# CHESSARENA: A CHESS TESTBED FOR EVALUATING STRATEGIC REASONING CAPABILITIES OF LARGE LANGUAGE MODELS

# **Anonymous authors**

Paper under double-blind review

#### **ABSTRACT**

Recent large language models (LLMs) have shown strong reasoning capabilities. However, a critical question remains: do these models possess genuine reasoning skills—particularly complex strategic reasoning—or are they primarily excelling at sophisticated pattern recognition within their training data? To address this question, this paper presents a chess testbed, ChessArena, to evaluate the strategic reasoning capabilities of LLMs. Chess requires complex strategic reasoning capabilities including long-term planning, strict rule comprehension, and multi-turn conversation memorization. Specifically, ChessArena is a competitive framework where LLMs play against each other, under four different play modes. The testbed is equipped with a ranking algorithm and a leaderboard. The testbed can also evaluate fine-grained capabilities including basic understanding, move selection, and puzzle solving. Over 13 LLMs with different modes are evaluated in ChessArena, playing over 800 games. The results reveal significant shortcomings in current LLMs: no model can beat Maia-1100 (a chess engine at human amateur level), while some even failed to defeat a random player that selects moves arbitrarily. We also present a strong baseline to the testbed: our fine-tuned Qwen3-8B substantially improved performance, approaching much larger state-of-the-art reasoning models.

#### 1 Introduction

Large language models (LLMs) have demonstrated remarkable capabilities across diverse domains, from code generation (Jimenez et al., 2023) to mathematical problem-solving (Cobbe et al., 2021). One significant contributing factor to the success is the availability of high-quality benchmarks such as LiveCodeBench (Jain et al., 2024) and AIME2025 (MAA, 2025).

As LLMs are increasingly applied in real-world problems, improving their strategic reasoning capability, i.e., the reasoning under dynamic environments and uncertain adversary actions (Gandhi et al., 2023; Duan et al., 2024), becomes an urgent demand. However, there is still a lack of well-established evaluation frameworks for effectively evaluating the strategic reasoning capabilities of LLMs. Existing evaluation frameworks typically focus on isolated question-answering tasks that may not capture the essential aspects of strategic reasoning (Lin et al., 2025; Kazemi et al., 2025; Dua et al., 2019; Chen et al., 2021). Additionally, current benchmarks (Austin et al., 2021; Sprague et al.) often suffer from data contamination, where test examples may have appeared in training data.

In this work, we choose chess as the testbed for evaluating the strategic reasoning capability of LLMs, as it provides an ideal environment requiring the ability to maintain coherent strategies across prolonged gameplay, follow complex instructions consistently, and adapt reasoning based on evolving contexts. Additionally, the vast state space of chess—with an estimated  $10^{47}$  possible board positions—virtually eliminates data contamination concerns.

Building on these advantages, we introduce ChessArena, a competitive platform where LLMs engage in complete chess games from opening to endgame. Our system implements a comprehensive ranking mechanism using approximately 30 games per model to ensure stable performance assessment. We evaluate models across four distinct play modes—Bullet, Blitz, Standard, and Blindfold—each designed to test different aspects of model capability, from rapid decision-making to memory retention in long-term contexts.

Our evaluation of over 13 state-of-the-art models, including O3, Gemini-2.5-Pro, and Doubao-Seed-1-6-Thinking, reveals their significant limitations. No model successfully defeated Maia-1100, a chess engine designed to play at a human amateur level, with some models losing even to a random player that arbitrarily selects a move from all legal moves. These results highlight three critical deficiencies: inconsistent instruction following (failure to maintain proper output formatting), weak tactic reasoning (selecting moves inferior to random choices), and limited multi-turn coherence (inability to maintain consistent play across multi-round games).

To investigate the underlying causes of these performance gaps, we developed three complementary evaluation tasks targeting specific reasoning components: basic rule understanding, single-move evaluation, and multi-step puzzle solving. These fine-grained assessments, combined with the competitive arena, provide comprehensive insight into model capabilities and limitations.

Finally, we demonstrate that post-training can address some of these deficiencies. Using high-quality gameplay data collected from ChessArena competitions, we fine-tuned Qwen3-8B through supervised learning followed by reinforcement learning. The resulting Qwen3-8B-Chess model shows substantial improvements in chess performance.

Our work makes three primary contributions:

- ChessArena Platform: We introduce a competitive evaluation framework for chess play. It is extensible, providing interfaces to any LLM participants. It supports the evaluation of complete games as well as fine-grained tasks targeted at specific reasoning components.
- Empirical Findings: Our testbed exposes critical gaps in current LLMs' strategic reasoning through over 800 systematic gameplays. Fine-grained evaluations also provide detailed insight into the sources of model limitations.
- Training Data and Model: We collect and curate high-quality strategic reasoning data from ChessArena, and demonstrate through the Qwen3-8B-Chess model that training on strategic reasoning data can improve performance.

#### 2 CHESSARENA

#### 2.1 OVERVIEW

As shown in Fig. 1, ChessArena is a simulation platform where LLMs compete against each other to acquire quantitative chess strength ratings. Each model operates independently, generating moves based solely on the current chessboard state, closely emulating human competitive play. During gameplay, models receive task instructions and board representations, analyze the position, and predict moves that iteratively update the chessboard state. Additionally, our ChessArena competitions offer high scalability, with unified interface management that facilitates easy integration of new LLMs without affecting existing rankings. We use Forsyth-Edwards Notation (FEN) (rec.games.chess, 1994) for board representation and support both Universal Chess Interface (UCI) (Kahlen, 2004) and Standard Algebraic Notation (SAN) (rec.games.chess, 1994) for move representations. For more information about these representations, please refer to Appendix B.4.

#### 2.2 PLAY MODES

To better evaluate the ability of LLMs, we design four play modes inspired by Lichess<sup>1</sup>. Each LLM player can be associated with one of the following four modes.

- **Bullet**: Given the chessboard state, the LLM must directly generate a move without any intermediate reasoning. Outputs containing any form of thinking process will be rejected.
- **Blitz**: Given the chessboard state, the LLM may optionally include a reasoning process before producing the move. This mode is designed specifically for non-thinking LLMs.
- **Standard**: Given the chessboard state, the LLM must generate a move accompanied by a chain-of-thought (CoT) reasoning process. This mode is designed specifically for thinking LLMs.

https://lichess.org/

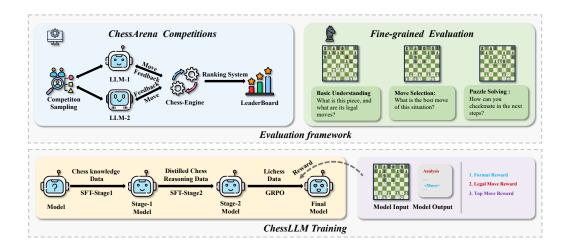


Figure 1: Overview of ChessArena competitions, fine-grained evaluation, and ChessLLM training. (1) An LLM can be integrated into ChessArena to compete against other models. After a certain number of competitions, each model is assigned a reliable Glicko rating and added to the leaderboard. (2) Three additional evaluation tasks are integrated into ChessArena to evaluate the chess capabilities at a fine-grained level. (3) We can extract high-quality chess reasoning data from the gameplay process, which can be used for training an LLM specially for chess.

• **Blindfold**: This mode represents the highest difficulty level. The model is provided with the move history from both players in the form of a multi-turn conversation. The LLM must reconstruct the chessboard state internally and produce a move with a thorough analysis.

#### 2.3 RANKING SYSTEM

Glicko Rating System. We adopted the Glicko rating system (Glickman, 1995) as our ranking algorithm. As an enhancement of the traditional Elo rating system, Glicko represents each player's chess strength using two parameters: the rating r (similar to traditional Elo) and the rating deviation d that reflects the uncertainty in the rating. A high d indicates that the player's rating is still unreliable and requires more matches to stabilize. In our scenarios, both parameters update after each competition, with d decreasing monotonically as the system gains confidence in the player's skill level.

Competition Sampling Strategy. We developed a competition sampling algorithm to accelerate the convergence of rating deviation (d). Mathematical analysis shows that d reduction is maximized when opponents have similar ratings (r) and low d values, as matches between players of comparable and established skill levels yield the most information. Complete details of the Glicko rating system and proofs regarding the competition sampling strategy are provided in Appendix C. Our algorithm enables new players to achieve reliable ratings (d < 100) within approximately 30 games.

#### 2.4 CHESS ENGINE

Regarding the chess engine, we chose Stockfish (Stockfish Development Team, 2016), which is currently the most powerful chess engine and has been widely used in chess analysis. We utilize the analysis results from Stockfish as a critical reference for subsequent evaluations. Specifically, given a search depth and a chessboard state, we use Stockfish to analyze the win rates of all legal moves for the current state. We consider moves with win rates in the top-3 as "top moves" for subsequent analysis. Additionally, we employed two supplementary engines as players in our testbed.

**Maia-1100**. To better understand the gap between the LLMs and human chess players, we incorporated Maia-1100 (McIlroy-Young et al., 2020), a chess AI with an Elo rating of approximately 1600 on real human chess platforms,<sup>2</sup>, which is roughly the average level for human players. Maia-1100 is specially developed for chess, and it is based on CNN and Monte Carlo Tree Search.

<sup>&</sup>lt;sup>2</sup>https://lichess.org/@/maia1

**Random Player**. We also included a random player, which chooses randomly from all legal moves on the board. Note that this player is not purely random, as we provide the legal moves to it.

#### 2.5 FINE-GRAINED EVALUATION TASKS

In addition to the overall Glicko rating, ChessArena also provides more comprehensive evaluations of the strategic reasoning capabilities of LLMs. We design three fine-grained tasks as follows.

**Basic Understanding**. This task evaluates models' basic understanding of chess rules and board states by testing their ability to generate legal moves. Given a chessboard state and a specific position, models must identify the piece of the given position (e.g., King or Queen) and generate all legal moves of this piece. We assess this basic understanding capability using three metrics: *Piece Match Accuracy (PMA)*, which measures the accuracy of piece identification, and *Precision* and *Recall*, which measure the accuracy of legal move prediction. To strengthen the evaluation, we introduce perturbations including empty squares and turn-mismatch scenarios (e.g., requesting a Black piece when it is White's turn). In such cases, the correct response should be no legal moves.

Move Selection. This task evaluates models' single-move chess-playing ability by requiring them to select optimal moves from a given board state. We assess performance using three metrics: Legal Rate (LR), Top Rate (TR), and Move Advantage Rate (MAR). LR quantifies the proportion of legal moves predicted by the model. TR evaluates whether the model's predictions are included in the "top moves" as determined by Stockfish. MAR measures the relative strength of a model's predicted move compared to all legal moves. Using Stockfish to evaluate win rates Q(FEN, move) for all legal moves from a given chessboard, we compute the Average Win Rate (AWR) as  $\frac{1}{M} \sum_{m=1}^{M} Q(\text{FEN}, \text{Move}_m)$ , where M is the number of legal moves. MAR is then calculated as:

$$\mathrm{MAR} = \frac{1}{N} \sum_{i=1}^{N} \frac{Q(\mathrm{FEN}_i, \mathrm{Move_{pred}}) - \mathrm{AWR}_i}{\mathrm{AWR}_i},$$

where N is the number of evaluation instances. For illegal moves, we set  $Q(\text{FEN}, \text{Move}_{\text{pred}}) = 0$ .

**Puzzle Solving**. In line with the work of Hwang et al. (2025) and Ruoss et al. (2024), we evaluate chess puzzle solving using the Lichess puzzle database<sup>3</sup>. Each puzzle begins from an initial board state and consists of k ground-truth sequential moves that represent the solution. At each step, we present the current board state and require the model to predict the optimal move. A puzzle is considered solved only if all k predicted moves match the ground truth exactly—even a single error in any step results in failure. We utilize the puzzle dataset from the Lichess puzzle database, where each puzzle is associated with an Elo rating ranging from 200 to 3000 on the Lichess platform. We use *Puzzle Solving Accuracy (PSA)* as our evaluation metric, which measures the percentage of puzzles that the model correctly solves.

#### 3 Post-train Qwen3 for ChessArena

To explore potential solutions to the observed strategic reasoning limitations exhibited by the models, we post-train an LLM (named Qwen3-8B-Chess) on Qwen3-8B, which is the weakest among the studied LLMs. Our post-training includes two stages of supervised fine-tuning (SFT) and one stage of group relative policy optimization (GRPO) (Shao et al., 2024).

**Supervised fine-tuning**. This phase aims to gain basic chess reasoning ability. It consists of two stages. In the first stage of SFT, we use chess-based dialogue data from ChessGPT (Feng et al., 2023), which covers discussions on basic chess rules, tactics, etc. This stage injects the background knowledge about chess into the model. In the second stage of SFT, we collect and filter data from games played on ChessArena, which is critical for endowing the model with fundamental chess reasoning skills.

**Group relative policy optimization**. In the following stage, we further enhance the chess ability through GRPO. GRPO has been demonstrated as an effective method for enhancing a model's reasoning capabilities, particularly when verifiable rewards (e.g., for mathematics or code generation)

<sup>3</sup>https://database.lichess.org

are employed (Guo et al., 2025). This is also the case for our chess scenario. Specifically, we utilize Stockfish to analyze the model's moves and generate verifiable reward signals, enabling the model to autonomously explore chess strategies through this feedback mechanism. We define three types of rewards: format reward, legal move reward, and top move reward. For more details about post-training (e.g., training data collection, reward design, and training hyperparameters), please refer to Appendix D.

# 4 Experiments

# 4.1 EXPERIMENTAL SETUP

**Evaluated Models.** We evaluated leading proprietary and open-source LLMs, including GPT-4.1 (OpenAI, 2025a), GPT-4o (OpenAI, 2024), O3 (OpenAI, 2025b), Claude-3-7-Sonnet (Anthropic, 2025), Gemini-2.5-pro (Google DeepMind, 2025), Qwen3-235B-A22B(Non-thinking) (Yang et al., 2025), DeepSeek-R1 (Guo et al., 2025), DeepSeek-V3 (Liu et al., 2024a), DeepSeek-V3.1(Non-thinking), and Doubao (Seed et al., 2025) series. We also include Qwen3-8B and our trained Qwen3-8B-Chess. All evaluated models are shown in Table 5.

**Ranking System**. We initialized our Glicko rating system with r=1500 and d=350, setting a minimum rating deviation of d=50 to ensure meaningful rating adjustments throughout the competition. Following Lichess, we display only players with  $d \leq 100$ , indicating statistically reliable ratings.

