Klassifikation*, 1995.
- [4] François Chollet. On the Measure of Intelligence. *arXiv:1911.01547*, 2019.
- [5] DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability with process-supervised reward models. *arXiv preprint arXiv:2501.12948*, 2025.
- [6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *ICML*, 2017.
- [7] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building Machines That Learn and Think Like People. *Behavioral and Brain Sciences*, 40:e253, 2017.
- [8] R. Lake and M. Hines. A world-championship-level othello program. *Artificial Intelligence*, 18(2):197–222, 1982.
- [9] I. Liao and A. Gu. Minizero: Comparative analysis of alphazero and muzero on go, othello, and atari. *arXiv preprint arXiv:2310.11305*, 2023.
- [10] Scott Reed, Konrad Zolna, Emilio Parisotto, Sergio Gomez Colmenarejo, Alexander Novikov, Gabriel Barth-Maron, Mai Gimenez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, Tom Eccles, Jake Bruce, Ali Razavi, Ashley Edwards, Nicolas Heess, Yutian Chen, Raia Hadsell, Oriol Vinyals, Mahyar Bordbar, and Nando de Freitas. A Generalist Agent. *Transactions on Machine Learning Research*, 2022.
- [11] Brian Rose. *Othello: A Minute to Learn... A Lifetime to Master*. Anjar Co., 2005.
- [12] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.

## A System Architecture

The Othello AI Arena’s modular architecture is designed for clarity, maintainability, and extensibility, facilitating the development and evaluation of intelligent systems. Figure 2 presents a comprehensive overview of the system’s key components and their hierarchical interactions, with a specific focus on how data and control flow through the various layers.

The architecture is broadly divided into several logical layers: the *Core Layer* handling fundamental game logic, the *UI Layer* for rendering and interaction, the *Strategy Layer* managing AI behaviors, and the *Competition Layer* orchestrating tournaments. These layers are all orchestrated by the main `app.js` and rely on a *Data Layer* for configurations and an *Entry Point* for application loading. This layered approach allows participants to focus on developing the Intelligent System (primarily interacting with the Strategy Layer) while the platform handles underlying complexities.

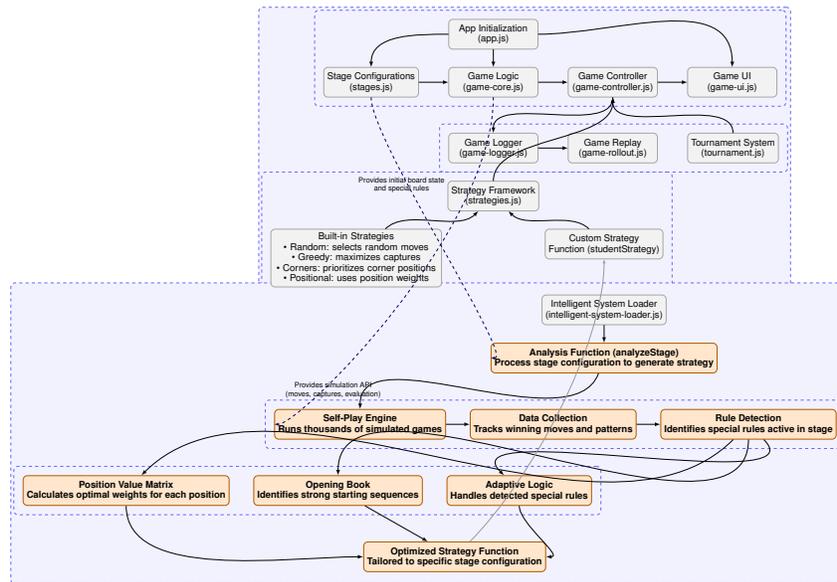


Figure 3: System Architecture of the Othello AI Arena (Gray part) and its interaction with the potential intelligent system (Orange part).

## B Core Modules and Implementation Details

The Othello AI Arena is built upon a modular JavaScript architecture, separating concerns into distinct modules that interact through well-defined interfaces. This design facilitates maintainability, extensibility, and clarity in understanding the system’s operation. Below, I detail the core modules that underpin the arena’s functionality, illustrating their responsibilities and key code components.

### B.1 Game Core Engine

The `game-core.js` module is the heart of the Othello Arena, implementing the fundamental game rules and variations. It manages the board state, validates moves, executes piece captures, and determines the next player based on the active stage’s rules. This module is critical as its behavior (e.g., how `isValidMove` or `makeMove` function under different rule variations) is what intelligent systems must infer through observation via the API.

- **Key Responsibilities:** Board initialization and management, move validation (`isValidMove`), move execution (`makeMove`), piece counting (`countDiscs`), and determining the next player (`determineNextPlayer`).
- **Illustrative Code Snippet (Move Execution Logic):** This snippet demonstrates how a move is executed, involving placing a piece and flipping captured pieces. The internal logic for flipping

pieces must implicitly handle stage-specific rules like 'ignore occlusion' without exposing these rules to the intelligent system API.

---

```
1 // Excerpt from game-core.js: Core move execution logic
2 function makeMove(row, col, player) {
3   if (!isWithinBoard(row, col) || board[row][col] !== GAME_CONSTANTS.EMPTY) return
4     ↪ false;
5   board[row][col] = player; // Place the player's piece
6
7   const opponent = player === GAME_CONSTANTS.BLACK ? GAME_CONSTANTS.WHITE :
8     ↪ GAME_CONSTANTS.BLACK;
9   const directions = [[-1, -1], [-1, 0], [-1, 1], [0, -1], [0, 1], [1, -1], [1,
10     ↪ 0], [1, 1]];
11
12 // Internal rule application based on currentStage configuration (not visible to
13 ↪ AI via API)
14 const stageConfig = currentStage || stages[0]; // Internal to game-core
15 const ignoreOcclusion = stageConfig.ignoreOcclusion || false; // Internal flag
16
17 const capturedPieces = [];
18 // Logic to find and flip pieces, considering ignoreOcclusion if true
19 // The AI observes the result via simulateMove, inferring the rule.
20 for (const [dr, dc] of directions) {
21   let r = row + dr;
22   let c = col + dc;
23   const toFlip = [];
24   let foundBlocked = false;
25   let foundOpponent = false;
26
27   while (isWithinBoard(r, c)) {
28     if (board[r][c] === opponent) {
29       toFlip.push([r, c]);
30       foundOpponent = true;
31     } else if (board[r][c] === GAME_CONSTANTS.BLOCKED) {
32       foundBlocked = true;
33       if (!ignoreOcclusion) { // If occlusion is not ignored, break
34         break;
35       }
36     } else if (board[r][c] === GAME_CONSTANTS.EMPTY) {
37       break;
38     } else if (board[r][c] === player) {
39       if (foundOpponent && toFlip.length > 0 && (!foundBlocked ||
40         ↪ ignoreOcclusion)) {
41         for (const [fr, fc] of toFlip) {
42           board[fr][fc] = player;
43           capturedPieces.push([fr, fc]);
44         }
45       }
46       break;
47     }
48     r += dr;
49     c += dc;
50   }
51 }
52 // ... (rest of makeMove function, e.g., logging) ...
53 return true; // Or false if move was invalid
54 }
```

---

**Note:** The `ignoreOcclusion` flag is an internal detail of the `game-core.js` module, used by the platform to define the specific behavior of `simulateMove` and `getValidMoves` for a given stage. Intelligent systems must infer the presence of such rules through observing the API's behavior, not by direct access to `stageConfig` flags.

## B.2 Intelligent System Execution Environment

The `intelligent-system-loader.js` module is responsible for securely loading, sandboxing, and executing submitted intelligent system code. It enforces the strict  $T_{\text{analysis}}$  time limit for strategy generation and provides the limited API through which intelligent systems interact with the environment.

- **Key Responsibilities:** Code compilation (`compileSystem`), secure execution in a Web Worker, time limit enforcement, and exposing the environment API (`getValidMoves`, `simulateMove`, `evaluateBoard`) to the intelligent system.
- **Illustrative Code Snippet (Worker-based Execution):** This shows the core mechanism for running student-submitted `analyzeStage` functions in a sandboxed environment, managing time, and passing the API.

