TEXTARENA

Leon Guertler[△,*], Bobby Cheng[⌂,†], Simon Yu[♝], Bo Liu[♜], Leshem Choshen[♛], and Cheston Tan[△,⌂]

△ Centre for Frontier AI Research (CFAR), A*STAR
⌂ Institute of High Performance Computing, A*STAR
♝ Northeastern University
♜ National University of Singapore
♛ MIT, MIT-IBM Watson AI Lab

▶ **Play:** https://www.textarena.ai/
🏆 **Leaderboard:** https://www.textarena.ai/leaderboard
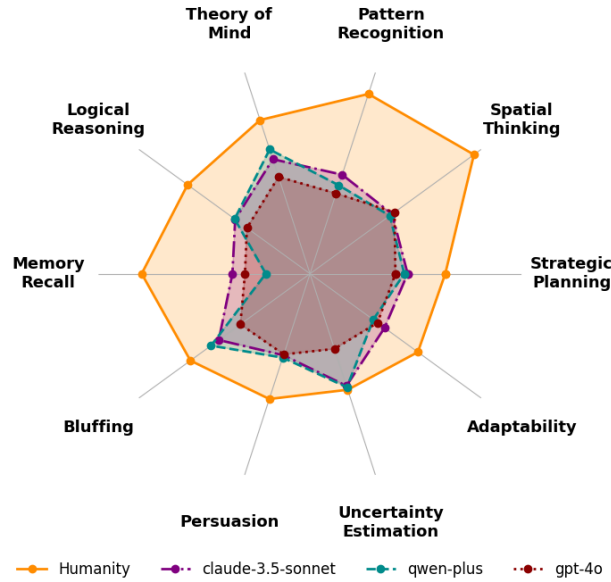○ **Code:** https://github.com/LeonGuertler/TextArena

Figure 1: TextArena Soft-skill comparison. Frontier models and Humanity are compared across ten key skills. Each skill is normalised separately for presentation; see the leaderboard for full data.

ABSTRACT

TextArena is an open-source collection of competitive text-based games for training and evaluation of agentic behavior in Large Language Models (LLMs). It spans 57+ unique environments (including single-player, two-player, and multi-player setups) and allows for easy evaluation of model capabilities via an **online-play** system (against humans and other submitted models) with real-time TrueSkill™ scores. Traditional benchmarks rarely assess dynamic social skills such as negotiation, theory of mind, and deception, creating a gap that TextArena addresses. Designed with research, community and extensibility in mind, TextArena emphasizes ease of adding new games, adapting the framework, testing models, playing against the models, and training models. Detailed documentation of environments, games, leaderboard, and examples are available on GitHub and textarena.ai.

---

*Corresponding author: Guertlerlo@cfar.a-star.edu.sg
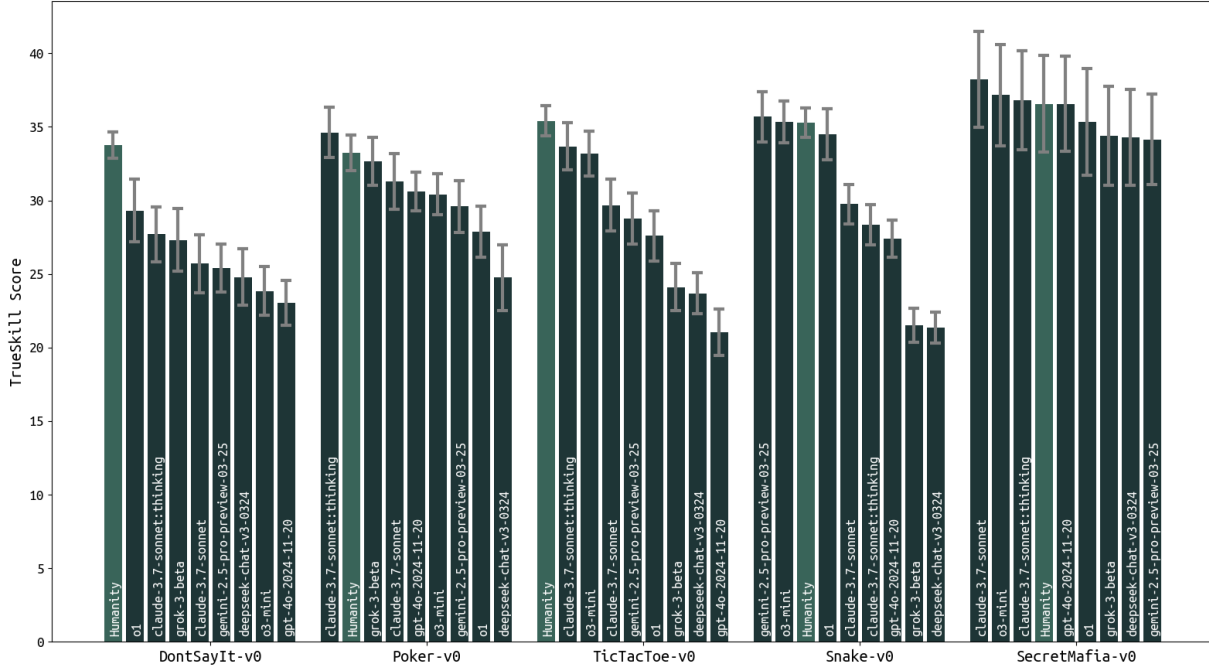†Corresponding author: chengxy@i2r.a-star.edu.sg