**Implementation Details**. We set the max new tokens to 4096 for non-thinking models and 16384 for thinking models. We set *temperature* as 0.2 and *top\_p* as 1 for all experiments, as we observed optimal model performance under these configurations. We evaluated the models under two distinct conditions: with and without the provision of legal moves. We provide the latest 10 moves as partial move history in UCI format to give models sufficient game context and prevent fivefold repetition draws that occurred frequently without this historical information. In ChessArena, the termination conditions adhere to standard chess rules, including checkmate, forfeit, stalemate, insufficient material, fivefold repetition, the 75-move rule, and move limit (please refer to Appendix B.3 for further details). Our experimental setup employs two deployment approaches: official model APIs for most models, and local deployment via vLLM (Kwon et al., 2023) for Qwen3-8B and its post-training versions.

For basic understanding and move selection evaluation, we constructed the evaluation data from actual competitions in ChessArena. There are 200 and 1000 samples in these two tasks, respectively. To ensure a fair comparison of the model's performance across the four play modes, we guaranteed that the evaluation data for these four play modes consisted of identical board states. For the Blindfold mode in the move selection experiment, the average number of conversation turns in our evaluation data is 47. For puzzle solving experiments, we use 1008 puzzles from the Lichess puzzle database. For more details about fine-grained evaluation data construction, please refer to Appendix E.

All of our prompt templates are shown in AppendixB.2, including prompt templates in ChessArena competitions of different play modes and fine-grained evaluation. In actual competitions, players compete in an even number of games (alternating between playing as White and Black) to balance the first-move advantage.

#### 

#### 4.2 EXPERIMENTAL RESULTS

**(A) LeaderBoard**. Our rating leaderboard is shown in Table 1. We provide legal moves to the vast majority of models, as they cannot effectively play without them.

*Key observations*. Among the models, we observe that thinking models such as O3, Doubao-Seed-1-6-Thinking, and Gemini-2.5-Pro currently lead the rankings. Among non-thinking models, GPT-4.1 and Claude-3-7-Sonnet show the strongest performance. When legal moves are provided, our trained Qwen3-8B-Chess achieves the best performance among all non-thinking models (including GPT-4.1). Notably, the untrained Qwen3-8B baseline ranks at the bottom of the leaderboard.

Comparing to Maia-1100 and Random Player. There remains a significant performance gap between LLMs and Maia-1100. Currently, no LLM has demonstrated the capability to defeat Maia-1100 in actual gameplay, which demonstrates the inadequacy of the model's strategic reasoning capabilities.

Table 1: Leaderboard of ChessArena. Thinking models generally perform better than non-thinking models, while all models are inferior to Maia-1100. When legal moves are provided, our post-trained Qwen3-8B-Chess outperforms other non-thinking models and is on par with thinking models. ('Legal Moves' stands for whether the set of legal moves are provided to the model, 'RD' means the rating deviation, and 'Interval' means the 95% confidence interval for the rating.)

Rank	Model	Mode	Legal Moves	Rating	RD	Interval	Games
1	Maia-1100	-	×	2220	82	(2058, 2382)	44
2	O3	Standard	×	1948	78	(1793, 2101)	28
3	Doubao-Seed-1-6-Thinking	Standard	✓	1830	50	(1729, 1929)	60
4	Gemini-2.5-Pro	Standard	✓	1819	81	(1659, 1979)	18
5	Qwen3-8B-Chess	Blitz	✓	1776	93	(1593, 1959)	16
6	Doubao-Seed-1-6-Thinking	Standard	×	1743	66	(1612, 1873)	36
7	GPT-4.1	Blindfold	✓	1699	50	(1601, 1797)	60
8	Doubao-Seed-1-6-Thinking	Blindfold	✓	1687	73	(1542, 1831)	24
9	GPT-4.1	Blitz	✓	1686	50	(1588, 1784)	182
10	Claude-3-7-Sonnet	Blitz	✓	1654	50	(1555, 1751)	74
11	Claude-3-7-Sonnet	Blindfold	✓	1625	66	(1493, 1756)	30
12	GPT-4.1	Blitz	×	1623	50	(1525, 1721)	106
13	Gemini-2.5-Pro	Standard	×	1616	74	(1469, 1762)	28
14	Qwen3-8B-SFT-Stage2	Blitz	✓	1612	56	(1501, 1721)	40
15	Claude-3-7-Sonnet	Blindfold	×	1588	72	(1445, 1729)	28
16	GPT-4.1	Bullet	✓	1583	50	(1485, 1681)	54
17	DeepSeek-V3	Blitz	✓	1553	50	(1454, 1650)	174
18	Random Player	-	✓	1524	50	(1425, 1621)	284
19	Qwen3-235B-A22B	Blitz	✓	1483	50	(1385, 1581)	146
20	DeepSeek-V3	Blitz	×	1482	58	(1367, 1597)	48
21	DeepSeek-V3	Blindfold	✓	1437	75	(1290, 1584)	24
22	DeepSeek-V3	Bullet	✓	1382	80	(1224, 1540)	22
23	Qwen3-235B-A22B	Bullet	✓	1369	54	(1261, 1476)	46
24	Qwen3-8B	Blitz	✓	1335	65	(1205, 1463)	32

Compared to the random player baseline, most models exhibit better performance. However, a few models still underperform the random player. This primarily occurs when models fail to generate legal moves due to the lack of instruction-following ability, resulting in forfeit losses. Table 11 presents metrics demonstrating substantial *parsing error rates* across models, indicating format noncompliance and instruction-following deficiencies. Several models also exhibit elevated *illegal move rates* even when legal moves are explicitly provided. While high *illegal move rates* are anticipated when legal moves are not provided, rates exceeding 5% in scenarios with provided legal moves warrant attention.

Different play modes. For the same model under different play modes, we observe that most models achieve their best performance in Blitz or Standard modes. This aligns with expectations, as these modes provide the model with the most direct board information while permitting reasoning. Under Blindfold conditions, O3, Doubao-Seed-1-6-Thinking, GPT-4.1, and Claude-3-7-Sonnet still demonstrate competent playing strength. They demonstrate stronger multi-turn memorization and long-term strategic reasoning capabilities than other models. However, in Bullet mode, nearly all models perform poorly. This suggests that prohibiting thought chain output (e.g., "Let me think ..." or reasoning steps) severely impairs the models' chess strategic reasoning capabilities.

- **(B) Basic Understanding.** Table 2 shows the results of the basic understanding task. It can be seen that thinking models (e.g., O3, Doubao-Seed-1-6-Thinking and DeepSeek-R1) have almost complete chessboard understanding capabilities, being able to identify pieces at specific positions on the board and generate related legal moves according to chess rules. Additionally, some strong non-thinking models, such as GPT-4.1 and Claude-3-7-Sonnet, also have relatively high *PMA*, *Precision* and *Recall*. Our trained Qwen3-8B-Chess shows improvement over Qwen3-8B on this task, even though we did not specifically train on this task.
- (C) Move Selection. Table 3 shows the results of move selection. We share our findings below.

*LLMs have significant room for improvement in the strategic reasoning of chess.* Among all LLMs we evaluated, thinking models such as O3, Gemini-2.5-Pro, and Doubao-Seed-1-6-Thinking performed the best, while GPT-4.1 and Qwen3-8B-Chess also showed relatively excellent performance. However,

Table 2: Basic understanding results. Thinking models such as O3 and Doubao-Seed-1-6-Thinking show strong chessboard understanding capabilities. Our post-training significantly improves the basic understanding capability.

Model	PMA (%)	Precision (%)	Recall (%)
GPT-4.1	98.0	89.3	92.1
O3	98.5	98.5	98.5
DeepSeek-V3	97.0	81.8	75.3
DeepSeek-V3.1	89.0	87.5	87.4
DeepSeek-R1	100.0	99.2	98.4
Doubao-1-5-Pro-32k	76.0	50.6	56.2
Doubao-1-5-Lite-32k	51.5	33.3	30.3
Doubao-1-5-Thinking-Pro	99.5	98.0	98.0
Doubao-Seed-1-6-Thinking	100.0	99.9	99.9
Qwen3-235B-A22B	80.5	50.7	49.3
Claude-3-7-Sonnet	98.0	87.6	87.3
Gemini-2.5-Pro	100.0	98.5	96.7
Qwen3-8B	36.0	14.1	18.8
Qwen3-8B-Chess-SFT-Stage1	63.5 (+31.5)	20.6 (+5.9)	29.5 (+14.3)
Qwen3-8B-Chess-SFT-Stage2	70.5 (+7.0)	51.9 (+31.3)	45.3 (+15.8)
Qwen3-8B-Chess (SFT+RL)	79.0 (+8.5)	52.6 (+0.7)	50.1 (+4.8)

their TP and MAR are far behind Maia-1100, which indicates that LLMs still have significant room for improvement in chess strategic reasoning capabilities. When legal moves are not provided, the performance of most models is even worse, as indicated by the negative MAR values.

Bullet and Blindfold chess games bring difficulties to LLMs. In terms of comparison among different play modes, the performance of LLMs (e.g, O3, Doubao-Seed-1-6-Thinking, DeepSeek-R1, GPT-4.1, Qwen3-235B-A22B, DeepSeek-V3.1) in Bullet or Blindfold mode is usually worse than in Blitz/Standard mode. Both Bullet (thinking content restricted) and Blindfold (multi-turn conversation reconstruction) pose certain difficulties for LLMs.

Thinking models tried to reconstruct the chessboard. For the Blindfold chess experiment, it appears that different models exhibit significant variations in performance. First of all, compared to Standard mode, DeepSeek-R1 and Doubao-Seed-1-6-thinking show a noticeable decline in performance in Blindfold mode. We manually checked their response and found that they were trying to reconstruct the chessboard, which brings much difficulty for them, especially when the number of conversation turns is large (i.e., more than 90 turns).

Non-Thinking models may be lazy in Blindfold chess games. GPT-4.1 and Qwen3-235B-A22B also demonstrate a significant drop when in Blindfold play mode compared to their performance in Blitz mode. For experiments Blindfold and without legal moves, we find that GPT-4.1, Qwen3-235B-A22B, and DeepSeek-V3 mostly base their responses on the last move in the conversation and continue from there, showing signs of laziness. They are more like guessing a move. Claude-3-7-Sonnet seems to reconstruct the chessboard genuinely and accomplishes this automatically without spending many response tokens. Overall, Blindfold chess poses significant challenges for models, revealing deficiencies in multi-turn reasoning capabilities.

Table 12 shows the average conversation turns for models predicting legal versus illegal moves in Blindfold mode without legal move provision. For non-thinking models, conversation turns had minimal impact on performance. However, thinking models required significantly fewer turns to predict legal moves compared to illegal ones, indicating that longer conversations impede board reconstruction. This disparity reveals the lazy behavior exhibited by non-thinking models. Notably, O3 maintains performance across more conversation turns than DeepSeek-R1 and Doubao-Seed-1.6-Thinking, demonstrating superior multi-turn memorization and reasoning capabilities. For more information about Blindfold mode analysis, please refer to Appendix G.1.

**(D) Puzzle solving.** Table 4 presents the main experimental results of puzzle solving, where we divide the puzzles according to their Elo ratings. Stockfish achieved an overall score of 98.4%, which aligns with expectations. Maia-1100 attained an overall score of 74.6%. Among all LLMs, the O3 model stands out remarkably, achieving a puzzle-solving rate of 55.6%. Other models all scored below 15%. Overall, thinking models outperformed non-thinking models. Our trained Qwen3-8B-Chess

Table 3: Move selection performance across four play modes with/without legal moves provision. We bold the highest LR, TR, and MAR within each group. LLMs still have significant room for improvement, especially when the legal moves are not provided.

Mode	Model or Engine	W	ith Legal N	<b>I</b> oves	Witl	nout Legal	Moves
Mouc	Widder of Engine	LR (%)	TR (%)	MAR (%)	LR (%)	TR (%)	MAR (%)
	Random Player	100.0	14.8	-1.1	/	/	/
	Maia-1100	/	/	/	100.0	78.3	+107.6
	GPT-4.1	97.5	25.9	+20.5	71.6	29.3	+6.2
	Claude-3.7-Sonnet	99.6	26.1	+25.6	68.4	18.2	-17.7
Blitz	DeepSeek-V3	99.1	18.5	+10.7	64.5	12.9	-27.7
	DeepSeek-V3.1	93.4	26.7	+18.6	63.7	16.9	-23.6
	Qwen3-235B-A22B	89.8	24.9	+29.0	64.2	17.0	-25.3
	Qwen3-8B	96.2	13.4	+1.8	9.8	2.1	-79.5
	Qwen3-8B-Chess-SFT-Stage1	86.8	13.6	-9.6	15.1	2.6	-74.9
	Qwen3-8B-Chess-SFT-Stage2	96.9	23.4	+15.1	66.3	13.3	-22.1
	Qwen3-8B-Chess (SFT+RL)	92.9	40.2	+41.1	87.6	20.2	-1.2
	GPT-4.1	98.7	25.0	+20.8	74.0	28.7	+5.7
	Claude-3.7-Sonnet	98.6	22.5	+16.8	75.2	17.9	-9.4
Bullet	DeepSeek-V3	98.9	18.8	+11.3	66.2	13.3	-21.8
	DeepSeek-V3.1	80.6	16.1	-8.0	56.3	12.7	-35.7
	Qwen3-235B-A22B	95.9	17.8	+4.5	69.1	15.9	-18.5
	DeepSeek-R1	100.0	32.7	+34.7	82.5	23.7	-1.0
	Doubao-1-5-Thinking-Pro	99.7	32.9	+35.4	78.0	24.8	+3.0
Standard	Doubao-Seed-1-6-Thinking	99.8	39.1	+53.7	90.7	36.0	+32.0
	Gemini-2.5-Pro	99.4	37.6	+46.5	85.5	40.5	+36.5
	O3	99.6	58.7	+80.1	98.0	62.0	+80.2
	GPT-4.1	96.8	20.1	+12.7	72.7	20.2	+1.2
	Claude-3.7-Sonnet	98.2	23.9	+21.5	77.3	18.9	-9.1
	DeepSeek-V3	95.1	19.2	+16.2	78.5	14.9	-7.8
Blindfold	DeepSeek-V3.1	96.5	26.0	+27.2	66.0	13.7	-18.0
Dinialola	DeepSeek-R1	94.7	22.7	+14.0	44.6	10.9	-36.9
	Qwen3-235B-A22B	96.1	19.9	+17.4	75.3	17.2	-10.4
	Doubao-Seed-1-6-Thinking	97.8	32.1	+36.5	43.6	12.9	-30.5
	Gemini-2.5-Pro	98.7	30.4	+23.5	68.7	21.5	-8.7
	O3	98.4	46.9	+63.2	86.9	43.5	+50.9

Table 4: Puzzle solving accuracy when legal moves are provided. LLMs perform relatively poorly, with O3 standing out as the strongest.