---

```
1 // Excerpt from intelligent-system-loader.js: Core analysis execution within a
  ↪ Worker
2 async analyzeStageWithSystem(systemName, stageConfig, systemCode) {
3   return new Promise((resolve, reject) => {
4     const worker = new Worker(URL.createObjectURL(new Blob([`
5       self.onmessage = function(e) {
6         const { code, stageConfig, initialBoard, validMoves, api } = e.data;
7         const startTime = performance.now();
8         const analyzeStage = new Function('stageConfig', 'initialBoard',
9           ↪ 'validMoves', 'api',
10          code + '\nreturn analyzeStage.apply(null, arguments);'
11        );
12        try {
13          const strategyFunction = analyzeStage(stageConfig, initialBoard,
14            ↪ validMoves, api);
15          self.postMessage({ status: 'success', strategyFunction:
16            ↪ strategyFunction.toString() });
17        } catch (error) {
18          self.postMessage({ status: 'error', message: error.message,
19            ↪ stack: error.stack });
20        }
21      });
22    `], { type: 'application/javascript' })));
23
24    // Time limit enforced externally (in loader, not worker code shown here)
25    const timeoutId = setTimeout(() => { worker.terminate(); reject(new
26      ↪ Error("Analysis timed out")); }, this.timeLimit);
27
28    worker.onmessage = (e) => {
29      clearTimeout(timeoutId);
30      if (e.data.status === 'success') { /* ... process strategy function ...
31        ↪ */ }
32      else { /* ... handle error ... */ }
33    };
34    worker.onerror = (error) => { clearTimeout(timeoutId); reject(error); };
35
36    worker.postMessage({ code: systemCode, stageConfig, initialBoard,
37      ↪ validMoves, api: this.api });
38  });
39 }
```

---

## B.3 Strategy Management

The `strategies.js` module manages both built-in AI strategies and custom strategies uploaded by participants. It handles the storage, retrieval, and runtime compilation of strategy code.

- **Key Responsibilities:** Defining standard Othello strategies (e.g., Random, Greedy, Positional), saving and loading custom strategy code (e.g., from `localStorage`), and providing

getCompiledStrategy to retrieve an executable function for a given strategy ID. This module serves the GameController by providing the actual move-making functions.

- **Illustrative Code Snippet (Strategy Retrieval and Compilation):** This demonstrates how the arena retrieves an executable strategy, prioritizing cached compiled functions for efficiency.

---

```
1 // Excerpt from strategies.js: Retrieving and compiling strategies
2 function getCompiledStrategy(controllerId, player) {
3     // Handle custom strategies (student-submitted)
4     if (controllerId.startsWith('custom_')) {
5         const strategyName = controllerId.replace('custom_', '');
6         // Prioritize cached compiled function (from intelligent-system-loader)
7         if (window.compiledIntelligentSystems[strategyName]) {
8             return window.compiledIntelligentSystems[strategyName]; // Function
9             ↪ stored by intelligent-system-loader
10        }
11        // Fallback: compile from stored code string if not cached
12        if (savedStrategies[strategyName]) {
13            const code = savedStrategies[strategyName];
14            try {
15                const compiledFunc = new Function('boardArg', 'playerArg',
16                    ↪ 'validMovesArg', 'makeMoveFunc',
17                    ↪ `${code}\nreturn studentStrategy(boardArg, playerArg,
18                    ↪ validMovesArg, makeMoveFunc);`);
19                compiledStudentAIs[strategyName] = compiledFunc; // Cache it
20                return compiledFunc;
21            } catch (e) { console.error(`Compile error for ${strategyName}:`, e);
22                ↪ return null; }
23        }
24    }
25    // Handle built-in strategies
26    else if (builtInStrategies[controllerId]) {
27        return builtInStrategies[controllerId]; // Directly use built-in functions
28    }
29    return null;
30 }
```

---

## B.4 Game Flow and UI Management

The game-controller.js and game-ui.js modules work in tandem to orchestrate the game flow and manage the web-based user interface.

- **GameController** (game-controller.js): Manages the overall game loop, including starting/resetting games, calling AI strategies for moves, and determining game-over conditions. It also tracks player time usage.
- **OthelloUI** (game-ui.js): Handles all aspects of the graphical user interface, including rendering the Othello board, updating scores and timers, displaying game status messages, and processing human player inputs.

## B.5 Logging and Replay System

The game-logger.js and game-rollout.js modules are dedicated to recording game history and enabling post-game analysis and visualization.

- **GameLogger** (game-logger.js): Provides a comprehensive system for recording every move, board state, player turn, and captured piece count during a game. It also stores game results for the leaderboard and supports saving data to localStorage.
- **GameRollout** (game-rollout.js): Utilizes the data from GameLogger to allow users to replay any past game, navigate through moves, and adjust playback speed. It dynamically reconstructs the board state at any point in the game.

## B.6 Tournament System

The `tournament.js` module orchestrates automated tournaments between submitted intelligent systems and built-in strategies.

- **Key Responsibilities:** Managing tournament rounds, recording match outcomes, calculating and displaying the leaderboard, and saving/loading tournament data. It provides the competitive environment for evaluating adaptive performance across private stages.

## C Othello Arena API Reference

The Othello AI Arena provides the following core API functions for developing intelligent systems, allowing them to probe and understand the environment:

- `getValidMoves(board, player)`: Returns a list of valid moves for a player on a given board state. This function dynamically applies the stage's underlying rules, meaning the intelligent system must observe its behavior to infer rule variations.
- `simulateMove(board, player, row, col)`: Simulates the outcome of making a specific move, returning the resulting board state and number of pieces captured. The behavior of this function (e.g., how pieces are captured) directly reflects the stage's rules, allowing the intelligent system to deduce capture mechanics without explicit rule disclosure.
- `evaluateBoard(board, player)`: Provides an evaluation of a given board state using pre-defined metrics (e.g., piece count, mobility, corner control). The intelligent system may need to learn or adjust the interpretation of these metrics based on observed game dynamics.

An intelligent system is implemented through the `analyzeStage` function, which must analyze the stage within a strict time limit (approximately 60 seconds) and return a tailored strategy function. The returned strategy function is then called during each turn of the game to determine a move.

---

```
1 // Basic Intelligent System Template
2 function analyzeStage(stageConfig, initialBoard, validMoves, api) {
3   // Perform stage analysis within the ~60-second limit
4   console.log("Analyzing stage: " + stageConfig.name);
5
6   // Identify static environmental features (e.g., board dimensions, presence of
7   // ↪ blocked cells)
8   const boardSize = initialBoard.length;
9   const hasBlockedCells = initialBoard.some(row =>
10     row.some(cell => cell === GAME_CONSTANTS.BLOCKED)); //
11     ↪ GAME_CONSTANTS.BLOCKED is a known constant
12
13   // The core challenge: Adapt strategy parameters based on inferred rules and
14   // ↪ board properties.
15   // This 'optimizeParameters' function (conceptual) would perform simulations and
16   // ↪ deductions.
17   const params = optimizeParameters(initialBoard, api);
18
19   // Return the strategy function.
20   // This function will be called during each game turn (cumulative ~10-second
21   // ↪ limit per game).
22   return function(board, player, validMoves) {
23     if (validMoves.length === 0) return null;
24
25     // This function decides the best move for the current turn,
26     // ↪ utilizing the parameters and understanding gained during analysis.
27     return selectBestMove(board, player, validMoves, params);
28   };
29 }
```

---

Key patterns for effectively utilizing the API include:

1. Thoroughly understanding board characteristics and inferring rule variations during the initial analysis phase through active probing via the API.
2. Efficiently using `simulateMove` for game tree exploration and evaluation function optimization, which helps in implicitly learning the game's dynamic rules.
3. Performing complex computations (like extensive simulations or rule induction) during the analysis phase and leveraging the results within the strategy function.
4. Employing memoization and optimization techniques to minimize execution time during gameplay.

### C.1 Example Intelligent System: Advanced Template

This example demonstrates a more sophisticated intelligent system that performs self-play simulations to learn about the stage, detect rules, and generate an adaptive strategy. It showcases how the provided API functions can be leveraged within the analysis time limit ( $T_{analysis}$ ) to synthesize a high-performing strategy for a novel environment.

---

```

1  /**
2   * Function that intelligent systems must implement to analyze a stage and generate
3   * ↪ a strategy
4   *
5   * @param {Object} stageConfig - Stage configuration object (contains name,
6   * ↪ boardSize, initial piece positions, but NO explicit rule flags).
7   * @param {Array<Array<number>>} initialBoard - 2D array representing the initial
8   * ↪ board state.
9   * @param {Array<Object>} validMoves - Valid moves for the first player (Black) on
10  * ↪ the initial board.
11  * @param {Object} api - Environment API with methods:
12  * - getValidMoves(board, player): Returns valid moves for a given board state.
13  * - simulateMove(board, player, row, col): Returns the resulting board and captured
14  * ↪ count for a move.
15  * - evaluateBoard(board, player): Provides evaluation metrics for a board.
16  * (These API methods internally implement the stage's TRUE rules,
17  * which the intelligent system must infer through observation).
18  *
19  * @returns {Function} Strategy function that will be called during gameplay.
20  */
21  function analyzeStage(stageConfig, initialBoard, validMoves, api) {
22    console.log("=== Starting Environmental Analysis and Strategy Adaptation ===");
23    console.log("Stage:", stageConfig.name, "Board Size:", stageConfig.boardSize);
24
25    // Setup simulation parameters for environmental probing and data collection
26    const NUM_SIM_GAMES = 3000;
27    const MAX_MOVES_PER_GAME = 60;
28    const ANALYSIS_TIME_LIMIT_MS = 55000; // Leave a few seconds buffer for final
29    ↪ strategy compilation
30
31    // Track the start time to adhere to the T_analysis limit
32    const startTime = Date.now();
33
34    // Data structures for storing observations and learned patterns
35    const gameSimulationLogs = []; // Stores detailed logs of simulated games
36    const positionWinRates = {}; // Tracks win rates for specific positions
37    const boardSize = stageConfig.boardSize;
38
39    // Initialize position win rates tracking
40    for (let r = 0; r < boardSize; r++) {
41      for (let c = 0; c < boardSize; c++) {
42        positionWinRates[`${r},${c}`] = { plays: 0, wins: 0, blackPlays: 0,
43        ↪ blackWins: 0, whitePlays: 0, whiteWins: 0 };
44      }
45    }
46  }