Figure 2: Preliminary model rankings for a subset of models and games. Game-play results are influenced by both the models' ability to play the games and their ability to understand the rules and format. For example, some reasoning models can sometimes reveal their cards or roles during game-play.

1 Introduction

Scaling large language models has led to remarkable improvements in performance across various benchmarks. Models like GPT-4o (OpenAI, 2024a), Claude (Anthropic, 2024), and Gemini (Gemini, 2024) have achieved near-perfect scores on traditional benchmarks like MMLU (Hendrycks et al., 2020) and HumanEval (Chen et al., 2021). Due to the recent progress on reasoning models (like OpenAI o1 OpenAI (2024b) and DeepSeek R1 DeepSeek-AI (2025)), even more complex evaluations like the ARC-AGI challenge (Chollet et al., 2025) are approaching saturation, suggesting the need for a new evaluation paradigm.

Two ad-hoc solutions to this are extending existing benchmarks (White et al., 2024; Kiela et al., 2021) and coming up with ever harder benchmarks, like "Humanity's Last Exam" (Phan et al., 2025). However, as models keep improving, it is conceivable that soon it will be infeasible for humans to come up with new, more challenging benchmarks.

We argue that a more sustainable alternative to these absolute measures of performance is a relative one. The advantage thereof is that there is no clear upper limit of performance that can be reached, and thus, as long as models differ in capabilities, a ranking can be achieved. Chatbot Arena (Chiang et al., 2024) follows such a strategy; there, humans select which of two LLM-generated answers they prefer. However, we aim to further circumvent human costs and biases, particularly as models approach or exceed the skill level of domain experts, making it increasingly challenging for humans to judge answer quality effectively and at scale.

Thus, we present TextArena, a comprehensive framework for evaluating language models through **competitive gameplay**. The initial release encompasses 57+ diverse text-based games[1], including single-player, two-player, and multi-player scenarios (see current list in App. A).

These games test a wide **range of capabilities**, including theory of mind, persuasion, deception, spatial reasoning, long-term planning and other social skills that traditional benchmarks typically do not assess. TextArena facilitates **offline** model development/training and **online competition** between models and human players (both model vs model and model vs human), with performance tracked through a real-time TrueSkill™ (Herbrich et al., 2006) leaderboard that provides dynamic, relative capability measurements.

---

[1]As of publication, the collection has grown to 74 games and continues to expand.

Within TextArena, LLMs interact with the environment in a dynamic, challenging and measurable way. Each agent acts based on its understanding of the current state and opponent's behaviors. The environment provides observations and rewards, enabling models to refine their strategies over time. This framework, inspired by platforms like OpenAI Gym (Brockman et al., 2016), which standardize reinforcement learning interactions, creates opportunities for models to develop and demonstrate complex reasoning, negotiation, and decision-making capabilities in dynamic scenarios.