Model or Engine	Puzzle Solving Accuracy (%)										
Model of Engine	200-600	600-1000	1000-1400	1400-1800	1800-2200	2200-2600	2600-3000	Overall			
Stockfish (Depth=20)	100.0	100.0	100.0	100.0	99.3	97.9	91.5	98.4			
Maia-1100	98.6	97.2	91.6	82.5	72.7	51.0	28.2	74.6			
Random Player	1.4	1.4	2.1	0.0	0.0	0.0	0.0	0.7			
GPT-4.1	18.9	14.0	8.4	4.9	1.4	2.8	0.0	7.2			
Claude-3-7-Sonnet	18.2	16.1	4.9	4.2	5.6	1.4	0.0	7.2			
DeepSeek-V3	11.9	7.7	2.1	0.7	0.0	0.7	0.0	3.3			
DeepSeek-V3.1	13.3	10.5	8.4	4.9	1.4	2.8	7.0	6.0			
Qwen3-235B-A22B	24.5	18.2	9.8	5.6	4.2	1.4	0.0	9.1			
Qwen3-8B	2.8	4.9	2.1	0.0	0.0	0.0	0.0	1.4			
Qwen3-8B-Chess	31.5	16.8	10.5	7.0	5.6	2.1	0.0	10.5			
O3	97.9	90.2	79.7	62.9	46.5	10.5	1.4	55.6			
Gemini-2.5-Pro	37.1	24.5	18.2	9.1	4.2	3.5	1.4	14.0			
Doubao-Seed-1-6-Thinking	27.3	23.8	11.9	7.7	4.2	1.4	2.1	11.2			
DeepSeek-R1	23.1	20.3	7.0	4.2	2.8	0.7	0.7	8.4			

achieved the highest score among non-thinking models. We also present the puzzle-solving results without legal moves in Table 13. As can be observed, stronger models such as O3, GPT-4.1, and Gemini-2.5-pro exhibit almost no performance degradation, whereas weaker models are significantly affected. The models' deficiencies in puzzle solving task indicate persistent limitations in long-term reasoning capabilities.

# 5 RELATED WORK

Chess Language Model. Recent studies have explored LLM applications to chess with interesting findings. Xiangqi-R1 (Chen et al., 2025) achieved strong performance in Chinese chess through SFT and GRPO training, while Hwang et al. (2025) encountered significant bottlenecks when applying RL methods to chess puzzle solving, which the authors attribute to the model's inadequate acquisition of chess-related knowledge during pretraining. Carlini (2023) discovered that GPT-3.5-turbo could play chess using PGN format, but deeper analysis revealed reliance on memorized patterns rather than genuine reasoning. In contrast, Chess Bench (Ruoss et al., 2024) achieved grandmaster-level performance using a 270M Transformer pre-trained through Stockfish knowledge distillation, though this represents a domain-specific architecture rather than a general language model. ChessGPT (Feng et al., 2023) represents a systematic approach, fine-tuning RedPajama-3B on web-scraped chess data to significantly outperform base models, while also contributing a high-quality chess-related training dataset. Similarly, Wang et al. (2025) fine-tuned LLaMA3-8B on expert-annotated datasets targeting tactics and strategy, achieving performance superior to GPT-40 on their benchmarks. However, their evaluation task—selecting the better move between two given options—is considerably less challenging than actual gameplay.

LLMs Evaluation Benchmark. ChatBot Arena (Chiang et al., 2024) introduced human preference-based evaluation using Elo rankings, grounding model assessment in naturalistic user interactions. SWE-Bench (Jimenez et al., 2023) evaluates LLMs on real-world software engineering tasks, while LiveCodeBench (Jain et al., 2024) provides continuously updated coding benchmarks from LeetCode and CodeForces to prevent data contamination. AIME2025 (MAA, 2025) assesses mathematical reasoning through 30 olympiad-level problems from the American Invitational Mathematics Examination. For strategic reasoning evaluation, GT-Bench (Duan et al., 2024) employs game-based scenarios to assess LLMs' strategic capabilities. ZebraLogic (Lin et al., 2025) tests logical reasoning through zebra puzzles of varying complexity. BBH (Kazemi et al., 2025) comprises 23 challenging multi-step reasoning tasks from BIG-Bench. Most relevant to our work, a concurrent work, GameArena (Lee et al., 2025), evaluated eight LLMs' chess abilities in blitz-style competitions. We introduce ChessArena as a comprehensive testbed with multiple gameplay scenarios and fine-grained studies to evaluate strategic reasoning capabilities in current language models.

#### 6 Discussion

**Generalization.** One interesting problem is whether models trained on chess-specific domains with enhanced strategic reasoning capabilities can be generalized to other domains. In Appendix G.2, we evaluate our chess-specific trained model's performance on other benchmarks. We found that our trained model demonstrated improvements on benchmarks such as AIME2025 and ZebraLogic, while maintaining comparable performance on other benchmarks.

**Limitations.** Our training data underwent outcome supervision filtering (evaluating only the quality of final moves without examining the reasoning process); this may result in training data containing cases where the reasoning process is flawed, but the final move is correct, potentially introducing noise into the dataset. This is a common issue shared across domains that rely on outcome supervision, such as code generation and mathematical reasoning. Besides, our trained model performs well under the "with legal moves" setting but poorly when such legal moves are not provided. This indicates that the model may still depend on memorization instead of developing genuine strategic understanding—a challenge potentially too demanding for 8B-parameter models. Alternatively, employing continued pre-training (Zhou et al., 2024) for the first stage of our SFT may be a viable option to improve its capability.

Conclusion. We introduce ChessArena, a competitive platform enabling large language models to play against each other in human-like chess competitions. Through authentic gameplay, we evaluate LLMs' strategic reasoning, instruction following, and multi-turn conversational memorization capabilities. Our analysis through ChessArena gameplay and fine-grained evaluation reveals substantial room for improvement in LLMs' chess strategic reasoning abilities. Observing deficiencies of current LLMs, we trained Qwen3-8B-Chess and achieved significant improvements in chess strategic reasoning capabilities. We hope our ChessArena platform, fine-grained evaluation tasks, and high-quality training datasets will contribute to future large language model research.

# ETHICS STATEMENT

ChessArena is constructed entirely from publicly available chess gameplay data and standard chess notation formats that permit open research usage, ensuring our contributions comply with established data usage protocols. During the collection and evaluation processes, we do not gather personal information about players or participants, and ChessArena task instances utilize only standard chess positions and moves that are part of the public domain of chess knowledge. Our contributions do not involve any human subject participation; we do not perform crowdsourcing or recruit human annotators for any component of ChessArena, including data collection, game execution, and evaluation procedures. ChessArena's model selection and evaluation criteria are based solely on objective performance metrics and do not implicitly or explicitly rely on any discriminative or biased heuristics for model assessment. For the dataset release, we plan to open-source the ChessArena competition platform, the ranking system implementation, competition sampling algorithms, fine-grained evaluation frameworks, experimental results, the training data used for fine-tuning the Qwen3-8B-Chess model, and the model weights. Following established best practices, we will provide comprehensive documentation describing each component and its usage, and establish accessible communication channels for soliciting community feedback to improve ChessArena. ChessArena does not present any immediately harmful applications, as chess gameplay represents a benign domain for evaluating strategic reasoning capabilities.

#### REPRODUCIBILITY STATEMENT

In our submitted source code, we provide all corresponding code for ChessArena competition, ranking system, competition sampling, and fine-grained evaluation, as well as chess training-related SFT and GRPO training data examples and code, enabling readers to reproduce our results. In the future, we plan to open-source all code with corresponding documentation. We also intend to release the trained Qwen3-8B-Chess model to facilitate subsequent research.

#### REFERENCES

- Anthropic. Claude 3.7 Sonnet and Claude Code. Technical report, Anthropic, 2025. URL https://www.anthropic.com/news/claude-3-7-sonnet.
  - Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- Nicholas Carlini. Language models can play chess: An analysis of GPT-3.5-turbo-instruct's chess playing ability. Personal Blog, September 2023. URL https://nicholas.carlini.com/writing/2023/chess-llm.html. Blog post analyzing how language models can play chess.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- Yuhao Chen, Shuochen Liu, Yuanjie Lyu, Chao Zhang, Jiayao Shi, and Tong Xu. Xiangqi-r1: Enhancing spatial strategic reasoning in llms for chinese chess via reinforcement learning. *arXiv* preprint arXiv:2507.12215, 2025.
- Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Banghua Zhu, Hao Zhang, Michael Jordan, Joseph E Gonzalez, et al. Chatbot arena: An open platform for evaluating llms by human preference. In *Forty-first International Conference on Machine Learning*, 2024.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Jianbo Dai, Jianqiao Lu, Yunlong Feng, Guangtao Zeng, Rongju Ruan, Ming Cheng, Dong Huang, Haochen Tan, and Zhijiang Guo. Mhpp: Exploring the capabilities and limitations of language models beyond basic code generation. *arXiv* preprint arXiv:2405.11430, 2024.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 2368–2378, 2019.
- Jinhao Duan, Renming Zhang, James Diffenderfer, Bhavya Kailkhura, Lichao Sun, Elias Stengel-Eskin, Mohit Bansal, Tianlong Chen, and Kaidi Xu. Gtbench: Uncovering the strategic reasoning capabilities of llms via game-theoretic evaluations. *Advances in Neural Information Processing Systems*, 37:28219–28253, 2024.
- Xidong Feng, Yicheng Luo, Ziyan Wang, Hongrui Tang, Mengyue Yang, Kun Shao, David Mguni, Yali Du, and Jun Wang. Chessgpt: Bridging policy learning and language modeling. *Advances in Neural Information Processing Systems*, 36:7216–7262, 2023.
- Kanishk Gandhi, Dorsa Sadigh, and Noah D Goodman. Strategic reasoning with language models. *arXiv preprint arXiv:2305.19165*, 2023.
- Mark E Glickman. The glicko system. *Boston University*, 16(8):9, 1995.
  - Google DeepMind. Gemini-2.5-Pro. Technical report, Google DeepMind, 2025. URL https://deepmind.google/models/gemini/.
  - Alex Gu, Baptiste Roziere, Hugh James Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida Wang. Cruxeval: A benchmark for code reasoning, understanding and execution. In *Forty-first International Conference on Machine Learning*.
  - Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

- Jian Hu. Reinforce++: A simple and efficient approach for aligning large language models. *arXiv* preprint arXiv:2501.03262, 2025.
  - Dongyoon Hwang, Hojoon Lee, Jaegul Choo, Dongmin Park, and Jongho Park. Can large language models develop strategic reasoning? post-training insights from learning chess. *arXiv* preprint *arXiv*:2507.00726, 2025.
  - Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024.
  - Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? *arXiv preprint arXiv:2310.06770*, 2023.
  - Stefan-Meyer Kahlen. Description of the universal chess interface (uci), 2004. URL https://www.wbec-ridderkerk.nl/html/UCIProtocol.html.
  - Mehran Kazemi, Bahare Fatemi, Hritik Bansal, John Palowitch, Chrysovalantis Anastasiou, Sanket Vaibhav Mehta, Lalit K Jain, Virginia Aglietti, Disha Jindal, Peter Chen, et al. Big-bench extra hard. *arXiv preprint arXiv:2502.19187*, 2025.
  - Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*, pp. 611–626, 2023.
  - Andrew Lee, Bo Chang Antonio Gulli, Bovard Doerschuk-Tiberi Bob Fraser, Chris Prichard Chad Woodford, Bob Fraser, et al. Chess text input. https://www.kaggle.com/benchmarks/kaggle/chess-text, 2025. Google DeepMind, Google Cloud, Kaggle.
  - Bill Yuchen Lin, Ronan Le Bras, Kyle Richardson, Ashish Sabharwal, Radha Poovendran, Peter Clark, and Yejin Choi. Zebralogic: On the scaling limits of llms for logical reasoning. *arXiv* preprint arXiv:2502.01100, 2025.
  - Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024a.
  - Siyao Liu, He Zhu, Jerry Liu, Shulin Xin, Aoyan Li, Rui Long, Li Chen, Jack Yang, Jinxiang Xia, ZY Peng, et al. Fullstack bench: Evaluating llms as full stack coder. *arXiv e-prints*, pp. arXiv–2412, 2024b.
  - Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*, 2025.
  - MAA. American invitational mathematics examination aime 2025, 2025. URL https://maa.org/math-competitions/american-invitational-mathematics-examination-aime.
    - Reid McIlroy-Young, Siddhartha Sen, Jon Kleinberg, and Ashton Anderson. Aligning superhuman ai with human behavior: Chess as a model system. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 1677–1687, 2020.
- OpenAI. Hello GPT-4o. Technical report, OpenAI, 2024. URL https://openai.com/index/hello-gpt-4o/.
- OpenAI. Introducing GPT-4.1 in the API. Technical report, OpenAI, 2025a. URL https://openai.com/index/gpt-4-1/.
  - OpenAI. Introducing-O3-and-O4-Mini. Technical report, OpenAI, 2025b. URL https://openai.com/index/introducing-o3-and-o4-mini/.

- rec.games.chess. Standard: Portable game notation specification and implementation guide, 1994. URL http://www.saremba.de/chessgml/standards/pgn/pgn-complete.htm.
  - Anian Ruoss, Grégoire Delétang, Sourabh Medapati, Jordi Grau-Moya, Li K Wenliang, Elliot Catt, John Reid, Cannada A Lewis, Joel Veness, and Tim Genewein. Amortized planning with large-scale transformers: A case study on chess. *Advances in Neural Information Processing Systems*, 37:65765–65790, 2024.
  - ByteDance Seed, Jiaze Chen, Tiantian Fan, Xin Liu, Lingjun Liu, Zhiqi Lin, Mingxuan Wang, Chengyi Wang, Xiangpeng Wei, Wenyuan Xu, et al. Seed1. 5-thinking: Advancing superb reasoning models with reinforcement learning. *arXiv* preprint arXiv:2504.13914, 2025.
  - Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv* preprint arXiv:2402.03300, 2024.
  - Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. In *Proceedings of the Twentieth European Conference on Computer Systems*, pp. 1279–1297, 2025.
  - Zayne Rea Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. Musr: Testing the limits of chain-of-thought with multistep soft reasoning. In *The Twelfth International Conference on Learning Representations*.
  - Stockfish Development Team. Stockfish 8. https://stockfishchess.org/, 2016. Open source chess engine.
  - Shu Wang, Lei Ji, Renxi Wang, Wenxiao Zhao, Haokun Liu, Yifan Hou, and Ying Nian Wu. Explore the reasoning capability of llms in the chess testbed. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pp. 611–622, 2025.
  - An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
  - Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025.
  - Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyan Luo, Zhangchi Feng, and Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. *arXiv* preprint arXiv:2403.13372, 2024.
  - Da-Wei Zhou, Hai-Long Sun, Jingyi Ning, Han-Jia Ye, and De-Chuan Zhan. Continual learning with pre-trained models: A survey. In *IJCAI*, 2024.
  - Terry Yue Zhuo, Vu Minh Chien, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widyasari, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, et al. Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions. In *The Thirteenth International Conference on Learning Representations*.