```

```

39
40 // Define simple, generic strategies for self-play simulations.
41 // These strategies do not have prior knowledge of rule variations,
42 // they just interact with the API's behavior.
43 const simulationStrategies = {
44   random: function (board, player, moves) {
45     if (moves.length === 0) return null;
46     return moves[Math.floor(Math.random() * moves.length)];
47   },
48   greedy: function (board, player, moves) {
49     if (moves.length === 0) return null;
50     let bestMove = moves[0];
51     let maxCaptures = -1;
52
53     for (const move of moves) {
54       const result = api.simulateMove(board, player, move.row, move.col);
55       if (result.valid && result.capturedCount > maxCaptures) {
56         maxCaptures = result.capturedCount;
57         bestMove = move;
58       }
59     }
60     return bestMove;
61   },
62   // A simple positional heuristic, which might be adapted later
63   basicPositional: function(board, player, moves) {
64     if (moves.length === 0) return null;
65     let bestMove = null;
66     let bestScore = -Infinity;
67
68     // Simple static positional values (can be adapted)
69     const staticWeights = {
70       '8': [
71         [120, -20, 20, 5, 5, 20, -20, 120],
72         [-20, -40, -5, -5, -5, -5, -40, -20],
73         [20, -5, 15, 3, 3, 15, -5, 20],
74         [5, -5, 3, 3, 3, 3, -5, 5],
75         [5, -5, 3, 3, 3, 3, -5, 5],
76         [20, -5, 15, 3, 3, 15, -5, 20],
77         [-20, -40, -5, -5, -5, -5, -40, -20],
78         [120, -20, 20, 5, 5, 20, -20, 120]
79       ],
80       '6': [
81         [50, -10, 10, 10, -10, 50],
82         [-10, -20, -5, -5, -20, -10],
83         [10, -5, 5, 5, -5, 10],
84         [10, -5, 5, 5, -5, 10],
85         [-10, -20, -5, -5, -20, -10],
86         [50, -10, 10, 10, -10, 50]
87       ]
88     };
89     const weights = staticWeights[board.length] || staticWeights['8']; //
90     ↪ Use 8x8 as default
91
92     for (const move of moves) {
93       const score = weights[move.row][move.col];
94       if (score > bestScore) {
95         bestScore = score;
96         bestMove = move;
97       }
98     }
99     return bestMove;
100   }
101 };
102 // Run self-play simulation games to gather data and observe behaviors

```

```

103 console.log(`Starting ${NUM_SIM_GAMES} simulation games (analysis time limit:
↳ ${ANALYSIS_TIME_LIMIT_MS}ms)...`);
104
105 let gamesCompleted = 0;
106 let totalMovesMade = 0;
107 let analysisTimeoutReached = false;
108
109 for (let gameNum = 0; gameNum < NUM_SIM_GAMES; gameNum++) {
110   if (Date.now() - startTime > ANALYSIS_TIME_LIMIT_MS) {
111     analysisTimeoutReached = true;
112     console.warn(`Analysis time limit reached (${ANALYSIS_TIME_LIMIT_MS}ms).
↳ Stopping simulations.`);
113     break;
114   }
115
116   const gameLog = {
117     moves: [],
118     finalBoard: null,
119     winner: null, // 1: Black, 2: White, 0: Tie
120     scores: { black: 0, white: 0 },
121     turnTransitions: [] // To observe if turns consistently alternate or
↳ not
122   };
123
124   let currentBoard = initialBoard.map(row => [...row]); // Deep copy
125   let currentPlayer = 1; // Black starts
126   let movesInCurrentGame = 0;
127   let consecutivePasses = 0;
128   let gameOver = false;
129
130   // Randomly select two simulation strategies for this game
131   const simStrategyPlayer1 =
↳ Object.values(simulationStrategies)[Math.floor(Math.random() *
↳ Object.keys(simulationStrategies).length)];
132   const simStrategyPlayer2 =
↳ Object.values(simulationStrategies)[Math.floor(Math.random() *
↳ Object.keys(simulationStrategies).length)];
133
134   while (!gameOver && movesInCurrentGame < MAX_MOVES_PER_GAME) {
135     if (Date.now() - startTime > ANALYSIS_TIME_LIMIT_MS) {
136       analysisTimeoutReached = true;
137       break;
138     }
139
140     const moves = api.getValidMoves(currentBoard, currentPlayer);
141
142     if (moves.length === 0) {
143       consecutivePasses++;
144       if (consecutivePasses >= 2) {
145         gameOver = true; // Game over if both players pass
146         break;
147       }
148       gameLog.turnTransitions.push({ from: currentPlayer, to:
↳ (currentPlayer === 1 ? 2 : 1), pass: true });
149       currentPlayer = (currentPlayer === 1 ? 2 : 1); // Switch player if
↳ current player has no moves
150       continue;
151     }
152
153     consecutivePasses = 0; // Reset pass counter
154
155     const chosenSimStrategy = (currentPlayer === 1) ? simStrategyPlayer1 :
↳ simStrategyPlayer2;
156     const chosenMove = chosenSimStrategy(currentBoard, currentPlayer,
↳ moves);

```

```

157
158     if (!chosenMove) { // Should ideally not happen if validMoves.length >
↳ 0
159         gameOver = true;
160         break;
161     }
162
163     const simResult = api.simulateMove(currentBoard, currentPlayer,
↳ chosenMove.row, chosenMove.col);
164
165     if (simResult.valid) {
166         const prevPlayer = currentPlayer;
167         currentBoard = simResult.resultingBoard; // Update board
168         gameLog.moves.push({ player: prevPlayer, position: chosenMove,
↳ capturedCount: simResult.capturedCount });
169
170         // Update position win rates data
171         const posKey = `${chosenMove.row},${chosenMove.col}`;
172         if (positionWinRates[posKey]) {
173             positionWinRates[posKey].plays++;
174             if (prevPlayer === 1) {
175                 positionWinRates[posKey].blackPlays++;
176             } else {
177                 positionWinRates[posKey].whitePlays++;
178             }
179         }
180
181         // Determine next player based on observed behavior (this is where
↳ inference matters)
182         const scores = countPieces(currentBoard); // Utility function below
183         let nextPlayerCandidate = (prevPlayer === 1) ? 2 : 1; // Standard
↳ alternate
184
185         // Attempt to infer 'fewer pieces continue' type rule by observing
↳ actual turn transitions
186         // We're looking for patterns where the current player plays again
187         const potentialValidMovesForNextPlayer =
↳ api.getValidMoves(currentBoard, nextPlayerCandidate);
188         if (potentialValidMovesForNextPlayer.length === 0) { // If next
↳ player cannot move, current player might continue or game ends
189             // This simple check doesn't infer complex rules, but models
↳ standard game flow.
190             // A more advanced system would track actual 'next player'
↳ behavior based on api.getValidMoves and piece counts.
191             // For instance, if the game *does* switch players, but then
↳ switches back quickly when piece counts are skewed.
192         }
193
194         // If the game rules cause the same player to move again, the API's
↳ getValidMoves and simulateMove
195         // will reflect that. The AI observes.
196         // For this generic template, we simply switch player. If the actual
↳ game logic allows consecutive turns,
197         // the `api.getValidMoves` for the *next* player would return an
↳ empty array if they were forced to pass,
198         // and the *current* player would get another chance.
199         // The key is that the AI *doesn't know* why it's happening, only
↳ *that* it's happening.
200         currentPlayer = nextPlayerCandidate; // Proceed with standard turn
↳ alternation for simulation logic
201
202         gameLog.turnTransitions.push({ from: prevPlayer, to: currentPlayer,
↳ pass: false, scoresAfterMove: scores });
203         movesInCurrentGame++;
204         totalMovesMade++;