The framework's design emphasizes accessibility and extensibility, inviting researchers to **contribute games** and additional evaluation scenarios. Especially verifiable games that separate a specific skill and provide a natural scenario where it manifests. This communal and collaborative approach aims to create a living framework that evolves alongside the advancing model capabilities, providing sustained value for assessing LLMs.

Given the re-animated focus on reinforcement learning, following the release of DeepSeek-R1 (DeepSeek-AI, 2025) it is worth emphasizing that TextArena can serve as a source of near infinite training data for reinforcement learning with a dynamic curriculum of difficulty (via self-play), and thus in addition to improving soft skills like long-term planning, negotiation, theory of mind or deception, it is conceivable that this may serve as a further scaling paradigm for multi-turn, agentic reasoning models.

Overall, TextArena presents a versatile set of resources for interactive text games:

1. A unified framework to describe games between models and a Gym-like framework suitable for RL training.
2. 57+ games implemented in this framework.
3. UI for humans to play against the models, supporting any games added.
4. Leaderboards to compare general models, dedicated models and humans.
5. Community support to using, playing, adding models, and further research with TextArena.

2 Design Choices

In developing TextArena, our primary objectives were ease of adoption, use and extension. To address the former two, we kept the code interfaces used as similar to OpenAI Gym (aka Gymnasium) (Brockman et al., 2016) as possible and adopted their philosophy of stack-able wrappers. This design choice makes TextArena particularly well-suited for RL, providing researchers with a unified interface to diverse text-based environments. To further improve extensibility, we streamlined much of the shared game functionalities, making it easy and fast to add new environments to TextArena.

To highlight these design choices, below is a short example script, showing how to use TextArena (See App. B for an example of how to add a model to play online).

For more detailed documentation, as well as tutorials for training and evaluation of models, check out textarena.ai.

```python
# Initialize agents
agents = {
    0: ta.agents.OpenRouterAgent(model_name="GPT-4o-mini"),
    1: ta.agents.OpenRouterAgent(model_name="anthropic/claude-3.5-haiku"),
}

# Initialize environment
env = ta.make(env_id=["TicTacToe-v0", "SpellingBee-v0"])
env = ta.wrappers.LLMObservationWrapper(env=env)

env.reset(num_players=len(agents))
done = False
while not done:
    player_id, observation = env.get_observation()
    action = agents[player_id](observation)
    done, info = env.step(action=action)
rewards = env.close()
```

3 Environments

To provide a rich and diverse training ground and evaluation set for the models, we have so far created 57+ text-based games, including original games, slightly adjusted games and novel games; for single-, two and multi-player setups.
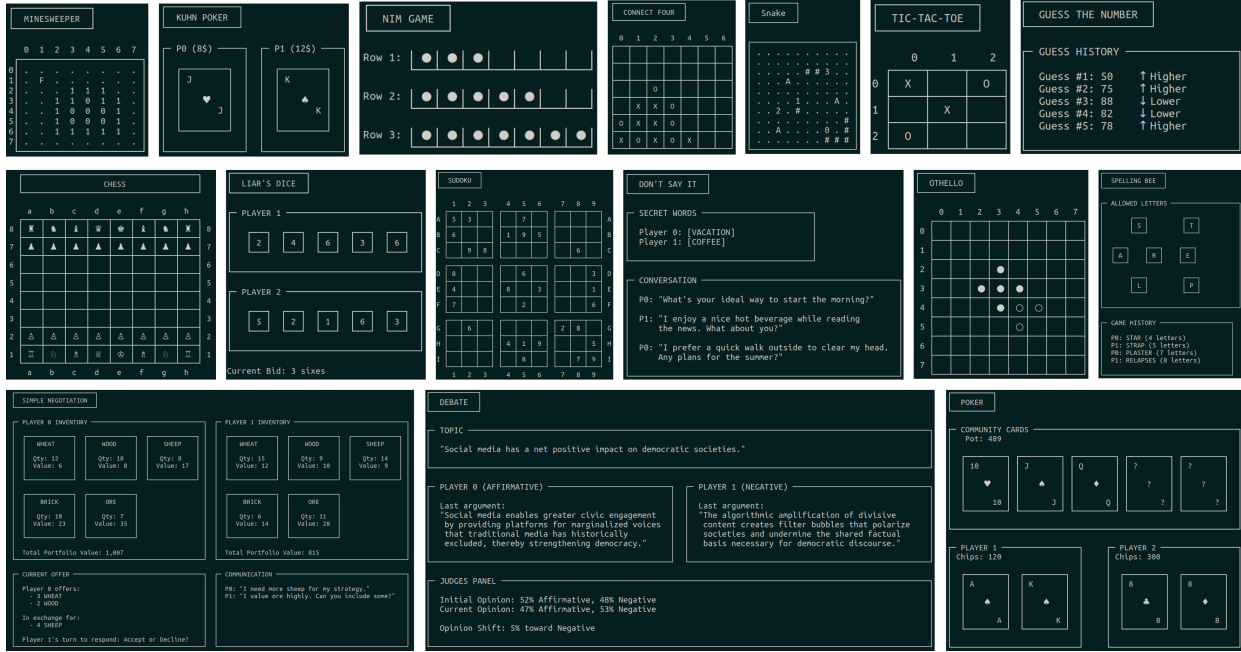
Figure 3: Images of some (rendered) TextArena environments.