#### APPENDIX The Usage of Large Language Models **B** More Implementation Details B.4 Difference Between Move Selection and Real Chess Competition . . . . . . . . . C Glicko Rating System & Competition Sampling Algorithm **D** Post-training Details **Fine-Grained Evaluation Dataset Construction Additional Results G** Analysis

In the Appendix, we provide detailed experimental settings, mathematical proofs for the competition sampling algorithm, post-training dataset construction and training details, and additional experimental results and analysis.

#### A THE USAGE OF LARGE LANGUAGE MODELS

We used large language models as assistant tools for two specific purposes in this work:

- Paper polishing: LLMs were employed to improve the clarity, grammar, and overall readability of the manuscript text.
- Table formatting: LLMs assisted in enhancing the visual presentation and formatting of tables to improve readability.

The use of LLMs was limited strictly to these two auxiliary functions. LLMs did not contribute to research ideation, experimental design, data analysis, result interpretation, or the generation of scientific content. All research ideas, methodologies, findings, and conclusions are entirely the work of the authors.

# B MORE IMPLEMENTATION DETAILS

#### B.1 EVALUATED MODELS

Our evaluated models are shown in Table 5.

Table 5: Large Language Models Evaluated in ChessArena

<b>Model Family</b>	Model Name	Type	Thinking
	GPT-4.1 (OpenAI, 2025a)	Proprietary	×
OpenAI	GPT-40 (OpenAI, 2024)	Proprietary	×
	O3 (2025-04-16) (OpenAI, 2025b)	Proprietary	$\checkmark$
	DeepSeek-V3 (0324) (Liu et al., 2024a)	Open Source	×
DeepSeek	DeepSeek-V3.1	Open Source	×
	DeepSeek-R1 (0120) (Guo et al., 2025)	Open Source	✓
	Doubao-1.5-Pro-32K	Proprietary	×
DritaDonas	Doubao-1.5-Lite-32K	Proprietary	×
ByteDance	Doubao-1.5-Thinking-Pro	Proprietary	$\checkmark$
	Doubao-Seed-1.6-Thinking (Seed et al., 2025)	Proprietary	$\checkmark$
Alibaba	Qwen3-235B-A22B (0514) (Yang et al., 2025)	Open Source	×
Alibaba	Qwen3-8B (0514)	Open Source	×
Anthropic	Claude-3.7-Sonnet (Anthropic, 2025)	Proprietary	×
Google	Gemini-2.5-Pro (Google DeepMind, 2025)	Proprietary	$\checkmark$
	Qwen3-8B-SFT-Stage1	Open Source	×
Ours	Qwen3-8B-SFT-Stage2	Open Source	×
	Qwen3-8B-Chess (SFT+RL)	Open Source	×

#### **B.2** PROMPT TEMPLATES

There are the prompt templates for chess competitions, designed for various play modes. Blitz and Standard allow the model to think, as shown in Figure 2. Bullet expects the model to output the answer directly without thinking, as shown in Figure 3. Blindfold is another mode, where the model is expected to reconstruct the board from the conversation history and play accordingly. We record the player's and opponent's moves in the conversation history. For details, please refer to Figure 4. For basic understanding, the prompt templates are shown in Figure 5. The prompt templates for move selection remain consistent with those of each play mode. The prompt template of puzzle solving is the same as Blitz/Standard play mode prompt template.

859

860

```
811
812
813
          Blitz/Standard prompt template
814
                 System:
815
                 You are an expert chess player. You are playing a game
816
                 of chess. You are playing as {White_or_Black}.
817
                 You must thoroughly analyze the position and play with
818
                 utmost caution. When you have the advantage, press it
819
                 relentlessly and aim for a swift checkmate. Carefully
                 evaluate every move to eliminate any chance of a
820
                 counterplay or draw by your opponent.
821
                 When at a disadvantage, strive to turn the tide and
822
                 win if possible. If victory is unattainable, exhaust
823
                 all possible means to force a draw.
824
                 Meticulously analyze legal moves, then select the
                 absolute best one. You need to determine whether
825
                 you are playing as Black or White. Then, you need to
826
                 observe the positions of your pieces and choose one
                 of your own pieces to move; make sure that your move
828
                 follows the rules of chess.
829
                 Considering the long-term strategy and short-term
                 tactic.Analyze the position carefully.You may think
830
                 through the position and consider multiple candidate
831
                 moves.
832
                 When you have decided on your final move, output it in
833
                 UCI notation (e.g., 'e2e4', 'g8f6' , 'e7e8q') in the
834
                 following format:
835
                 <answer>
836
                 <move>
837
838
                 </answer>
839
                 For example:
840
                 e2e4
841
843
                 Reminder of chess rules:
844
                    - Bishops move diagonally.
845
                    - Rooks move horizontally or vertically.
                    - Knights jump in an L-shape.
846
                    - Queens combine rook and bishop movement.
847
                    - Kings move one square in any direction.
848
                    - Pawns move forward, capture diagonally, and can
849
                 promote.
850
                 You can think and reason as much as you want (step by
851
                 step), but your final move must be formatted exactly
852
                 as shown above.
                 User:
853
                 The current FEN: {<fen>}
854
                 (Optional) Legal moves in UCI notation: {<uci_1>
855
                 <uci_2> ...<uci_L>}.
856
                 What is the best move?
857
858
```

Figure 2: Input prompt format for Blitz and Standard chess competition. Whether to provide legal moves is optional.

882 883

884 885

886

887

888

889

890

891

892

893

894

895

897

898

899900901

902

903 904

905

906907908

909

910

911

912

913914915

916

```
864
          Bullet prompt template
865
866
                 System:
867
                 You are an expert chess player. You are playing a game
                 of chess. You are playing as White/Black.
868
                 Same as Blitz. Omitted.
869
                 You must give me your answer directly without using
870
                 any other words. I will not accept your answer if there
871
                 are any other words. Only output your move content. Your
872
                 final move must be formatted exactly as shown above.
873
874
                 User:
875
                 The current FEN: {<fen>}
                 (Optional) Legal moves in UCI notation: {<uci_1>
876
                 <uci_2> ...<uci_L>}.
877
                 What is the best move?
878
879
```

Figure 3: Input prompt format for Bullet chess competition. Whether to provide legal moves is optional. Thinking is forbidden.

```
Blindfold prompt template
       System:
       You are an expert chess player. You are playing a game
       of chess. You are playing as White/Black.
       We have the move history of you
       and your opponent. You must reconstruct the game and
       analyze the best move on the chessboard.
       Same as Blitz. Omitted.
       You can think and reason as much as you want (step by
       step), but your final move must be formatted exactly
       as shown above.
       User(First Turn):
       This is the beginning of the game.
       Assistant(First Turn):
       <move_1>
       User:
       Your opponent's last move is <move_2>.
       Assitant:
       <move_3>
       Multi-Turns
       User: Your opponent's last move is {<move_k>}.
       (Optional) Legal moves in UCI notation: {<uci_1>
       <uci_2> ...<uci_L>}.
       What is the best move?
```

Figure 4: Input prompt format for Blindfold chess competition. Whether to provide legal moves is optional. This is a multi-round conversation template. LLMs should reconstruct the chessboard from the conversation history.

967 968

```
920
921
         Basic understanding evaluation prompt template
922
923
               System:
924
               You are an expert chess player. I need you to help
925
               me model a chessboard. The specific steps are as
               follows:
926
               I will provide you with a FEN string representing
927
               the current board state, and then give you a
928
               position. You need to identify the piece at that
929
               position from the FEN and output all legal moves
930
               for that piece.
931
               You must carefully analyze the board, consider
932
               the rules of chess, and provide the final answer.
933
934
               Your answer should be format as follows (output a
935
               json):
936
               ```json
937
                   "piece": <piece symbol>,
938
                   "legal moves": [<list of legal moves>]
939
940
941
942
               For example:
943
               FEN: rnbqkbnr/pppppppppppp/8/8/8/8/PPPPPPPPPPRNBQKBNR
944
               w KQkq - 0 1
945
               Position:g1
946
               Answer:
947
                ``json
948
949
                   "piece":
                              "N",
                   "legal moves":
                                   ["g1h3", "g1f3"]
951
952
               Note:
953
               If the given position has no piece, directly
954
               output empty(i.e.,None), and the corresponding
955
               legal moves should also be empty(i.e.,[]).
956
               When it's White's turn to move, if the position
957
               contains a Black piece, you should identify the
958
               piece, but its legal moves must be empty (and
959
               vice versa for Black's turn).
960
               You can think and reason as much as you want (step
961
               by step), but your final answer must be formatted
962
               exactly as shown above.
963
               User:
964
               Current board position in FEN notation: {<fen>}
               Position:{<pos>}
965
966
```

Figure 5: Input prompt format for basic understanding

#### **B.3** TERMINATION CONDITIONS

Following official chess rules, our games terminate under these conditions:

- Checkmate: A player delivers checkmate, winning the game.
- **Forfeit**: A player fails to generate a legal move after multiple attempts (indicating either instruction-following deficits or board analysis failures), awarding victory to the opponent. We will give an instruction if an LLM fails to give a legal move. We allow a player to retry 5 times.
- Stalemate: Draw declared when the active player has no legal moves but is not in check.
- **Insufficient material**: Draw due to neither player having sufficient pieces to force checkmate (e.g., king vs. king).
- **Fivefold repetition**: Draw triggered by the same position recurring five times.
- 75-move rule: Draw if 75 consecutive moves occur without pawn advances or captures.
- Move limit: Draw if the total move count exceeds the maximum move count. We set it to 200 moves.

#### **B.4** CHESS NOTATION

**Board Representation**. We adopt the Forsyth-Edwards Notation (FEN) rec.games.chess (1994) as our chessboard representation Standard. FEN is a widely recognized notation system that encodes a chess position into six space-delimited fields, comprehensively capturing the game state (e.g., piece placement, active color, castling rights, en passant targets, half move clock, and full move number). This notation is supported by the Python-Chess library and provides LLMs with an unambiguous, machine-parseable representation of board states, where each unique chess position maps to a distinct FEN string.

**Move Representation**. For move encoding, we implement the Universal Chess Interface (UCI) Kahlen (2004) Standard, which specifies moves in coordinate notation (e.g., "e2e4" for pawn advance). UCI's start-to-end positional format ensures deterministic move interpretation. Additionally, we maintain compatibility with Standard Algebraic Notation (SAN) to accommodate alternative LLM outputs. Our system automatically normalizes all move representations into a canonical form, enabling robust analysis regardless of the LLM's native output format.

In ChessArena, we first prefer to have models output UCI notation moves. If UCI notation moves cannot be extracted, we will extract SAN moves. We support both move notations. Regarding chessboard representations, we know that besides FEN representation, there is also Portable Game Notation (PGN) rec.games.chess (1994) representation. However, PGN representation shows the move history of a game and cannot directly reveal the piece arrangement on a board, so we use FEN representation, which is much more direct and clear for models. In basic understanding experiments, GPT-4.1, Claude-3-7-sonnet, and Doubao-Seed-1-6-thinking all showed a high piece match rate, precision, and recall, indicating they have understanding capabilities for FEN board representation, but their actual chess gameplay performance still has considerable space for improvement.

# B.5 DIFFERENCE BETWEEN MOVE SELECTION AND REAL CHESS COMPETITION

Move selection evaluation results show an overall consistent trend with the ChessArena Leaderboard. ChessArena competition is more complex than single move selection. Additionally, LLMs also have opportunities to adjust themselves. LLMs must try their best in a pressure situation. Move selection offers a straightforward and efficient method for assessing an LLM's chess strategic reasoning. In contrast, ChessArena competition provides a more accurate and engaging evaluation by requiring models to participate in extended game sessions.

#### C GLICKO RATING SYSTEM & COMPETITION SAMPLING ALGORITHM

#### C.1 GLICKO RATING SYSTEM

In the Glicko ranking system, each player is assigned two values: the rating r and the rating deviation RD. In the original Glicko paper Glickman (1995), these values are updated after a certain period; In

our scenario, we assume they are updated after each competition. The updated values, denoted as r' and RD', are given by the following formulas:

$$r' = r + \frac{q}{\frac{1}{RD^2} + \frac{1}{d^2}} g(RD_o) \left( s_o - E(s \mid r, r_o, RD_o) \right)$$
 (1)

$$RD' = \sqrt{\left(\frac{1}{RD^2} + \frac{1}{d^2}\right)^{-1}}$$
 (2)

where

$$q = \frac{\ln 10}{400} \approx 0.0057565 \tag{3}$$

$$g(RD) = \frac{1}{\sqrt{1 + \frac{3q^2RD^2}{\pi^2}}}$$
(4)

$$E(s \mid r, r_o, RD_o) = \frac{1}{1 + 10^{-g(RD_o)(r - r_o)/400}}$$
 (5)

$$d^{2} = \left(q^{2} \left(g(RD_{o})\right)^{2} E(s \mid r, r_{o}, RD_{o}) \left(1 - E(s \mid r, r_{o}, RD_{o})\right)\right)^{-1}$$
(6)

where s represents the competition result (i.e., 1 for a win, 0.5 for a draw, and 0 for a loss),  $r_o$  and  $RD_o$  denote the opponent's rating and rating deviation, respectively. These calculations are performed for each player participating in the rating period.

#### C.2 COMPETITION SAMPLING

**System Objectives and Optimization Criteria** In equation equation 2, we can see that RD will definitely decrease as matches progress, indicating that a player's rating becomes increasingly reliable. So under what conditions does a player's RD decay faster, enabling the player to converge most quickly? We provide a mathematical analysis in this section. The core goal of this matching system is to accelerate the convergence rate of player ratings, specifically by maximizing the reduction rate of rating deviation (RD). The optimization objective function is defined as:

$$\arg\max_{r_i, r_j, RD_i, RD_j} \Delta RD_i + \Delta RD_j \tag{7}$$

where  $\Delta RD_i$  and  $\Delta RD_j$  represent the changes in rating deviation for player i and player j after matching, respectively.