```

```

205         } else {
206             console.warn("Invalid move in simulation (should not happen for
                ↪ valid moves)");
207             gameOver = true;
208             break;
209         }
210     }
211
212     // Game finished - record final state and winner
213     const finalScores = countPieces(currentBoard);
214     gameLog.finalBoard = currentBoard;
215     gameLog.scores = finalScores;
216
217     if (finalScores.black > finalScores.white) { gameLog.winner = 1; }
218     else if (finalScores.white > finalScores.black) { gameLog.winner = 2; }
219     else { gameLog.winner = 0; } // Tie
220
221     // Update position win rates based on game winner
222     for (const move of gameLog.moves) {
223         const posKey = `${move.position.row},${move.position.col}`;
224         const posStats = positionWinRates[posKey];
225         if (posStats && gameLog.winner === move.player) {
226             posStats.wins++;
227             if (move.player === 1) { posStats.blackWins++; } else {
                ↪ posStats.whiteWins++; }
228         }
229     }
230
231     gameSimulationLogs.push(gameLog);
232     gamesCompleted++;
233 }
234
235 console.log(`Completed ${gamesCompleted} games with ${totalMovesMade} total
    ↪ moves in ${Date.now() - startTime}ms`);
236
237 // --- Phase 2: Environmental Behavior Analysis and Rule Inference ---
238 // Instead of looking for specific rule names, we analyze observed behaviors for
    ↪ anomalies.
239 const observedBehaviors = analyzeEnvironmentBehaviors(gameSimulationLogs,
    ↪ initialBoard, api);
240 console.log("Observed Environmental Behaviors:", observedBehaviors);
241
242 // --- Phase 3: Adaptive Strategy Synthesis ---
243 // Generate a position value matrix based on observed win rates
244 const adaptedPositionValueMatrix = generatePositionValueMatrix(positionWinRates,
    ↪ boardSize, observedBehaviors);
245
246 // Create an opening book from successful simulated games
247 const derivedOpeningBook = buildOpeningBook(gameSimulationLogs);
248
249 // Return the final, adapted strategy function
250 return createStrategyFunction(adaptedPositionValueMatrix, derivedOpeningBook,
    ↪ observedBehaviors, boardSize, api);
251
252 // ===== Utility Functions (Internal to analyzeStage)
    ↪ =====
253
254 // Counts pieces on a given board state
255 function countPieces(board) {
256     let black = 0, white = 0, empty = 0, blocked = 0;
257     for (let r = 0; r < board.length; r++) {
258         for (let c = 0; c < board[r].length; c++) {
259             if (board[r][c] === 1) black++;
260             else if (board[r][c] === 2) white++;
261             else if (board[r][c] === 0) empty++;

```

```

262         else if (board[r][c] === 3) blocked++;
263     }
264 }
265 return { black, white, empty, blocked };
266 }
267
268 /**
269  * Analyzes simulation logs to infer environmental behaviors (rules)
270  * without knowing specific rule names beforehand.
271  */
272 function analyzeEnvironmentBehaviors(simLogs, initialBoard, api) {
273     const behaviors = {
274         // High-level observations of turn dynamics
275         consecutiveTurnObserved: false,
276         consecutiveTurnFrequency: 0,
277         consecutiveTurnPlayerBias: { black: 0, white: 0 }, // If one player gets
                ↪ more consecutive turns
278
279         // High-level observations of capture mechanics
280         captureThroughBlockedCellsObserved: false, // Did a capture happen over
                ↪ a blocked cell?
281         unusualCapturePatternObserved: false, // More general, did captures seem
                ↪ off?
282
283         // Observations about winning conditions
284         winConditionReversedObserved: false, // Did lowest score win
                ↪ unexpectedly?
285         // ... add more abstract behavioral observations
286     };
287
288     let totalTurnTransitions = 0;
289     let consecutiveTurnsCount = 0;
290     let blackConsecutiveTurns = 0;
291     let whiteConsecutiveTurns = 0;
292
293     for (const game of simLogs) {
294         if (game.moves.length < 2) continue;
295
296         // Analyze turn dynamics
297         for (let i = 1; i < game.moves.length; i++) {
298             const prevPlayer = game.moves[i - 1].player;
299             const currPlayer = game.moves[i].player;
300             totalTurnTransitions++;
301             if (prevPlayer === currPlayer) {
302                 consecutiveTurnsCount++;
303                 if (prevPlayer === 1) blackConsecutiveTurns++;
304                 else whiteConsecutiveTurns++;
305             }
306         }
307
308         // Analyze capture mechanics (basic heuristic)
309         // Look for any capture that crosses a blocked cell in the final board
                ↪ state
310         // This is still a simple heuristic; robust detection would need more
                ↪ targeted probes.
311         for (const move of game.moves) {
312             const startR = move.position.row;
313             const startC = move.position.col;
314             const player = move.player;
315             const opponent = (player === 1) ? 2 : 1;
316
317             // Re-simulate this move on the board *before* the move to check
                ↪ intermediate states
318             // This is computationally intensive. A simpler approach is to look
                ↪ at final board and piece positions.

```

```

319         // However, a real intelligent system would use API to actively
320         ↪ probe
321     // Create a board state before this move occurred if possible from
322     ↪ logs, or approximate.
323
324     // A more direct probing strategy:
325     // Construct a small, controlled scenario to test a specific
326     ↪ behavior pattern.
327     // This is hard to do generically without knowing board structure.
328     // e.g., if (initialBoard has blocked cells) { try to place a piece,
329     ↪ then simulate. }
330     // For instance: pick a random empty cell, place current player
331     ↪ piece.
332     // Does it flip pieces beyond a blocked cell?
333     // This would be done with direct calls to api.simulateMove on
334     ↪ crafted micro-boards.
335
336     // Example: Try to detect 'capture through blocked cells' by
337     ↪ crafting a minimal board and probing
338     // This would need to happen at the very beginning of analyzeStage,
339     ↪ before main loop.
340     // This template will stick to observations from general self-play
341     ↪ for simplicity.
342     // A very simple check (less robust): Check if any flipped piece in
343     ↪ the log has a blocked cell between it and the origin move.
344     // This is difficult without reconstructing the exact path and board
345     ↪ state for each capture.
346 }
347
348 // Analyze winning condition
349 // If the game winner is the player with FEWER pieces, this suggests a
350 ↪ reversed winning condition
351 const finalScores = game.scores;
352 if (finalScores.black > finalScores.white && game.winner === 2) { //
353     ↪ Black has more pieces, but White wins
354     behaviors.winConditionReversedObserved = true;
355 } else if (finalScores.white > finalScores.black && game.winner === 1) {
356     ↪ // White has more pieces, but Black wins
357     behaviors.winConditionReversedObserved = true;
358 }
359 }
360
361 if (totalTurnTransitions > 0) {
362     behaviors.consecutiveTurnFrequency = consecutiveTurnsCount /
363     ↪ totalTurnTransitions;
364     if (behaviors.consecutiveTurnFrequency > 0.05) { // Threshold for
365     ↪ significance
366         behaviors.consecutiveTurnObserved = true;
367         if (blackConsecutiveTurns > whiteConsecutiveTurns * 1.5) { //
368         ↪ Significant bias towards black
369             behaviors.consecutiveTurnPlayerBias.black =
370             ↪ blackConsecutiveTurns / consecutiveTurnsCount;
371         } else if (whiteConsecutiveTurns > blackConsecutiveTurns * 1.5) { //
372         ↪ Significant bias towards white
373             behaviors.consecutiveTurnPlayerBias.white =
374             ↪ whiteConsecutiveTurns / consecutiveTurnsCount;
375         }
376     }
377 }
378
379 // Example: Probe for ignoreOcclusion at analysis start with a crafted
380 ↪ board
381 // This is a more direct test rather than observing during self-play.
382 // It's still a heuristic, not guaranteed to find all types of hidden
383 ↪ rules.