**Mathematical Derivation Process** According to the update rules of the Glicko-1 system, the change in rating deviation can be expressed as:

$$\Delta RD_i = RD_i - \sqrt{\left(\frac{1}{RD_i^2} + \frac{1}{d_i^2}\right)^{-1}} = RD_i - \sqrt{\frac{RD_i^2}{1 + \frac{RD_i^2}{d_i^2}}}$$
(8)

$$\Delta RD_j = RD_j - \sqrt{\left(\frac{1}{RD_j^2} + \frac{1}{d_j^2}\right)^{-1}} = RD_j - \sqrt{\frac{RD_j^2}{1 + \frac{RD_j^2}{d_j^2}}}$$
(9)

To maximize  $\Delta RD_i + \Delta RD_j$ , we need to minimize  $d_i^2$  and  $d_i^2$ . According to equation 6:

$$d_i^2 = \left(q^2 \left(g(RD_j)\right)^2 E(s \mid r_i, r_j, RD_j) \left(1 - E(s \mid r_i, r_j, RD_j)\right)\right)^{-1}$$
(10)

Therefore, minimizing  $d_i^2$  and  $d_j^2$  is equivalent to maximizing:  $q^2g(\mathrm{RD}_i)^2E_i(1-E_i)$  and  $q^2g(\mathrm{RD}_j)^2E_j(1-E_j)$ , where:

$$E_i = \frac{1}{1 + 10^{-g(\text{RD}_j)(r_i - r_j)/400}}, \quad E_i = 1 - E_j$$
(11)

$$g(\text{RD}) = \frac{1}{\sqrt{1 + \frac{3q^2\text{RD}^2}{\pi^2}}}, \quad q = \frac{\ln 10}{400} \approx 0.0057565$$
 (12)

Based on the above derivation, the optimization objective can be transformed into:

$$\operatorname{argmax}_{r_i, r_j, \operatorname{RD}_i, \operatorname{RD}_j} E_i (1 - E_i) \left[ g(\operatorname{RD}_i)^2 + g(\operatorname{RD}_j)^2 \right] \tag{13}$$

**Key Conclusions** From equation (7), we can draw the following important conclusions: When  $r_i = r_j$  (i.e., the two players have the same rating), then  $E_i = E_j = 0.5$ , at which point  $E_i(1 - E_i)$  reaches its maximum value of 0.25. Meanwhile, g(RD) is a decreasing function of RD, meaning that smaller RD results in larger g(RD).

Therefore, the optimal matching strategy is:

- Prioritize matching players with similar ratings  $(r_i \approx r_j)$
- Under the premise of similar ratings, select players with smaller rating deviations (RD)

This strategy ensures maximum information gain for both players in the competition, thereby accelerating rating convergence.

**Algorithm Premises** A minimum rating deviation threshold min\_RD (a hyperparameter) is set. When a player's RD  $\leq$  min\_RD, their rating deviation no longer decreases.

#### **Competition Sampling Process**

- A player initiates a match request, and the system records their current rating r and rating deviation RD
- 2. The system searches for potential opponents in the match pool and calculates the matching score:

$$score(i, j) = E_i(1 - E_i) \left[ g(RD_i)^2 + g(RD_i)^2 \right]$$

- 3. The opponent with the highest matching score is prioritized
- 4. For players with high RD, the system prioritizes matching them with opponents who have low RD and similar ratings
- 5. After the opponent accepts the match, the match begins
- 6. After the match, both players' r and RD are updated based on the results

#### **ChessArena Matching System Variants** The system supports two startup modes:

- 1. Random startup mode:
  - (a) A player is randomly selected from the player pool
  - (b) The selected player automatically initiates a match request
  - (c) Steps 2-6 of the Competition Sampling process are executed
- 2. Specified startup mode:
  - (a) An initial player is specified by a human
  - (b) The specified player initiates a match request
  - (c) Steps 2-6 of the Competition Sampling process are executed

# D POST-TRAINING DETAILS

#### D.1 SFT DATA COLLECTION

**ChessGPT** ChessGPT Feng et al. (2023) has open-sourced a text pre-training dataset and a post-training SFT dataset related to chess. These datasets include conversational data about chess, covering topics such as basic rules and tactical discussions. We sampled chess-related portions (GPT-4-Chess, Chess-Forums, and Chess-Modeling) from this dataset as part of our SFT data.

**Distillation** We distilled data from non-thinking models: GPT-4.1, DeepSeek-V3, Qwen3-235b-a22b, and Claude-3-7-Sonnet; Thinking models: Doubao-Seed-1-6-thinking and DeepSeek-R1. The input prompt format resembles the Blitz play mode prompt template, and the output includes the model's analysis of the chessboard and the final move selection. We used Stockfish to ensure the quality of the distilled data. We only retained data where the final move was among the top three moves analyzed by Stockfish. We filtered the data whose response length is less than 100. The characteristics of the distilled dataset are shown in the Table 6. We use the tiktoken(cl100k-base)<sup>4</sup> tokenizer to estimate the length of the distilled dataset.

What's more, in ChessArena, LLMs may initially fail to provide a legal move in the first round but correct themselves in subsequent attempts. We also extract such data for training, as it helps the model learn multi-turn correction capabilities. There are 652 samples in the multi-turn correction dataset.

Table 6: Characteristics of the Main Distilled Dataset(excluding multi-turn correction data)

Туре	Count	Prompt Length (avg.)	Resp. Len. (avg.)	TOP1	TOP2-3
Non-Thinking	21,278	575	527	10,273	11,005
Thinking	3,399	582	5,014	1,862	1,537
Multi-Round Correction	652	1693	343	276	376

Table 7: SFT data summary

Da	itaset	Count	Description	Prompt Length (avg.)	Resp. Len. (avg.)
ChessGPT(Stage1)	GPT4-Chess	3908	Chess-related synthesized data from GPT-4	41	38
, ,	Chess Forums	5395	Chess-related dialogues data from online platform	245	178
	Chess Modeling	3000	chessboard understanding data like PGN to FEN, FEN to UCI et al.	116	65
Distilled Data(Stage2)	Move Selection Multi-Round Correction	21278 652	Distilled single turn data Multi-Round Correction data	524 1693	527 343

#### D.2 RL DATA COLLECTION

Theoretically, our RL training data is virtually unlimited, as only a single chessboard state is required to conduct RL training. Accordingly, we extracted board state data from the Lichess database and constructed our RL training dataset. All board states corresponding to FEN positions used in Fine-Grained studies were filtered out, resulting in a final set of 56,000 training samples. We ensured a balanced distribution of board states across the opening games, middle games, and end games to facilitate comprehensive learning by the model. Although experiments with larger datasets were attempted, no further improvement in model performance was observed.

#### D.3 REWARD

Our reward function consists of three components: 1) Format reward; 2) Legal move reward; 3) Top move reward. Format reward guides the model to follow certain formats; Legal move reward guides

<sup>4</sup>https://cookbook.openai.com/examples/how\_to\_count\_tokens\_with\_tiktoken

the model to infer legal moves from the chessboard; Top move reward guides the model to acquire chess strategy reasoning capabilities. Top moves are analyzed by Stockfish, and we pre-process to obtain top moves before training to avoid Stockfish consuming excessive CPU resources during training.

**Format Reward** If the model's output follows the specified format (i.e., it is contained within the prescribed block), the format reward  $reward_f = 1$ ; otherwise,  $reward_f = 0$ .

**Legal Move Reward** If the model's predicted move is among the legal moves for the current board position, the legal moves reward  $reward_l = 1$ ; otherwise,  $reward_l = 0$ .

**Top Move Reward** If the model's predicted move matches one of the top moves pre-analyzed by Stockfish, the top moves reward  $reward_t = 1$ ; otherwise,  $reward_t = 0$ .

The final reward is calculated as a weighted sum of these three rewards:

```
Reward = \epsilon_f \times reward_f + \epsilon_l \times reward_l + \epsilon_t \times reward_t
```

where  $\epsilon_f$ ,  $\epsilon_l$ , and  $\epsilon_t$  are the corresponding weight coefficients.

#### D.4 TRAINING HYPER-PARAMETERS

**Supervised Fine-tuning** We train our models using the LlamaFactory Zheng et al. (2024) framework. Our hyper-parameters are shown in Table 8. In our training Qwen3-8B-Chess process, to utilize more data (though mixing thinking and non-thinking data for training may cause issues), we use the non-thinking distilled dataset to train our Qwen3-8B-Chess. So, our Qwen3-8B-Chess is trained in non-thinking mode. For the multi-turn correction dataset, we only train our model on the final turn response.

**Reinforcement Learning** We train our models using the verl (Sheng et al., 2025) framework. All of our training experiments are finished on 8 NVIDIA H800 80GB GPUs. A single training experiment takes approximately 60 hours. To enhance model performance, we incorporated methodologies from DR. GRPO Liu et al. (2025), Reinforce++ Hu (2025), and DAPO Yu et al. (2025). It was observed that the integration of these techniques contributed to improved model performance.

# E FINE-GRAINED EVALUATION DATASET CONSTRUCTION

We extracted the FEN of board states that actually occurred, thereby minimizing the risk of data contamination to the greatest extent possible.

ChessBoard Extraction In each game of ChessArena, numerous board states are generated (averaging 40 chessboard states per game). However, many of these board states may be duplicated, particularly in the opening phase. We extract distinct board states from this data while ensuring a balanced distribution across the opening-game, middle-game, and end-game stages (with the middle-game slightly outnumbering both the opening and endgame phases). A total of 79,441 FENs are collected. FENs already present in the training dataset are removed to prevent data contamination of the trained models. After this filtering, 57,511 FENs remain. These retained chessboard states can be utilized to construct the subsequent evaluation set.

**Basic Understanding** The FEN data acquired from the Chessboard Extraction step is subsequently utilized for further dataset construction. In the basic understanding evaluation dataset construction process, each position is sampled according to the following distribution: with an 85% probability, a position containing one of the player's own pieces is selected; with a 7% probability, a position from the opponent's pieces is chosen; and with an 8% probability, an empty square is selected. Each board FEN is used to construct one basic understanding data instance, resulting in a total of 57,511 instances. To facilitate efficient evaluation, a subset of 200 instances is selected for assessment. Consistent experimental outcomes have been observed across subsets of varying sizes, including 200, 500, 1000, and larger.

Table 8: Training hyper-parameters for Qwen3-8B-Chess.

Supervised Fine-tun	ing Stage 1	Reinforcement Learning			
Hyperparameter	Value	Hyperparameter	Value		
Training Steps Optimizer Learning Rate Global Batch Size Epochs Warmup Ratio Lr scheduler type Max tokens	1038 AdamW 5e-6 32 3 0.1 Cosine 4096	Training Configuration Training Steps Optimizer Learning Rate Global Batch Size Mini Batch Size Epochs Max Tokens	1750 AdamW 1e-6 128 64 4 4096		
Supervised Fine-tun Hyperparameter	ing Stage 2 Value	GRPO Configuration Number of Rollouts Rollout Temperature	8 1.0		
Training Steps Optimizer Learning Rate Global Batch Size Epochs Warmup Ratio Lr scheduler type Max tokens	2130 AdamW 5e-6 32 3 0.1 Cosine 4096	Rollout Topp  KL Loss Coefficient Entropy Coefficient Clip High Clip Low  Reward Configuration $\epsilon_f$	1.0 0 0 0.28 0.2		

**Move Selection** Following the acquisition of the FEN board state data from the initial step, theoretically, all instances could serve as evaluation data for move selection. However, to ensure both accuracy and efficiency in the evaluation process, we performed a rollout using Qwen3-8B-Chess on each data instance. This procedure ensured a balanced distribution of easy, medium, and difficult problems. From this processed set, 1,000 instances were selected to constitute the evaluation dataset for the move selection fine-grained experiment. For Blindfold play mode, we use the real move history that happened in ChessArena as the conversation history. However, it is important to note that a comprehensive evaluation of models' chess capabilities can be achieved without relying on our provided FEN board representations, for instance, by extracting board states from the Lichess database. Furthermore, given the vast search space of chess and the virtually infinite number of possible board configurations, the risk of data contamination is negligible.

**Puzzle Solving** We retrieved puzzle data from the Lichess database and randomly selected 1,008 samples to form the evaluation set. Subsequently, we partitioned the data into seven segments based on Elo rating intervals of 400 points, with each segment containing exactly 144 puzzle instances. The dataset is sufficiently large to allow discernible observation of the differences in puzzle-solving capabilities among the LLMs.

**Datasets Distribution** As mentioned before, we performed a preemptive rollout procedure using Qwen3-8B-Chess on the move selection evaluation dataset to categorize the difficulty levels. The rollout was conducted 8 times per instance with hyper-parameters set to temperature 1.0 and top-p 0.95. We defined an instance as easy if Qwen3-8B-Chess selected a top-3 move in all 6-8 rollouts, medium if it did so in 3-5 rollouts, and hard if it never selected a top-3 move. The overall difficulty distribution and statistics of fine-grained evaluation are illustrated in Table 9.

Table 9: Statistics of fine-grained evaluation datasets.

Task	Category	Count	Percentage (%)
	Normal Positions	144	72.0
Pagia Undangtanding	Empty Positions	19	9.5
<b>Basic Understanding</b>	Opponent Positions	37	18.5
	Total	200	100.0
	Phase - Early (0-20)	241	24.1
	Phase - Middle (20-60)	472	47.2
	Phase - Late (>60)	287	28.7
<b>Move Selection</b>	Difficulty - Easy	215	21.5
	Difficulty - Medium	187	18.7
	Difficulty - Hard	598	59.8
	Total	1000	100.0
	Mate	308	30.6
	Cruising	405	40.2
Puzzle Solving	Advantage	275	27.3
	Others	20	2.0
	Total	1008	100.0

# F ADDITIONAL RESULTS

#### F.1 WHOLE LEADERBOARD

Our complete rating leaderboard is presented in Table 10, which includes additional models that do not affect the conclusion analysis, as well as models with RD values exceeding 100. Table 11 reports secondary metrics from ChessArena competitions, encompassing win-loss number, instruction following metrics (parsing errors, forbidden moves, legal moves), and move quality measures (top moves).

# F.2 MOVE HISTORY AFFECTION

To explore the performance difference of the model when provided with or without move history, we conducted an extra evaluation. In Blitz mode, the model's performance difference is when provided with or without move history, as can be seen from Table 14. Providing or not providing move history does not significantly affect the model's performance, with evaluation metrics fluctuating within the range of 1% to 5%. As we previously mentioned, providing move history primarily serves to allow the model to access historical information, thereby helping to avoid fivefold repetition draw. Without legal move constraints, PGN notation improves model performance, likely because PGN notation aligns better with the model's training corpus. The model is more familiar with this method of representing move history. In our ChessArena, as long as both sides are provided with consistent information, it does not introduce unfairness that might arise from differences in move history representation.

#### G ANALYSIS

#### G.1 BLINDFOLD ANALYSIS

We discussed the different behaviour between thinking and non-thinking LLMs in Blindfold chess games. Non-thinking models frequently exhibit laziness and shortcut-taking behavior, while thinking models genuinely attempt to reconstruct the board, but this task proves excessively difficult. Table 15 is a case study of different models' responses in a Blindfold chess game. Thinking models such as Doubao-Seed-1-6-thinking always try to reconstruct the chessboard state, but it's too difficult. GPT-4.1 often takes shortcuts by directly following the last move made on its own side in the conversation history. Claude-3-7-Sonnet does not explicitly reconstruct the chessboard state but still retains some

Table 10: Whole Rating Leaderboard

Rank	Model	Туре	Legal Moves	Rating	RD	Interval	Game
1	Maia-1100	-	l ×	2220	82	(2058, 2382)	44
2	O3	Standard	×	1948	78	(1793, 2101)	28
3	Doubao-Seed-1-6-Thinking	Standard	✓	1830	50	(1729, 1929)	60
4	Gemini-2.5-Pro	Standard	✓	1819	81	(1659, 1979)	18
5	Qwen3-8B-Chess	Blitz	✓	1776	93	(1593, 1959)	16
6	Doubao-Seed-1-6-Thinking	Standard	×	1743	66	(1612, 1873)	36
7	GPT-4.1	Blindfold	✓	1699	50	(1601, 1797)	60
8	Doubao-Seed-1-6-Thinking	Blindfold	✓	1687	73	(1542, 1831)	24
9	GPT-4.1	Blitz	✓	1686	50	(1588, 1784)	182
10	Claude-3-7-Sonnet	Blitz	✓	1654	50	(1555, 1751)	74
11	O3	Blindfold	×	1636	115	(1409, 1861)	16
12	Claude-3-7-Sonnet	Blindfold	✓	1625	66	(1493, 1756)	30
13	GPT-4.1	Blitz	×	1623	50	(1525, 1721)	106
14	Gemini-2.5-Pro	Standard	×	1616	74	(1469, 1762)	28
15	Qwen3-8B-SFT	Blitz	✓	1612	56	(1501, 1721)	40
16	Doubao-1-5-Thinking-Pro	Standard	✓	1598	63	(1473, 1723)	32
17	Claude-3-7-Sonnet	Blindfold	×	1588	72	(1445, 1729)	28
18	DeepSeek-R1	Standard	✓	1587	50	(1487, 1686)	54
19	GPT-4.1	Bullet	✓	1583	50	(1485, 1681)	54
20	GPT-40	Bullet	✓	1568	80	(1409, 1725)	28
21	DeepSeek-V3	Blitz	✓	1553	50	(1454, 1650)	174
22	Doubao-1-5-Pro	Blitz	✓	1539	58	(1423, 1654)	42
23	Random Player	-	✓	1524	50	(1425, 1621)	284
24	Doubao-1-5-Lite	Blitz	✓	1509	78	(1354, 1662)	28
25	Qwen3-235B-A22B	Blitz	✓	1483	50	(1385, 1581)	146
26	DeepSeek-V3	Blitz	×	1482	58	(1367, 1597)	48
27	Qwen3-8B-Chess	Blitz	×	1472	88	(1297, 1645)	16
28	Claude-3-7-Sonnet	Bullet	✓	1452	59	(1334, 1569)	34
29	DeepSeek-V3	Blindfold	✓	1437	75	(1290, 1584)	24
30	GPT-40	Blindfold	✓	1402	81	(1241, 1561)	20
31	DeepSeek-V3	Bullet	✓	1382	80	(1224, 1540)	22
32	Qwen3-235B-A22B	Bullet	✓	1369	54	(1261, 1476)	46
33	Qwen3-8B	Blitz	✓	1335	65	(1205, 1463)	32
34	Doubao-Seed-1-6-Thinking	Blindfold	×	1276	90	(1097, 1453)	24
35	GPT-4.1	Blindfold	×	1237	160	(922, 1550)	8

al

ability to analyze the game situation. Overall, chess gameplay under blindfold play modes proves excessively challenging, making it difficult for models to reconstruct the chessboard state and analyze potential moves accurately. Non-thinking models demonstrate virtually no capability for genuine board reconstruction. Among thinking models, only the O3 model exhibits satisfactory performance. This highlights substantial room for improvement in models' multi-turn conversation memorization capabilities.

Table 11: ChessArena Competition Results and Performance Metrics. Mode: Play Modes(Blitz/Bullet/Standard/Blindfold); Legal: Whether legal moves were provided; Wins/Losses: Number of games won/lost; Parsing Err%: Invalid format rate; Illegal Mv%: Illegal move rate; Forbidden%: Illegal thought rate in Bullet play mode; Legal Mv%: Legal move rate; Top Mv%: Top move rate. Due to the existence of draws, the sum of wins and losses does not equal the total number of games played by the model.

Rank	Model	Mode	Legal	Wins	Losses	Parsing Err%	Illegal Mv%	Forbidden%	Legal Mv%	Top Mv%
1	Maia-1100	-	×	44	0	0	0	0	100	87.5
2	O3	Standard	×	15	13	48.7	1.5	0	45.1	78.6
3	Doubao-Seed-1-6-Thinking	Standard	✓	32	7	0.9	0.1	0	40.6	51.4
4	Gemini-2.5-Pro	Standard	✓	10	2	28.7	0.5	0	61.1	61.8
5	Qwen3-8B-Chess	Blitz	✓	7	0	0.2	0.2	0	99.6	44.3
6	Doubao-Seed-1-6-Thinking	Standard	×	15	17	13.9	7.2	0	61.1	51.4
7	GPT-4.1	Blindfold	✓	18	12	0.7	3.7	0	94.6	51.4
8	Doubao-Seed-1-6-Thinking	Blindfold	✓	6	1	1	1.3	0	96.6	55.1
9	GPT-4.1	Blitz	✓	54	44	4.7	1.5	0	87.5	53.4
10	Claude-3-7-Sonnet	Blitz	✓	13	18	0.3	1.8	0	97.1	52
11	O3	Blindfold	×	16	0	6.9	1.9	0	85	77.2
12	Claude-3-7-Sonnet	Blindfold	✓	4	6	0.1	1.3	0	98.6	53.7
13	GPT-4.1	Blitz	×	67	38	9.8	18.4	0	71.8	59.4
14	Gemini-2.5-Pro	Standard	×	16	12	37.3	9	0	50.7	73.6
15	Qwen3-8B-SFT	Blitz	✓	10	7	2.2	2.3	0	95.5	38.7
16	Doubao-1-5Thinking-Pro	Standard	✓	4	2	29.6	2.4	0	67.2	53.4
17	Claude-3-7-Sonnet	Blindfold	×	9	17	1	17.6	0	52.6	58.5
18	DeepSeek-R1	Standard	✓	10	6	30.7	1.4	0	67.7	51
19	GPT-4.1	Bullet	✓	8	6	12.5	1.3	0	78.6	45
20	GPT-4o	Bullet	✓	2	9	0	1	0	98.1	34.4
21	DeepSeek-V3	Blitz	✓	34	42	0.5	0.6	0	97.9	45.9
22	Doubao-1-5-Pro	Blitz	✓	10	6	0.1	5.9	0	93.6	32.2
23	Random Player	-	✓	0	3	0	0	0	100	40.3
24	Doubao-1.5-Lite	Blitz	✓	4	11	23.2	3.3	0	73.4	33.2
25	Qwen3-235B-A22B	Blitz	✓	27	41	8.4	3.6	0	87.8	39.3
26	DeepSeek-V3	Blitz	×	34	14	1.4	40.8	0	57.7	43.8
27	Qwen3-8B-Chess	Blitz	×	7	9	0	30.7	0	69.3	33.8
28	Claude-3-7-Sonnet	Bullet	✓	0	6	25.3	0.6	0	74.1	34.6
29	DeepSeek-V3	Blindfold	✓	6	6	4.4	9.4	0	86.2	33.7
30	GPT-40	Blindfold	✓	4	1	0	1.8	0	97.1	37.4
31	DeepSeek-V3	Bullet	✓	2	5	0	2.5	3.4	96.9	32.9
32	Qwen3-235B-A22B	Bullet	✓	6	13	0	2.7	0.5	96.8	35.3
33	Qwen3-8B	Blitz	✓	6	15	1.6	1.3	0	97.1	32.9
34	Doubao-Seed-1-6-Thinking	Blindfold	×	1	23	1.7	23.7	0	35.1	54
35	GPT-4.1	Blindfold	×	0	8	2.8	34.6	0	62.6	62.2

Table 12: Average Conversation Turn Count in Blindfold move selection (Without Legal Move Provision)

Model	Thinking	Successful Turn Count	Failed Turn Count
GPT-4.1	×	94	94
DeepSeek-V3	×	94	88
Qwen3-235B-A22B	×	96	86
Claude-3.7-Sonnet	×	92	94
DeepSeek-R1	<b>√</b>	72	112
Doubao-Seed-1.6-Thinking	$\checkmark$	70	112
Gemini-2.5-Pro	$\checkmark$	88	105
O3	$\checkmark$	88	129

Table 13: Puzzle Solving Accuracy: Blitz/Standard Prompt Template and Legal Moves not Provided

Model or Engine	Puzzle Solving Accuracy (%)								
	200-600	600-1000	1000-1400	1400-1800	1800-2200	2200-2600	2600-3000	Overall	
GPT-4.1	44.1	29.4	18.2	12.6	4.2	2.8	0.0	15.9	
Claude-3-7-Sonnet	18.2	16.1	4.9	4.2	5.6	1.4	0.0	7.2	
DeepSeek-V3	2.1	2.1	2.1	2.8	0.7	1.4	0.0	1.6	
DeepSeek-V3.1(Non-thinking)	9.8	7.7	4.9	2.1	2.8	1.4	0.7	4.2	
Qwen3-235B-A22B(Non-thinking)	16.7	12.5	7.2	4.5	5.0	4.2	0.0	7.1	
Qwen3-8B	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
Qwen3-8B-Chess(SFT+RL)	7.0	8.4	4.2	2.8	0.7	0.7	0.0	3.4	
O3	95.8	93.0	87.4	68.5	44.8	16.8	4.2	58.6	
Gemini-2.5-Pro	45.2	39.2	19.6	14.0	2.8	3.5	0.7	19.9	
Doubao-Seed-1-6-Thinking	21.0	22.4	13.3	4.9	4.9	0.7	0.0	9.6	
DeepSeek-R1	18.9	13.3	11.2	1.9	2.1	1.4	0.7	7.1	

Table 14: GPT-4.1 Performance in Blitz Mode: Impact of Move History

<b>Evaluation Condition</b>	LR (%)	TR (%)	
With Legal Move Constraints			
Without Move History	98.2	25.0	
With Move History(List of UCI)	97.2	29.0	
With Move History(PGN)	97.0	28.7	
Without Legal Move Constraints			
Without Move History	68.4	28.4	
With Move History(List of UCI)	70.2	26.2	
With Move History(PGN)	75.4	34.5	

Table 15: A Case Study of Blindfold chess game of different LLMs

# **Query Details:**

System prompt is omitted.

#### **Conversation History (Recent Moves):**

80+ Conversation Turns are omitted.

- **user**: The opponent's last move is b7c7.
- assistant: My move: e3e7
- **user**: The opponent's last move is c7c6.
- assistant: My move: e7c7
- user: The opponent's last move is c6c7.
- assistant: My move: e8e7
- user: The opponent's last move is c7b6.
- assistant: My move: f1h3
- **user** (**Final Query**): The opponent's last move is b6a6. What is the best Move? Think it step by step.

# Case 1: Doubao-Seed-1-6-thinking Failed to Reconstruct the Chessboard: It has fallen into the repeater problem.

```
Got it, let's try to figure out the current position first. I need to recap the move history to get the pieces on the board. First, I'm White, since the first move was g1f3 (White's knight). Let's list the key moves for
```

First, I'm White, since the first move was glf3 (White's knight). Let's list the key moves for both sides, focusing on piece locations:

```
White's moves (me): - Knights: Started on g1 and b1. g1f3 -> f3g5 -> g5f7 -> f7h8 -> h8f7 -> f
```

Case 2: GPT-4.1 takes a shortcut: It continues moving from the landing point(i.e., h3) of the last piece in the Query.

```
Table 15 (Continued): Blindfold Model Performance

Table 15 (Continued): Blindfold Model Performance

My move:

1517

1518

h3e6

1519
```

Case 3: Claude-3-7-Sonnet analyzes the chessboard and tries to give a better move.

I'll analyze the position carefully to find the best move.

**Step 1: Confirm Current Material and Position** - White's pieces: Based on moves, we have a bishop on h3 (f1h3), a rook/queen on e7 (e8e7, previously e3e7 -> e7c7 -> e8e7), and other pieces (knights, remaining pawns) whose positions need to be inferred. - Black's pieces: Only the king is confirmed on a6 (b6a6), and other pieces (like pawns on a6's original position) have been captured or moved. - Material advantage: White has at least a bishop + rook/queen vs. Black's lone king—clear advantage, goal is to checkmate.

Step 2: Evaluate Candidate Moves 1. e7a7: Directly attacks the Black king on a6. If e7 is a queen/rook, e7a7 would be check. But need to confirm if a7 is occupied—no Black pieces are known to be on a7, so this could be checkmate. 2. h3c8: Moves the bishop to c8, controlling b7 (a square the Black king might flee to if attacked). This prepares for a subsequent checkmate but is not immediate. 3. e7e6: Moves the e7 piece to e6, controlling d5 and f5—limits the king's potential escape to the center, but is slower than direct attack. 4. b4b5: If White has a pawn on b4, b4b5 would check the king (attacking a6), but there's no record of a b4 pawn in the move history, so this is risky.