```

```

362     const probeBoard = [
363         [0,0,0,0,0,0,0,0],
364         [0,0,0,0,0,0,0,0],
365         [0,0,2,3,1,0,0,0], // Opponent, Blocked, Player (try to capture over
           ↪ blocked)
366         [0,0,0,0,0,0,0,0],
367         [0,0,0,0,0,0,0,0],
368         [0,0,0,0,0,0,0,0],
369         [0,0,0,0,0,0,0,0],
370         [0,0,0,0,0,0,0,0]
371     ];
372     // Scale probe board to current board size if needed, or create a generic
           ↪ smaller one.
373     // Assume probe board is compatible with boardSize or it's a fixed small
           ↪ test.
374
375     // Simulate placing a piece that should capture over the blocked cell (if
           ↪ rules allow)
376     // If current player is 1 (Black), they want to capture White (2) over
           ↪ Blocked (3).
377     const simulatedProbeResult = api.simulateMove(probeBoard, 1, 2, 0); // Try
           ↪ to capture from (2,0) over (2,1),(2,2) to (2,4)
378     if (simulatedProbeResult.valid && simulatedProbeResult.capturedCount > 0) {
379         // If a capture occurred, and there's a blocked cell in the line
           ↪ between the start and end piece,
380         // then it's likely 'ignoreOcclusion' or similar.
381         // This needs more robust check: is there a blocked cell between the
           ↪ origin and the captured pieces?
382         // A very simple proxy: if the move *would* be invalid in standard
           ↪ Othello due to a block, but here it's valid.
383         // This requires knowing standard Othello rules internally.
384         // For simplicity, we just mark it as "unusual capture pattern" if the
           ↪ probe yields unexpected captures.
385         behaviors.captureThroughBlockedCellsObserved = true; // Still a
           ↪ simplification for the template
386     }
387
388
389     return behaviors;
390 }
391
392
393 // Generates a position value matrix, potentially adapted by observed behaviors
394 function generatePositionValueMatrix(posRates, boardSize, behaviors) {
395     const matrix = Array(boardSize).fill().map(() => Array(boardSize).fill(0));
396
397     // Start with a generic positional weight matrix (can be a standard Othello
           ↪ heuristic)
398     const baseWeights_8x8 = [
399         [120, -20, 20, 5, 5, 20, -20, 120],
400         [-20, -40, -5, -5, -5, -5, -40, -20],
401         [20, -5, 15, 3, 3, 15, -5, 20],
402         [5, -5, 3, 3, 3, 3, -5, 5],
403         [5, -5, 3, 3, 3, 3, -5, 5],
404         [20, -5, 15, 3, 3, 15, -5, 20],
405         [-20, -40, -5, -5, -5, -5, -40, -20],
406         [120, -20, 20, 5, 5, 20, -20, 120]
407     ];
408     const baseWeights_6x6 = [
409         [50, -10, 10, 10, -10, 50],
410         [-10, -20, -5, -5, -20, -10],
411         [10, -5, 5, 5, -5, 10],
412         [10, -5, 5, 5, -5, 10],
413         [-10, -20, -5, -5, -20, -10],
414         [50, -10, 10, 10, -10, 50]

```

```

415 ];
416 const baseWeights = (boardSize === 6) ? baseWeights_6x6 : baseWeights_8x8;
417
418 for (let r = 0; r < boardSize; r++) {
419   for (let c = 0; c < boardSize; c++) {
420     matrix[r][c] = baseWeights[r] ? baseWeights[r][c] : 0; // Handle
421     ↪ different board sizes if needed
422   }
423 }
424
425 // Adjust based on observed win rates from simulations
426 for (let r = 0; r < boardSize; r++) {
427   for (let c = 0; c < boardSize; c++) {
428     const posKey = `${r},${c}`;
429     const posStats = posRates[posKey];
430
431     if (posStats && posStats.plays >= 3) { // Require minimum plays for
432     ↪ confidence
433       const confidence = Math.min(1, posStats.plays / 10);
434       const winRateAdjustment = (posStats.wins / posStats.plays - 0.5)
435       ↪ * 100 * confidence;
436       matrix[r][c] = Math.round(matrix[r][c] * 0.7 + winRateAdjustment
437       ↪ * 0.3); // Blend with base
438     }
439   }
440 }
441
442 // Adapt matrix based on observed behaviors
443 if (behaviors.winConditionReversedObserved) {
444   console.log("Adapting position matrix for reversed winning condition.");
445   // Invert the values to prioritize positions that lead to fewer pieces
446   for (let r = 0; r < boardSize; r++) {
447     for (let c = 0; c < boardSize; c++) {
448       matrix[r][c] = -matrix[r][c];
449     }
450   }
451 }
452 // Add more adaptations based on other behaviors...
453
454 return matrix;
455 }
456
457 // Builds an opening book based on successful simulated games
458 function buildOpeningBook(simLogs) {
459   const openings = {};
460
461   simLogs.forEach(game => {
462     if (game.moves.length < 4) return; // Opening sequence too short
463
464     const openingMoves = game.moves.slice(0, Math.min(6,
465     ↪ game.moves.length));
466     const openingKey = openingMoves.map(m =>
467     ↪ `${m.player}:${m.position.row},${m.position.col}`).join('|');
468
469     if (!openings[openingKey]) {
470       openings[openingKey] = { sequence: openingMoves, wins: 0, losses: 0,
471       ↪ games: 0 };
472     }
473     openings[openingKey].games++;
474     if (game.winner === 1) openings[openingKey].wins++;
475     else if (game.winner === 2) openings[openingKey].losses++;
476   });
477
478 // Select the best openings (e.g., top 3 by win rate, with minimum plays)
479 const goodOpenings = Object.values(openings)

```

```

473         .filter(o => o.games >= 2 && (o.wins / o.games) > 0.6)
474         .sort((a, b) => (b.wins / b.games) - (a.wins / a.games))
475         .slice(0, 3);
476
477     return goodOpenings.map(o => ({
478         sequence: o.sequence.map(m => ({ player: m.player, position: { row:
479             ↪ m.position.row, col: m.position.col } })),
480         winRate: o.games > 0 ? o.wins / o.games : 0
481     }));
482 }
483
484 /**
485  * Creates the final strategy function that the Othello Arena will execute
486  * ↪ during gameplay.
487  * This function should be fast and rely on pre-computed data from the analysis
488  * ↪ phase.
489  */
490 function createStrategyFunction(positionValues, openingBook, behaviors,
491 ↪ boardSize, api) {
492     console.log("\n=== Final Strategy Function Created ===");
493     console.log(`Leveraging ${openingBook.length} opening sequences and adapted
494     ↪ position values.`);
495     console.log("Adapting based on observed behaviors:", behaviors);
496
497     return function (board, player, validMoves) {
498         // Handle no valid moves (should be handled by GameController for
499         ↪ passes)
500         if (validMoves.length === 0) return null;
501
502         // --- Phase 1: Opening Book (if applicable) ---
503         const pieceCount = countPieces(board).black + countPieces(board).white;
504         if (pieceCount < 12 && openingBook.length > 0) { // Still early game
505             // Attempt to follow an opening sequence from the derived book
506             for (const opening of openingBook) {
507                 // This matching logic is simplified; a real system would hash
508                 ↪ board states
509                 // or track sequence of moves to find exact matches.
510                 // For now, it just checks if the *next move in an opening* is
511                 ↪ valid.
512                 for (const moveData of opening.sequence) {
513                     if (moveData.player === player) {
514                         const openingMove = { row: moveData.position.row, col:
515                             ↪ moveData.position.col };
516                         if (validMoves.some(m => m.row === openingMove.row &&
517                             ↪ m.col === openingMove.col)) {
518                             return openingMove;
519                         }
520                     }
521                 }
522             }
523         }
524
525         // --- Phase 2: Positional Evaluation with Lookahead ---
526         let bestMove = null;
527         let bestScore = -Infinity;
528
529         for (const move of validMoves) {
530             // Start with the position's base value from the adapted matrix
531             let score = positionValues[move.row][move.col];
532
533             // Simulate the move using the API to observe its direct impact
534             const result = api.simulateMove(board, player, move.row, move.col);
535
536             if (result.valid) {
537                 // Basic capture bonus

```

```

528     score += result.capturedCount * 5;
529
530     // Evaluate the resulting board state using the API's evaluator
531     ↪ (reflects true rules)
532     const evaluationAfterMove =
533     ↪ api.evaluateBoard(result.resultingBoard, player);
534
535     // Incorporate mobility score, adapted for potential reversed
536     ↪ win condition
537     let mobilityContribution = evaluationAfterMove.mobilityScore *
538     ↪ 10;
539     if (behaviors.winConditionReversedObserved) {
540         // If reversed, mobility for opponent might be good for us
541         ↪ (fewer pieces)
542         // This is a complex adaptation; for now, maybe reduce
543         ↪ mobility focus or even invert.
544         // A simpler adaptation: focus on reducing opponent's moves
545         ↪ if reversed
546         mobilityContribution = -evaluationAfterMove.mobilityScore *
547         ↪ 5; // Less mobility is good
548     }
549     score += mobilityContribution;
550
551     // Add a bonus for corners (if applicable, values from
552     ↪ positionValues already handle this)
553     // (Corner values are usually high in positionValues, so direct
554     ↪ addition might be redundant)
555
556     // Simple 1-ply minimax-like lookahead using API's
557     ↪ evaluateBoard
558     // Simulate opponent's best response to this move
559     const opponent = (player === 1) ? 2 : 1;
560     const opponentValidMoves =
561     ↪ api.getValidMoves(result.resultingBoard, opponent);
562
563     if (opponentValidMoves.length > 0) {
564         let worstCaseOpponentScore = Infinity; // Opponent tries to
565         ↪ maximize their score
566
567         for (const oppMove of opponentValidMoves) {
568             const oppSimResult =
569             ↪ api.simulateMove(result.resultingBoard, opponent,
570             ↪ oppMove.row, oppMove.col);
571             if (oppSimResult.valid) {
572                 const oppEvaluation =
573                 ↪ api.evaluateBoard(oppSimResult.resultingBoard,
574                 ↪ player).totalScore;
575                 // We want to minimize opponent's gain, so we look
576                 ↪ for the worst-case for us
577                 if (oppEvaluation < worstCaseOpponentScore) {
578                     worstCaseOpponentScore = oppEvaluation;
579                 }
580             }
581         }
582         // Penalize our score by opponent's potential worst-case
583         ↪ (for us) outcome
584         score += worstCaseOpponentScore * -0.5; // Multiply by a
585         ↪ negative factor
586     } else {
587         // If opponent has no moves after our move, it's generally
588         ↪ very good!
589         score += 100;
590     }
591 }
592
593 }
594
595 }
596
597 }
598
599 }
600
601 }

```

```

572         if (score > bestScore) {
573             bestScore = score;
574             bestMove = move;
575         }
576     }
577 }
578 return bestMove;
579 };
580 }
581 }

```

---

## D Board Variation Types

The Othello AI Arena introduces a diverse set of environmental variations to rigorously evaluate the adaptive and generalization capabilities of intelligent systems. These variations, which remain unseen to the AI until the analysis phase, probe different cognitive aspects crucial for artificial general intelligence. They are broadly categorized as follows:

### D.1 Structural Variations

Structural variations modify the physical layout of the Othello board, challenging an AI's spatial reasoning, pathfinding, and dynamic re-evaluation of positional values.