**Step 3: Select the Best Move** e7a7 is the most efficient: it directly targets the Black king, and since there are no known Black pieces blocking a7, this move will put the king in check. If the king has no legal moves to escape (a6 can only move to a5, b5, or b6—b6 is empty but controlled by White's pieces), this becomes checkmate.

```
My move:
```

e7a7

1512

1520

1521 1522

1523

1525

1526

1527

1529

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541 1542

1543

1545 1546

1547 1548

1549

1550

1551

1552

1553

1554

1555

1556

1557

1558

1560

1561

1563

1564

1565

#### G.2 THE GENERALIZATION OF CHESS REASONING TRAINING

In Table 16, we present the performance of Qwen3-8B and its training versions on other important benchmarks, including LiveCodeBench-v5 (Jain et al., 2024), AIME2025 (MAA, 2025), ZebraLogic (Lin et al., 2025), BigCodeBench (Zhuo et al.), MHPP (Dai et al., 2024), FullStack-Bench (Liu et al., 2024b), CruxEval (Gu et al.) and DROP (Dua et al., 2019)). They are well-known benchmarks in code generation, logical reasoning, and mathematics domains. These models were evaluated under non-thinking (cause our model is trained under a non-thinking chat template) conditions with a 4k context length, temperature of 0.6, and top-p of 0.95. The results demonstrate that our chess domain-trained Qwen3-8B model maintains comparable performance across other benchmarks without significant score degradation. Notably, we observe performance improvements on AIME2025 and ZebraLogic, indicating that chess domain-specific training exhibits positive transfer effects to other domains. Furthermore, the cross-domain performance improvements appear to stem primarily from RL, demonstrating the effectiveness of chess RL training. We observed substantial improvements in the model's performance on Zebralogic tasks  $(25.9 \rightarrow 48)$ , with score increases across all puzzle categories: small puzzles, medium puzzles, large puzzles, and XL puzzles  $(70.3/12.1/0.0/0.0 \rightarrow 92.2/55.4/14.5/0.5)$ . Models trained on chess demonstrate significant performance gains on puzzle-solving and related reasoning tasks. We analyzed the model's responses for a case study to investigate the underlying factors contributing to the performance improvements, as shown in Table 17. After chess training, the model became significantly more systematic and logically clear in solving such problems, with a final verification process that markedly improved the model's success rate in addressing these problems.

Table 16: Performance on other important benchmarks of Qwen3-8B and its chess training versions

Model	LiveCodeBench	AIME2025	ZebraLogic	BigCodeBench	MHPP	FullStackBench	CruxEval	DROP
Qwen3-8B	25.19	21.56	25.9	41.32	43.81	45.82	73.25	85.15
Qwen3-8B-Chess-SFT-Stage2	27.48	18.85	30.4	41.4	41.43	44.46	68	82.83
Qwen3-8B-Chess (SFT+RL)	25.19	23.54	48	39.82	44.29	46.59	72.25	83.58

Table 17: ZebraLogic Case Study

#### **Puzzle**

#### **Puzzle to Solve:**

There are 4 houses, numbered 1 to 4 from left to right, as seen from across the street. Each house is occupied by a different person. Each house has a unique attribute for each of the following characteristics:

- Each person has a unique name: 'Alice', 'Eric', 'Arnold', 'Peter'
- Each person has an occupation: 'artist', 'engineer', 'teacher', 'doctor'
- People have unique favorite book genres: 'fantasy', 'science fiction', 'mystery', 'romance'
- People use unique phone models: 'google pixel 6', 'iphone 13', 'oneplus 9', 'samsung galaxy s21'

**Clues:** 1. The person who is an engineer is directly left of the person who uses a Samsung Galaxy S21.

- 2. The person who loves fantasy books is in the second house.
- 3. Alice is not in the second house.
- 4. Eric is the person who is a teacher.
- 5. The person who uses a Samsung Galaxy S21 is the person who loves fantasy books.
- 6. The person who uses an iPhone 13 is the person who loves science fiction books.
- 7. The person who loves science fiction books is somewhere to the left of the person who uses a OnePlus 9.
- 8. The person who uses a OnePlus 9 is Arnold.
- 9. The person who is a doctor is the person who loves mystery books.
- 10. The person who uses an iPhone 13 is the person who is a teacher.

# **Ground-Truth**

```
{
    "House 1": { "Name": "Alice", "Occupation": "engineer", "BookGenre": "romance",
    "PhoneModel": "google pixel 6" },
    "House 2": { "Name": "Peter", "Occupation": "artist", "BookGenre": "fantasy",
    "PhoneModel": "samsung galaxy s21" },
    "House 3": { "Name": "Eric", "Occupation": "teacher", "BookGenre": "science fiction",
    "PhoneModel": "iphone 13" },
    "House 4": { "Name": "Arnold", "Occupation": "doctor", "BookGenre": "mystery",
    "PhoneModel": "oneplus 9" }
}
```

# Case 1: Response from Qwen3-8B. The reasoning process lacks clarity and violates clue 3,9

#### Reasoning:

Starting with Clue 2, we know the person who loves fantasy books is in House 2. Clue 5 says the person who uses a Samsung Galaxy S21 is the one who loves fantasy books, so House 2 uses the Samsung Galaxy S21. Clue 1 says the engineer is directly left of the person using a Samsung Galaxy S21, so the engineer must be in House 1. Clue 10 states the person using an iPhone 13 is the teacher, and Clue 4 says Eric is the teacher, so Eric uses an iPhone 13. Clue 6 says the person using an iPhone 13 loves science fiction books, so Eric loves science fiction books. Clue 7 says the person who loves science fiction books is to the left of the person using a OnePlus 9, so Eric must be in a house to the left of the OnePlus 9 user. Continued on the next page.

# Table 17 (Continued): ZebraLogic Case Study

#### **Case 1: Continue**

1620

1621 1622

1623

1624

1625

1626

1627

1628

1629

1630

1633

1634

1635

1637

1638 1639 1640

1641

1642 1643 Clue 8 says the person using a OnePlus 9 is Arnold, so Arnold is in a house to the right of Eric.

Clue 3 says Alice is not in the second house, and we already know House 2 has the person who loves fantasy books. Clue 9 says the doctor loves mystery books. Clue 10 already assigns Eric as the teacher. Clue 4 confirms Eric is the teacher. Clue 10 also says the person using an iPhone 13 is the teacher, so Eric uses an iPhone 13. Using this information, we can assign the rest of the attributes.

Solution

{

"House 1": { "Name": "Peter", "Occupation": "engineer", "BookGenre": "mystery", "PhoneModel": "google pixel 6" },

"House 2": { "Name": "Alice", "Occupation": "doctor", "BookGenre": "fantasy", "PhoneModel": "samsung galaxy s21" },

"House 3": { "Name": "Eric", "Occupation": "teacher", "BookGenre": "science fiction", "PhoneModel": "iphone 13" },

"House 4": { "Name": "Arnold", "Occupation": "artist", "BookGenre": "romance", "PhoneModel": "oneplus 9" }

Case 2: Response from Qwen3-8B-Chess. The entire reasoning process is divided into 6 systematic steps with clear logical structure, including a final verification process.

```
1644
          Reasoning:
1645
          Let's break down the clues step by step to solve the puzzle.
1646
          Step1: Key Information - Names: Alice, Eric, Arnold, Peter
1647
          - Occupations: artist, engineer, teacher, doctor
1648
          - BookGenres: fantasy, science fiction, mystery, romance
1649
          - PhoneModels: google pixel 6, iphone 13, oneplus 9, samsung galaxy s21
1650
          Step 2: Clues Analysis
          Clue 1: The person who is an engineer is directly left of the person who uses a Samsung Galaxy
1651
          S21. - This means the engineer is in a house with a number one less than the house of the
          Samsung Galaxy S21 user.
          Clue 2: The person who loves fantasy books is in the second house. - House 2's BookGenre is
1654
          fantasy.
1655
          - Eric's Occupation = teacher.
1656
          <For brevity and readability, the following of this step has been omitted.>
1657
          Step3: Step-by-Step Deduction
1658
          From Clue 5 and 2: - House 2's PhoneModel = samsung galaxy s21 (since BookGenre =
1659
          fantasy).
          From Clue 1: - The engineer is directly left of the Samsung Galaxy S21 user (House 2).
1661
          - So the engineer must be in House 1 (since only House 1 is directly left of House 2).
          From Clue 10 and 4: - Eric is a teacher and uses an iPhone 13.
1662
          - So Eric's PhoneModel = iphone 13, and from Clue 6, his BookGenre = science fiction.
1663
          <Omitted.>
1664
          Step 4: Now, let's determine the positions
1665
          House 1: - Occupation = engineer
1666
          - Name = ? (Not Eric, not Arnold, not Alice in House 2)
          - Possible names: Alice or Peter
1668
          House 2: - BookGenre = fantasy
1669

    PhoneModel = samsung galaxy s21

          - Name = ? (Not Alice, not Eric)
1671

    Possible names: Arnold or Peter

1672
          Continued on the next page.
1673
```

```
1674
                                    Table 17 (Continued): ZebraLogic Case Study
1675
1676
         Case 2: Continue.
1677
           House 3:
1678
           - Possible: Eric (if he is not in House 2 or 1)
1679
           - But Eric's PhoneModel = iphone 13, and BookGenre = science fiction.
1680
           House 4:
1681
           - Possible: Arnold (if he is not in House 3)
1682
           - But Arnold's PhoneModel = oneplus 9.
1683
           Let's try to place Eric:
1684
           Eric must be in a house where:
1685
           - PhoneModel = iphone 13
           - BookGenre = science fiction
1687
           - To the left of Arnold (oneplus 9 user)
           Possible positions for Eric:
1688
           - If Eric is in House 2: But House 2's BookGenre = fantasy, not science fiction (contradicts Clue
1689
           - If Eric is in House 3: Then Arnold must be in House 4 (since 3 < 4).
           - If Eric is in House 1: Then Arnold must be in House 2 (but House 2's PhoneModel = samsung
           galaxy s21, not oneplus 9).
1693
           So the only valid position for Eric is **House 3**.
           <Omitted.>
1695
           Step5: Let's check if we missed something(confirm):
           Let's re-examine the occupations:
           - Occupations: artist, engineer, teacher, doctor
           - We have: - House 1: engineer (from Clue 1) - House 3: teacher (Eric) - So remaining: artist
1698
1699
           and doctor must be in Houses 2 and 4.
           House 2: - BookGenre = fantasy (not mystery) - So not the doctor.
1700
           House 4: - BookGenre = romance (not mystery) - So not the doctor.
1701
           This is a problem. Let's check if we made a mistake.
1702
           Let's go back to the initial assumption about the engineer.
1703
           Clue 1: The person who is an engineer is directly left of the person who uses a Samsung Galaxy
1704
           S21.
1705
           We assumed the engineer is in House 1 (since only House 1 is directly left of House 2).
1706
           <Omitted.>
1707
           Step 6: Verification
           - Clue 1: Engineer (House 1) is directly left of Samsung Galaxy S21 user (House 2) \rightarrow \checkmark
1709
           - Clue 2: BookGenre = fantasy in House 2 \rightarrow \checkmark
          - Clue 3: Alice \neq House 2 \rightarrow \checkmark (Alice is in House 1)
1710
           - Clue 4: Eric is a teacher \rightarrow
1711
           - Clue 5: Samsung Galaxy S21 user = fantasy lover \rightarrow \checkmark (House 2)
1712
           - Clue 6: iPhone 13 user = science fiction lover \rightarrow \checkmark (Eric, House 3)
1713
           - Clue 7: Science fiction lover (House 3) is left of OnePlus 9 user (House 4) \rightarrow \checkmark
1714
           - Clue 8: Arnold uses OnePlus 9 \rightarrow \checkmark
1715
           - Clue 9: Doctor = mystery lover \rightarrow \checkmark (Arnold, House 4)
1716
           - Clue 10: iPhone 13 user = teacher \rightarrow \checkmark (Eric, House 3)
1717
           All clues are satisfied.
1718
           Solution: {
1719
                 "House 1": { "Name": "Alice", "Occupation": "engineer", "BookGenre": "romance",
           "PhoneModel": "google pixel 6" },
                 "House 2": { "Name": "Peter", "Occupation": "artist", "BookGenre": "fantasy",
           "PhoneModel": "samsung galaxy s21" },
1722
                 "House 3": { "Name": "Eric", "Occupation": "teacher", "BookGenre": "science fiction",
1723
           "PhoneModel": "iphone 13" },
                 "House 4": { "Name": "Arnold", "Occupation": "doctor", "BookGenre": "mystery",
1725
           "PhoneModel": "oneplus 9" }
1726
```

#### G.3 LEGAL MOVES AS POTENTIAL CONSTRAINTS

In the aforementioned evaluation, we observed superior model performance under the "with legal moves" configuration, which aligns with theoretical expectations. This approach enables the model to select exclusively from predetermined legal moves, thereby virtually eliminating the possibility of generating illegal moves. Barring fundamental deficiencies in instruction-following capabilities, models can consistently select appropriate moves from the provided legal move set. However, this configuration may inadvertently constrain the model's full potential. We discovered that under the "with legal moves" condition, models frequently exhibit diminished analytical engagement, characterized by the tendency to bypass deliberative reasoning processes and directly output moves without substantive thought. Conversely, when the same board state (FEN) is evaluated under "without legal moves" constraints, models often generate superior move selections. Under these unrestricted conditions, models engage in comprehensive tactical analysis, systematically considering viable moves for each piece and conducting thorough evaluations of all possible move qualities. This phenomenon manifests consistently across various non-reasoning models, including Qwen3-235b-a22b (non-thinking) and Deepseek-V3. Illustrative examples of this behavioral pattern are presented in Table 18 and Figure 6.

Table 18: With legal moves, the model's potential may not be fully realized

# Case 1: Qwen3 outputs a move directly (legal moves provided)

```
1747
1748
          Query:
1749
          System prompt is omitted.
1750
          User:
1751
          Current board position in FEN notation: rnbqkbnr/pp2pppp/3p4/8/3pP3/5N2/PPP2PPP/
1752
          RNBOKB1R w KOkg - 04
1753
          Legal moves in UCI notation: f3g5, f3e5, f3h4, f3d4, f3d2, f3g1, h1g1, f1a6, f1b5, f1c4, f1d3,
          fle2, e1e2, e1d2, d1d4, d1d3, d1e2, d1d2, c1h6, c1g5, c1f4, c1e3, c1d2, b1c3, b1a3, b1d2, e4e5,
1754
          h2h3, g2g3, c2c3, b2b3, a2a3, h2h4, g2g4, c2c4, b2b4, a2a4
1755
1756
```

```
Response from Qwen3-235b-a22b(Without any reasoning process, providing its answer directly, exhibiting certain lazy behavior)

e4e5
```

Case 2: Owen3 thinks carefully and outputs a better move (no legal moves provided)

1782 Table 18 (Continued): With legal moves, the model's potential may not be fully realized 1783 1784 1785 Query: 1786 System prompt is omitted. 1787 User: 1788 Current board position in FEN notation: rnbqkbnr/pp2pppp/3p4/8/3pP3/5N2/PPP2PPP/ 1789 RNBQKB1R w KQkq - 04 1790 1791 1792 1793 Response from Qwen3-235b-a22b 1794 Let's analyze the position carefully. 1795 1796 Position in FEN: rnbqkbnr/pp2pppp/3p4/8/3pP3/5N2/PPP2PPP/RNBQKB1R w KQkq - 0 4 1797 **Interpretation:** 1798 - Side to move: White 1799 - Opening: This looks like a **Scotch Game** or a variation of it, with an early d4 by Black. Material is equal 1801 - White has a central pawn on e4, and Black has pawns on d6 and d4, indicating a possible gambit or aggressive central play by Black. - White has not yet developed much, but has the advantage of a tempo lead. 1803 - Black's queen is still on d8, and Black's kingside is intact. - White's knight is already on f3, and the e4 pawn is protected. 1805 - White's light-squared bishop and queen are still on their starting squares. 1806 1807 **Candidate Moves:** 1808 1. **e4-e5** – Pushing the pawn to open the center and gain space. 1809 2. **Qxd4** – Capturing the pawn with the queen. 1810 3. **Nxd4** – Capturing the pawn with the knight. 1811 4. **Bc4** – Developing with a strong bishop eyeing f7. 1812 5. c3 – Supporting a later d4 or preparing for development. 1813 **Evaluation:** - e5 is a strong move, gaining space and limiting Black's options in the center. 1815 - **Qxd4** is playable, but exposes the queen early and may lead to early exchanges. 1816 - Nxd4 is also playable, but after Qxd4, White must move the knight or queen again, giving 1817 Black a tempo advantage. 1818 - **Bc4** is a good developing move, preparing castling and exerting pressure on the kingside. 1819 - c3 is a bit slow here; White should focus on development and center control. 1820 1821 **Best Plan:** Developing with **Bc4** is the most natural and strong move. It develops a piece, pressures the f7 square (which is always a weak point in the opening), and keeps all options open for White in the center. 1824 1825 **Final Move:** 1826 1827 f1c4 1828

1830 1831



(a) Move from Qwen3 when providing legal moves. Qwen3 choose e4e5. It's a normal move.



(b) Move from Qwen3 when not providing legal moves. After careful consideration, Qwen3 choose f1f4. This move keeps options open for White in the center.

Figure 6: Providing legal moves may lock model's potential

#### G.4 RL TRAINING ANALYSIS

**Training Curve** Our RL training mean-reward and response length are shown in Figure 7, which were downloaded from WanDB<sup>5</sup>. Both mean reward and response length exhibit upward trends, indicating that the model's chess strategic reasoning capabilities are progressively improving.

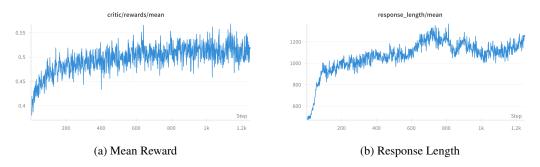


Figure 7: Mean Reward and Response Length Curve of RL Training

Rollout We conducted a comprehensive analysis of the evolution in Qwen3-8B-Chess's board evaluation reasoning throughout reinforcement learning training. By examining the model's cognitive processes for identical chess positions across different training phases, we traced the development of its analytical capabilities. Cases 1 and 2 in Table 19 exemplify this progression. During the early training phase, while Qwen3-8B-Chess demonstrated fundamentally sound analysis, its strategic depth remained limited. Specifically, the model should have prioritized material capture over king threats, as such threats lacked tactical substance in the given position. In contrast, during the late training phase, the model exhibited a marked preference for material acquisition and engaged in systematic verification processes, critically evaluating the validity of its own reasoning. Furthermore, in Case 3 of Table 19, Qwen3-8B-Chess demonstrated enhanced analytical rigor through dual verification mechanisms and extended tactical calculations. The model began simulating move sequences systematically, employing this methodological approach to identify what it determined to be the optimal solution.

<sup>5</sup>https://wandb.ai/

1890 Table 19: Qwen3-8B-Chess thought difference between RL training early stage and late stage 1892 Case 1: Owen3-8B thought process of early training stage 1894 System prompt is omitted. User: Current board position in FEN notation: 4k3/2R5/8/P6K/2p4b/2P2P2/P7/8 w - - 5 69 1898 Legal moves in UCI notation: c7c8, c7h7, c7g7, c7f7, c7e7, c7d7, c7b7, c7a7, c7c6, c7c5, c7c4, 1899 h5h6, h5g6, h5h4, h5g4, a5a6, f3f4, a2a3, a2a4 1900 1901 1902 Response from Qwen3-8B-Chess in early RL training stage 1903 1904 Let's analyze the position carefully. 1905 **FEN:** 4k3/2R5/8/P6K/2p4b/2P2P2/P7/8 w - - 5 69 This is a late-middlegame or early-endgame position. White has a rook on c7, a king on h5, 1907 pawns on a5, c3, f3, and a2. Black has a king on e8, a bishop on h4, and a pawn on c4. 1908 **Kev Observations:** - White Rook on c7 is active and can potentially promote if supported. - Black Bishop on h4 is currently attacking the pawn on g3 (which doesn't exist yet), but more 1910 importantly, it's eyeing the long diagonal. 1911 - Black King on e8 is somewhat passive, but could become active depending on how the game 1912 1913 - White King on h5 is quite advanced and could support a pawn push on the kingside. 1914 - Pawn Structure: 1915 - White has a passed pawn on a5, which could be pushed. 1916 - Black has a pawn on c4, which is isolated and potentially weak. 1917 - White has pawns on c3 and f3, which are mostly defensive. 1918 - White has a potential passed pawn on a2, but it's not very advanced. 1919 **Candidate Moves:** 1920 From the list of legal moves: - Rook moves from c7: 1921 - 'c7c8': Attacks the black king directly. Could lead to a check. - 'c7h7', 'c7g7', 'c7f7', 'c7e7', 'c7d7', 'c7b7', 'c7a7': Various rook moves along the 7th rank. 1923 - 'c7c6', 'c7c5', 'c7c4': Moving the rook down the file. 1924 - King moves from h5: 1925 - 'h5h6', 'h5g6', 'h5h4', 'h5g4': King moves toward the center or supports pawn advances. 1926 - Pawn moves: - 'a5a6': Pushing the a-pawn forward. 1927

- 'f3f4': Pushing the f-pawn.
  - 'a2a3', 'a2a4': Pushing the a-pawn from a2.

#### Evaluation:

The most forcing move is 'c7c8', which puts the rook on the same file as the black king and gives check. This forces the black king to move, and White can then possibly follow up with more active play.

- 1933 Let's look at that line:
  - 1. 'c7c8+' (Rook checks the king) Black must move the king: 'Ke8-d7' or 'Ke8-e7'
  - 2. Then White can continue to press with the rook, possibly supporting the a-pawn or advancing the king.
  - Alternatively, moving the king ('h5g6') could also be useful to support the center, but it's less immediate.
- Pushing the a-pawn ('a5a6') is also a good idea, but it's slower and doesn't create an immediate threat.
  - 40 Best Move:
- The most active and forcing move is:
- 1942

1929

1930

1931

1932

1934

1935

1943 c7c8

• • • •

```
1944
                      Table 19 (Continued): Qwen3-8B-Chess's thought during RL training
1945
1946
        Case 2: Qwen3-8B thinks with confirm process.
1947
1948
          Response from Qwen3-8B-Chess in the late RL training stage
1949
          Let's analyze the position carefully.
1950
          FEN: 4k3/2R5/8/P6K/2p4b/2P2P2/P7/8 w - - 5 69
1951
          This is a late-middlegame or early-endgame position. White has a rook on c7, a king on h5,
1952
          pawns on a5, c3, f3, a2. Black has a king on e8, a bishop on h4, and a pawn on c4.
1953
          Key Observations:
1954
          - White's rook is active on c7, but not yet fully optimized.
1955
          - Black's bishop is active on h4, controlling some diagonals.
          - Black's king is in the center, but not immediately vulnerable.
1957
          - White's king is advanced on h5, which could be useful for supporting pawn advances or
1958
          attacking.
          - Pawn structure: White has a passed pawn on a5, and Black has a passed pawn on c4. Both
          sides have pawns that could become targets.
          Candidate Moves:
1961
          From the list of legal moves: - 'c7c8': Rook moves to c8 (could be useful if aiming to control
1962
          the back rank)
1963
          - 'c7h7': Rook moves to h7 (could support king on h5)
1964
          - 'c7g7': Rook moves to g7 (could control g-file)
1965
          - 'c7f7': Rook moves to f7 (could control f-file)
1966
          - 'c7e7': Rook moves to e7 (could control e-file)
1967
          - 'c7d7': Rook moves to d7 (could control d-file)
1968
          - 'c7b7': Rook moves to b7 (could control b-file)
1969
          - 'c7a7': Rook moves to a7 (could support a-pawn)
          - 'c7c6': Rook moves to c6 (could attack c4 pawn)
1970
          - 'c7c5': Rook moves to c5 (could control c-file)
1971
          - 'c7c4': Rook captures on c4 (could win a pawn)
1972
          - 'h5h6': King moves forward (could support pawn advance)
1973
          - 'h5g6': King moves toward center (could be risky)
1974
          - 'h5h4': King captures bishop (if bishop is not protected)
1975
          - 'h5g4': King moves toward center (could be risky)
          - 'a5a6': Pawn advances (could create a passed pawn)
          - 'f3f4': Pawn advances (could control center)
          - 'a2a3': Pawn advances (could support a5 pawn)
          - 'a2a4': Pawn advances (could create a passed pawn)
          Evaluation:
          Most Promising Moves: - 'c7c4': Capturing the pawn on c4 would give White a material
1981
          advantage (rook for pawn). But we need to check if it's safe.
1982
          - 'h5h4': Capturing the bishop on h4. This would be a major material gain (king captures
          bishop).
1984
           But is the bishop protected? Let's see: - After 'h5h4', Black can respond with 'g5g4' (if
          there's a pawn on g5), but there isn't. So the bishop is not protected. So this is a free bishop.
1986
          So 'h5h4' is a **very strong move**, winning a bishop.
1987
          Let's confirm: - Is the bishop on h4 protected? No.
1988
          - Is the king on h5 safe after capturing? Yes, because the bishop is not protected, and no
1989
          immediate threats.
          So 'h5h4' is a **free bishop**, which is a major material advantage.
          Conclusion:
          The best move is to capture the bishop with the king.
1992
1993
          h5h4
1996
```

Case 3: Qwen3-8B thinks with some double-check and moving sequence analysis process.

1998 Table 19 (Continued): Qwen3-8B-Chess's thought during RL training 1999 2000 Let's analyze the position carefully. 2002 FEN: 2bq2k1/5p2/p4rn1/1pb1pNQ1/2p1P1P1/2P5/PPB2PN1/4R1K1 w - - 1 30 2003 You are playing as White. 2004 **Key Observations:** King Safety: 2006 - White king is on g1, relatively safe for now. Black king is on g8, also fairly safe. Piece Activity: 2007 - White's knight on f5 is active. White's queen on g5 is active and could be dangerous. Black's knight on 2008 g6 is defending h8 and f8 squares. Black's bishop on c5 is active. Black's rook on f6 is active. White's 2009 bishop on c2 is somewhat passive but can become active. White's rook on e1 is on an open file (e-file). 2010 **Pawn Structure:** 2011 - White has a pawn on e4, which is advanced. Black has a pawn on c4, which is also advanced. There are 2012 some isolated and potential weak pawns. Threats: 2013 - White's queen on g5 is eyeing the g-file and could potentially threaten the black king if supported. The knight on f5 could jump to d6 or h6. Black's knight on g6 is protecting the king and may be involved in 2015 defense. Black's bishop on c5 is targeting the e3 square and could be part of a diagonal. 2016 **Candidate Moves:** Let's look at some promising moves for White: 2017 1. f5d6 2018 - Knight moves to d6, attacking the bishop on c8 and possibly threatening to fork. Also attacks the pawn on 2019 e5. Could disrupt Black's structure. 2021 - Knight moves to h6, attacking the rook on f6. Could lead to a tactical opportunity if the rook moves. But Black can just move the rook. 3. g5g6 2023 - Queen moves to g6, attacking the knight on g6. If knight moves, White can possibly capture on f6. But 2024 Black can play ...Rh8 or ...Qe8 to defend. 2025 4. c2a4 2026 - Bishop moves to a4, targeting the pawn on c4. Could put pressure on Black's queenside. 2027 - Knight moves to g7, attacking the pawn on g7. But that's not useful since it's a pawn and not a piece. 2028 6. g5f6 2029 - Queen captures rook on f6! This is a very tempting move. Let's check if it's safe: 2031 - After 'g5f6', Black must recapture with the knight ('g6f6'). - Then White can play 'f5d6', attacking the bishop on c8 and threatening to win material. 2032 - Also, White gains a rook for a queen? No, wait: Queen takes rook, then knight takes queen. 2033 - So White wins a rook and loses a queen? That's a bad trade. - Unless there's a follow-up. 2035 Wait, let's double-check: 2036 - 'g5f6' (Queen takes rook). Black must respond with 'g6f6' (knight takes queen). Now White can play 'f5d6', attacking the bishop on c8 and the pawn on e5. Bishop on c8 is defended by queen (which is gone), 2037 so not immediately vulnerable. But White now has a knight on d6, which is strong. 2038 But is this a good trade? Queen for rook and knight? 2039 Yes! Because Black's knight was defending the king, and now it's gone. Also, White's knight on d6 is 2040 very active. 2041 So this seems like a **good tactical opportunity**. Let's confirm: 'g5f6' (Qxf6) -> 'g6f6' (NxQ) -> 'f5d6' (Ng6-e5 or Ng6-c7?) 2042 But after 'f5d6', Black can play 'c8d7' (Bxd7), and White can play 'd6e8' (Nxe8), winning the queen! 2043 Wait, let's walk through: 2044 1. 'g5f6' (Qxf6) 2. 'g6f6' (NxQ) 3. 'f5d6' (Ng6-e5 or Ng6-c7?) 4. 'c8d7' (Bxd7) 5. 'd6e8' (Nxe8) — White wins the queen! 2046 Yes! That's a winning sequence. So the move 'g5f6' leads to a forced win of the queen. Is there any way Black can avoid this? No, because after 'g5f6', Black must take with the knight. If they don't, White keeps the rook. 2049 So the best move is: 2050 g5f6 ` This is a tactical shot that wins material and gives White a clear advantage.