- **Board Size Alterations:** The game board deviates from the standard  $8 \times 8$  grid (Stage 1), including smaller (e.g.,  $6 \times 6$ , Stage 2) or larger dimensions. This tests an AI's ability to generalize its search strategies, positional evaluations, and heuristics to new scales.
- **Blocked Cells:** Impassable cells are introduced on the board, fundamentally altering connectivity and control areas. This necessitates dynamic re-evaluation of move validity, region control, and strategic pathfinding in a constrained space. An example of such a variation is the " $8 \times 8$  (Partial C-Squares-cw)" stage (Stage 3), where specific cells are blocked.

Figure 4 illustrates some fundamental structural variations designed within our framework. Specifically, the "C-square" variation (Figure 4c) introduces blocked cells that significantly impact movement patterns and strategic value maps.

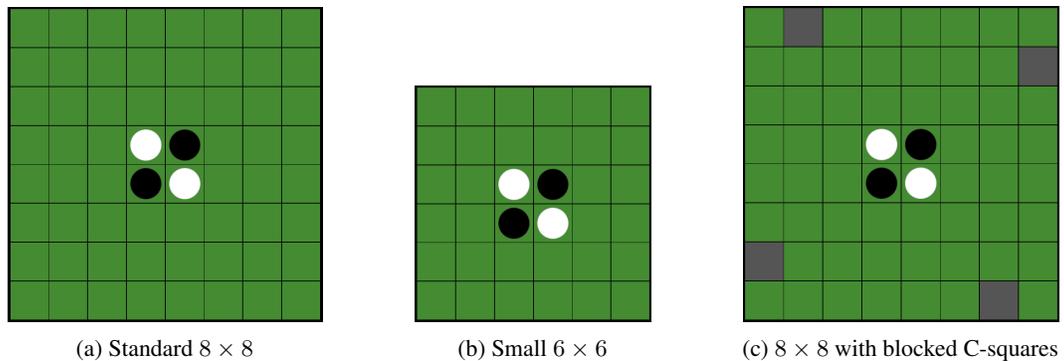


Figure 4: Basic spatial board variations with different board sizes and blocked cell configurations.

### D.2 Rule Variations

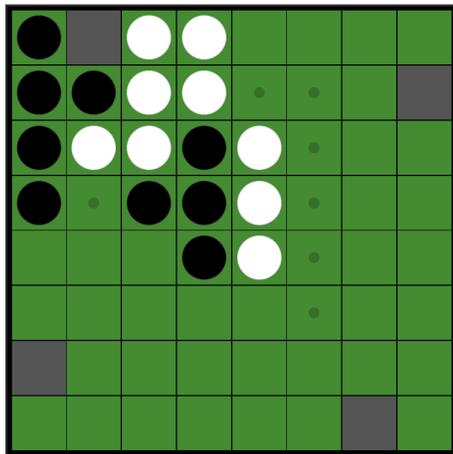
Rule variations fundamentally alter the game mechanics, demanding rapid rule induction and flexible adaptation of strategic thinking. These are particularly challenging as the AI must infer the very laws governing the environment.

- **Capture Mechanics Modifications:** The core "sandwich" rule for capturing pieces can be altered. For instance, in an 'ignore occlusion' variant, blocked cells do not halt capture lines, allowing

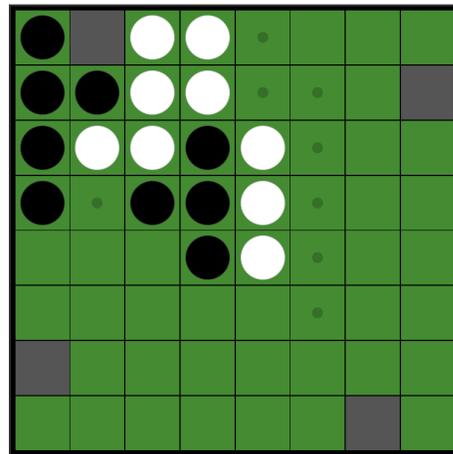
pieces to be flipped over obstacles. This requires an AI to quickly update its move simulation and evaluation logic based on observed outcomes. A notable example is the "8 × 8 (C-Squares Occlusion Agnostic)" stage (Stage 12), which features this rule in our benchmark.

- **Turn Dynamics Alterations:** The rules governing whose turn it is can change. An example is the 'fewer pieces continue' rule (Stage 13), where the player with fewer pieces takes consecutive turns, profoundly affecting temporal planning and requiring an understanding of game flow based on piece counts.
- **Winning Conditions Changes:** The objective of the game can be modified. For example, in "Reverse Othello," the player with the *least* pieces at the end wins, directly challenging goal reorientation and counter-intuitive strategic thinking.
- **Move Restrictions:** Certain conditions might impose limitations on move validity, requiring dynamic re-evaluation of the action space.

Figure 5 visually demonstrates how a subtle rule change, like 'ignore occlusion,' impacts gameplay and player perception. The intelligent system must discern this alteration through interaction with the API's `simulateMove` function.



(a) Standard Capture Mechanics (as typically found in Stage 3): For instance, placing a piece at e1 would not capture pieces beyond b1, as b1 is a blocked cell. This illustrates the typical logic where blocked cells halt capture lines.



(b) Occlusion-Agnostic Capture (Stage 12): Placing a piece at e1 *can* capture pieces (like c1, d1) even across a blocked cell (b1). An intelligent system must infer this altered capture mechanic by observing the API's behavior.

Figure 5: Advanced spatial and rule variations demonstrating complex environmental modifications. The AI must adapt its strategy based on inferred rule changes.

### D.3 Initial State Variations

These variations present non-standard starting configurations, which invalidate reliance on fixed opening strategies and demand dynamic early-game adaptation.

- **Non-standard Initial Placement:** Pieces may begin in unusual positions, requiring the AI to dynamically generate early-game strategies instead of using predefined opening books.
- **Pre-placed Special Pieces:** (Potential future extension) Introducing pieces with unique properties or interactions from the outset.
- **Randomized Initialization:** Some initial board states might include random elements, adding to the novelty.

Table 1 summarizes the impact of each variation type on game dynamics and required strategic adaptations.

Table 1: Impact of variation types on game dynamics and strategic adaptation.

Variation Type	Impact on Game Dynamics	Strategic Implications
Structural Variations	Changes move patterns, accessibility of specific areas	Re-evaluation of positional values, shifts in area control strategies
Rule Variations	Alters capture mechanisms, modifies turn dynamics	Changes risk-reward balance, requires adaptation of tempo strategies
Initial State Variations	Modifies early game dynamics, affects opening patterns	Invalidates standard opening strategies, necessitates adaptive early-game play

## E Time Management Strategies in Intelligent System Development

Participants develop intelligent systems for this benchmark must strategically allocate computational resources across three distinct temporal phases, each characterized by different constraints and optimization objectives.

### E.1 Leveraging $T_{\text{development}}$ : The Extended Preparation Phase

The development period  $T_{\text{development}} \approx 30$  days provides participants with substantial computational budget to design and optimize their intelligent systems before deployment.

**Foundation Building and API Mastery** During  $T_{\text{development}}$ , participants must achieve comprehensive understanding of the arena architecture and API functionality through systematic experimentation on public stages  $S_{\text{public}} = \{s_1, s_2, s_3\}$ . The optimization objective is to develop an intelligent system  $I$  such that the generated strategies  $f_{s_i} = I(s_i, \text{API})$  outperform baseline strategies across all  $s_i \in S_{\text{public}}$ . This process enables participants to infer the benchmark’s underlying design principles and adaptation requirements.

**Generalization vs. Overfitting Trade-off** While participants have access to  $|S_{\text{public}}| = 3$  training stages and knowledge that evaluation will occur on  $|S_{\text{private}}| = n$  hidden stages, they are encouraged to hypothesize potential configurations and construct augmented stages  $S_{\text{aug}}$  for validation. However, the development approach must avoid creating systems that memorize solutions across  $S_{\text{public}} \cup S_{\text{aug}}$ , as this constitutes data interpolation rather than genuine adaptation capability. The fundamental challenge requires developing adaptation mechanisms that generalize effectively when  $S_{\text{train}} \cap S_{\text{test}} = \emptyset$ .

**Temporal Resource Allocation Optimization** A critical aspect of  $T_{\text{development}}$  involves optimizing the allocation strategy for  $T_{\text{analysis}}$ . Let  $\alpha, \beta, \gamma$  represent the proportion of analysis time dedicated to environment discovery, strategy adaptation, and parameter optimization respectively, where  $\alpha + \beta + \gamma = 1$ . An exemplary allocation might set  $\alpha = 0.2$  for winning condition and dynamics identification,  $\beta = 0.3$  for adapting existing high-performance strategies, and  $\gamma = 0.5$  for self-play-based parameter refinement among top-performing agents. Since optimizing  $(\alpha, \beta, \gamma)$  requires extensive experimentation and cannot be performed within the constrained  $T_{\text{analysis}}$  window, these hyperparameters must be determined during  $T_{\text{development}}$ .

Similarly, participants must predetermine  $T_{\text{game}}$  utilization strategies. Given that typical 8x8 Othello games require approximately 30 moves for completion, the challenge becomes distributing the total budget  $T_{\text{game}} \approx 10$  seconds across individual move decisions  $t_i$  such that  $\sum_{i=1}^{30} t_i \leq T_{\text{game}}$  while maximizing strategic performance.

### E.2 Optimizing $T_{\text{analysis}}$ : Rapid Environment Modeling

The analysis phase  $T_{\text{analysis}} \approx 60$  seconds requires efficient algorithms for environment discovery and strategy synthesis when presented with novel stage configuration  $s \in S_{\text{private}}$ .

**Environment Dynamics Identification** Upon encountering stage  $s$  with unknown rule variations, the intelligent system must rapidly identify the winning condition  $W(s)$  and transition dynamics  $\mathcal{T}(s)$ .

This process leverages the provided API functions: `getValidMoves( $b, p$ )`, `simulateMove( $b, p, r, c$ )`, and `evaluateBoard( $b, p$ )`, where  $b$  represents board state,  $p$  denotes player, and  $(r, c)$  specifies position coordinates. The system must systematically explore the action space  $\mathcal{A}(b, p)$  to construct an implicit model of environment dynamics within the temporal constraint.

**Strategic Hypothesis Formation** Following environment model construction, the system must synthesize strategy function  $f_s : (\mathcal{B} \times \mathcal{P} \times \mathcal{A}) \rightarrow \mathcal{A}$  that maps board states, player identifiers, and valid action sets to selected actions. This synthesis process typically involves rapid simulation-based learning, where the system executes approximately 3000 self-play games using the API to gather statistical evidence for position evaluation and opening sequence optimization.

### E.3 Managing $T_{\text{game}}$ : Constrained Real-Time Execution

The game execution phase imposes the constraint  $\sum_{\text{moves}} \text{Time}(f_s(b, p, M)) \leq T_{\text{game}} \approx 10$  seconds, requiring strategic time distribution analogous to bullet chess optimization [2].

**Theoretical Foundations** World-championship Othello programs demonstrate efficient search budget allocation under extreme temporal constraints [8]. Progressive simulation techniques applied to perfect-information games optimize computational resources through adaptive budget adjustment [9].

**Optimal Time Allocation Strategies** Empirically validated time management follows structured distribution patterns:

- **Phase-dependent Allocation:** Let  $t_{\text{opening}}, t_{\text{midgame}}, t_{\text{endgame}}$  denote per-move time budgets for different game phases:
  - Opening:  $t_{\text{opening}} \in [0.5, 1.0]$  seconds for rapid development
  - Midgame:  $t_{\text{midgame}} \in [1.0, 2.0]$  seconds with higher allocation for tactical critical points
  - Endgame:  $t_{\text{endgame}} < 1.0$  second based on remaining search space
- **Premove Optimization:** Pre-computation of likely responses to minimize execution latency while maintaining adaptability to unexpected opponent moves
- **Heuristic Approximation:** Replacement of computationally expensive search procedures with rapid pattern-based evaluation focusing on corner control and edge stability
- **Memoization Strategies:** Caching of pre-computed positional value matrices and tactical sequences during  $T_{\text{analysis}}$  for  $O(1)$  lookup during  $T_{\text{game}}$

## F Game Logging and Data Format

The Othello AI Arena provides a comprehensive game logging system that records the details of each game played, including both standard matches and tournament games. This logging functionality is crucial for post-hoc analysis of adaptation strategies, debugging intelligent systems, and building valuable datasets for future research.

The platform allows users to save game logs in two primary formats: a human-readable text format and a machine-readable JSON format. This can be done via a "Save Log" function available in the user interface after a game or tournament is completed. These logs capture the full sequence of events and states, enabling participants and researchers to reconstruct games and analyze player behaviors and environmental interactions.

### F.1 Human-Readable Text Log

The text log provides a simple, turn-by-turn record of the game in a format easily understandable by humans. It includes basic information about the game setup (strategies, stage) and the sequence of moves made by each player, along with the final score and winner. This format is particularly useful for quickly reviewing game flow and identifying specific sequences of moves.

A representative snippet of the human-readable text log format is shown below:

---

```
1  === Game 1 ===
2  Game started: Corners(B) vs Greedy(W) on Stage: Standard 8x8
3  Corners(B): d3
4  Greedy(W): c3
5  Corners(B): b3
6  Greedy(W): b2
7  Corners(B): b1
8  Greedy(W): e3
9  Corners(B): f3
10 Greedy(W): a1
11 Corners(B): c4
12 Greedy(W): g3
13 Corners(B): h3
14 Greedy(W): e2
15 Corners(B): d1
16 Greedy(W): a3
17 Corners(B): a2
18 Greedy(W): b4
19 Corners(B): a4
20 Greedy(W): a5
21 Corners(B): b5
22 Greedy(W): c5
23 Corners(B): a6
24 Greedy(W): c2
25 Corners(B): c1
26 Greedy(W): e1
27 Corners(B): f1
28 Greedy(W): d2
29 Corners(B): e6
30 Greedy(W): e7
31 Corners(B): f4
32 Greedy(W): g4
33 Corners(B): g2
34 Greedy(W): g1
35 Corners(B): h1
36 Greedy(W): f2
37 Corners(B): e8
38 Greedy(W): f5
39 Corners(B): g5
40 Greedy(W): h5
41 Corners(B): c6
42 Greedy(W): h2
43 Corners(B): h4
44 Greedy(W): b7
45 Corners(B): a8
46 Greedy(W): d7
47 Corners(B): h6
48 Greedy(W): h7
49 Corners(B): h8
50 Greedy(W): g6
51 Corners(B): c8
52 Greedy(W): f6
53 Corners(B): d6
54 Greedy(W): d8
55 Corners(B): f7
56 Greedy(W): f8
57 Corners(B): g8
58 Greedy(W): g7
59 Corners(B): c7
60 Greedy(W): b8
61 Corners(B): b6
62 Greedy(W): a7
63 Game over: Final score 27-37
```

```
64 White wins!  
65 % ... (additional games if multiple are logged)
```

---

## F.2 Machine-Readable JSON Log

The JSON log provides a structured, machine-readable representation of the game data. This format is ideal for automated analysis, data processing, and integration with external tools for in-depth research. It contains detailed information for each game, including metadata, the initial board state, a sequence of moves with associated board states after each move, captured piece counts, time usage, and potentially analysis-specific data from intelligent systems.

A simplified structure and a representative snippet of the JSON log format are shown below. The full JSON file contains entries for each move, detailing the board state, player, position, captured pieces, and time taken for that move.

---

```
1  [  
2  {  
3    "metadata": {  
4      "timestamp": "2025-05-16T07:04:50.912Z",  
5      "stageId": "stage-0",  
6      "stageName": "Standard 8x8",  
7      "blackStrategy": "Corners",  
8      "whiteStrategy": "Greedy",  
9      "blackScore": 27,  
10     "whiteScore": 37,  
11     "winner": 2, // 1 for Black, 2 for White, 0 for Tie  
12     "gameLength": 60 // Number of moves  
13   },  
14   "initialBoard": [  
15     // 2D array representing the board (0: Empty, 1: Black, 2: White, 3: Blocked)  
16     [0, 0, 0, 0, 0, 0, 0, 0],  
17     // ... more rows  
18     [0, 0, 0, 0, 0, 0, 0, 0]  
19   ],  
20   "moves": [  
21     {  
22       "player": 1, // 1 for Black, 2 for White  
23       "position": {"row": 2, "col": 3},  
24       "capturedCount": 1,  
25       "timeSpent": 127, // Time taken for this move in milliseconds (example  
26       ↪ field)  
27       "boardAfter": [  
28         // 2D array of the board state after this move  
29         [0, 0, 0, 0, 0, 0, 0, 0],  
30         // ... more rows  
31         [0, 0, 0, 0, 0, 0, 0, 0]  
32       ]  
33     },  
34     // ... additional move objects  
35   ],  
36   "analysisData": { // Optional: Data generated by the intelligent system during  
37     ↪ analysis  
38     "blackSystem": {  
39       "analysisTime": 58432, // Total analysis time in milliseconds  
40       "exploredPositions": 2743, // Example analysis metric  
41       // ... other analysis-specific data  
42     }  
43   }  
44 ]  
45 ]
```

---

The availability of both detailed, machine-readable JSON data and easily reviewable text logs significantly enhances the research utility and educational value of the Othello AI Arena. Researchers can use the JSON data for large-scale quantitative analysis and model training, while participants can use the text logs and visual replay (discussed in Section 2.1) to understand game flow and debug their strategies.

## G Human Adaptation to Novel Board Games: Observations

To better understand the benchmark’s relevance to human intelligence, I analyzed how humans adapt to novel board games using our time constraint framework, by observing an experienced board game player (one of our graduate student with chess ELO 1800+).

### G.1 Adaptation in Othello Arena Stages

I examined adaptation patterns by having the participant navigate through progressively complex Othello Arena stages. Despite limited prior Othello experience beyond basic concepts like corner advantages, this provided an ideal case study for rapid environmental adaptation within our benchmark framework.

**Stage 3 to Stage 12: Rule Variation Discovery** The transition from Stage 3 (blocked cells with standard capture rules) to Stage 12 (identical board layout but modified capture mechanics allowing pieces to jump over obstacles) revealed sophisticated rule inference mechanisms. When initially attempting to play Stage 12 with Stage 3 strategies, the participant quickly recognized discrepancies between expected and actual game outcomes. Within two games, they formulated a crucial insight: "Stage 12’s evaluation criteria are closer to Stage 1 (standard  $8 \times 8$ ) than Stage 3," demonstrating rapid similarity matching between environmental variants.

The participant’s rule discovery process followed a systematic three-step cognitive pattern. First, they conducted internal simulations, stating "If I place here, the opponent will likely respond there," followed by intuitive strategic evaluation based on domain knowledge accumulated from previous games. Second, they analyzed opponent moves to infer underlying strategic principles, noting "I try to understand what the opponent wants to achieve by observing their moves." Third, they relied on visual similarity assessment, explaining that board layouts that appear similar often share similar rule structures, allowing for rapid strategic transfer.

**Domain Knowledge Transfer and Linguistic Abstraction** Throughout the adaptation process, the participant leveraged abstract strategic principles commonly used in perfect-information games. They employed heuristic guidelines such as "occupy the center first" and "secure corners early," demonstrating how linguistic abstractions facilitate rapid strategy transfer across game variants. Remarkably, the participant noted that when teaching complete novices, strategic concepts are often communicated through everyday language expressions like "good shape" or "strong position," suggesting that natural language serves as a bridge between complex strategic reasoning and rapid knowledge transfer.

When encountering rule ambiguity, the participant explicitly articulated their uncertainty management strategy: "In situations where I’m unsure about the rules, I tend to adopt more conservative, broadly applicable strategies rather than specific optimizations. If I know I have multiple attempts, I’ll use the first game for exploration—trying various moves to understand the dynamics—while in subsequent games I can optimize based on learned patterns."

**Stage 13: Turn Dynamics and Hypothesis Refinement** Stage 13 introduced a novel turn mechanism where players with fewer pieces on the board could take consecutive turns. This created a more complex adaptation challenge, as the participant initially formed an incorrect hypothesis about the underlying mechanics. They interpreted the rule as "players who move quickly can take multiple consecutive turns," leading to a time-pressure based strategy rather than the actual piece-count based mechanic.

Despite the misinterpretation, the participant demonstrated sophisticated hypothesis testing behavior. They noted, "When I suspect the rules have changed, I engage in more exploratory play during the first game, deliberately trying various moves to understand the new dynamics. This is different from

optimization-focused play where I'm trying to win—here I'm trying to learn." The participant's approach shifted toward broader, more generalizable strategies when facing uncertainty, prioritizing information gathering over immediate tactical advantage.

**Cognitive Mechanisms and Strategic Flexibility** The adaptation process revealed several key cognitive mechanisms that distinguish human learning from current AI approaches. The participant demonstrated within-episode learning, continuously refining their environmental model during gameplay rather than requiring separate training phases. They exhibited rapid transfer learning, immediately drawing analogies to familiar game patterns and applying relevant strategic principles. Most notably, they showed meta-cognitive awareness, explicitly reasoning about their own learning process and adjusting their exploration-exploitation balance based on the perceived novelty of the environment.

When asked about their approach to reward function estimation in novel environments, the participant explained: "I rely heavily on domain knowledge from previous games. When I encounter a new situation, I first look for similar patterns I've seen before, then gradually adjust my evaluation based on what works and what doesn't. The key is having these broad strategic principles that work across different games, then fine-tuning them to the specific rules I discover."

This observation suggests that effective adaptation in complex environments may require both rapid pattern recognition capabilities and flexible strategic frameworks that can be quickly reconfigured based on environmental feedback—characteristics that current AI systems struggle to replicate efficiently.

## G.2 Adaptation in YINSH Game

As he demonstrated the ability to defeat intermediate AI opponents and human players in both Othello Arena's hidden stages and Ataxx within 2–3 games. To explore the limits of rapid human adaptation, I introduced him to YINSH<sup>4</sup>, more complex abstract strategy game that combines elements of Othello and Gomoku with additional strategic layers involving ring management and multi-objective optimization.

**Rapid Environment Modeling and Rule Discovery** During the first game (~30 moves, ~15 minutes), the player demonstrated remarkable efficiency in discovering environment dynamics within a single episode. They correctly identified the core winning condition (graduating three rings by forming consecutive lines of five pieces), understood the dual mechanics of piece placement and ring movement, and began forming hypotheses about capture mechanisms. Notably, approximately 70% of the game's fundamental mechanics were grasped within this first encounter. However, some misconceptions emerged initially—for instance, believing that forming exactly five pieces (rather than five or more) was required for ring graduation, and uncertainty about which ring to remove when multiple options were available. This rapid yet imperfect rule acquisition represents an extremely compressed learning timeline compared to traditional reinforcement learning approaches, which typically require hundreds or thousands of episodes for comparable environmental understanding.

**Strategic Hypothesis Formation and Transfer Learning** The second game revealed sophisticated strategy development through active hypothesis testing and knowledge transfer. The player immediately drew analogies to familiar games, recognizing blocking and connection patterns similar to Go and Gomoku. They applied prior strategic knowledge about piece activity, stating that "in perfect information games, it's important to keep all pieces active" and began implementing multi-step planning with explicit discussion of opponent modeling ("they're attacking, so I need to defend" and "this move is safe for now"). The player exhibited within-episode learning, continuously refining their understanding of positional value, ring positioning, and tactical combinations as the game progressed. This demonstrates the human capacity for rapid transfer learning and strategic reasoning under minimal data exposure—a stark contrast to typical online reinforcement learning which requires numerous episodes for comparable strategic development.

**Time Constraint Analysis** Interpreting human adaptation through our framework reveals interesting parallels and contrasts with AI systems. Human learning effectively combines  $T_{\text{analysis}}$  and

---

<sup>4</sup>[www.yinsh.net](http://www.yinsh.net)

$T_{\text{game}}$  in a fluid, interleaved manner similar to ARC-PRIZE’s structure, where each 15-minute game serves simultaneously as analysis time (discovering rules and forming strategies) and execution time (implementing tactical decisions). The key insight is that humans achieve intermediate-level play within merely 2–3 attempts (30–45 minutes total), while maintaining the ability to deliberate and refine hypotheses in real-time during gameplay. Unlike our benchmark’s clear separation between analysis and execution phases, humans seamlessly transition between environment modeling, strategy formation, and tactical execution within the same episode.

**Implications for AGI Development** This observation pattern aligns with recent developments in test-time compute scaling, where systems like GRPO [5] and test-time reinforcement learning [12] demonstrate that complex novel tasks require active hypothesis building and refinement during inference, rather than relying solely on pre-trained knowledge. The human adaptation pattern suggests that effective artificial general intelligence may need similar capabilities for dynamic reasoning and rapid environmental modeling under minimal data exposure. Furthermore, the stark efficiency gap between human adaptation (2–3 games) and current AI approaches (typically requiring thousands of simulations as seen in our template system) highlights a fundamental challenge: while our benchmark systems can eventually match or exceed human performance through extensive simulation within  $T_{\text{analysis}}$ , they lack the elegant efficiency of human learning that combines rapid rule discovery, strategic transfer, and real-time adaptation within a unified cognitive process.