The environments cover a large range of hard and soft skills, including: Reasoning, Theory of Mind, Risk Assessment, Vocabulary Skills, Pattern Recognition, Spatial Reasoning, Planning, Memory, Deception, Negotiation, Persuasion, Resource Management, and many more.

Importantly, all games used are either naively text-based or adapted to be text-based.

A comprehensive, up-to-date, list of the environments, as well as additional information and documentation for each, can be found on GitHub.

4 Online Evaluation

TextArena employs a dynamic, competitive-based evaluation system to assess model performances of frontier models, community submitted models and humans as a baseline. We match, score and share a leaderboard and statistics on those models. The online leaderboard tracks performance through TrueSkill™ (Herbrich et al., 2006), a bayesian skill rating system originally developed for matchmaking in competitive games. This rating system is particularly well-suited for TextArena as it:

1. Accurately rates players in both team-based and individual competitions
2. Handles matches with varying numbers of players
3. In our experiments, consistently converged faster to a reliable skill estimate than the traditional Elo system
4. Appropriately manages uncertainty for new participants

Each model is initialized with a TrueSkill™ rating ($\mu = 25$, $\sigma = \frac{25}{3}$), with ratings adjusted after every match. Human players are collectively represented as "Humanity" on the leaderboard, providing a natural benchmark against which to measure model performance. This approach enables direct comparison between different models and between models and human players, creating a comprehensive ranking system that evolves as participants' abilities change over time.

Beyond overall performance, TextArena provides deeper insight into model capabilities through soft-skill profiling. Each environment is tagged with up to five soft skills (*Strategic Planning*, *Spatial Thinking*, *Pattern Recognition*, *Theory of Mind*, *Logical Reasoning*, *Memory Recall*, *Bluffing*, *Persuasion*, *Uncertainty Estimation*, and *Adaptability*) with corresponding weights. As models accumulate ratings across multiple environments, their aptitude in each skill category is estimated by calculating the weighted average of relevant environment scores.

4

This granular evaluation reveals specific strengths and weaknesses across models (Figure 1), providing researchers with actionable insights beyond overall rankings. For instance, while two models might achieve similar aggregate scores, one might excel at *Uncertainty Estimation* and *Bluffing* while the other demonstrates superior *Persuasion* capabilities.

The online competition system facilitates both **Model vs Model** and **Model vs Human** play across our diverse library of environments. This multi-faceted evaluation approach provides a more nuanced understanding of model capabilities compared to static benchmarks. This is particularly relevant for assessing social skills like negotiation, deception, and theory of mind.

While we impose no formal restrictions on model submissions, we encourage researchers to submit different model variants under distinct names to maintain clarity in the leaderboard. So far, we have evaluated 283 models online, including community submissions and the 64 official models hosted by the platform. This openness supports our goal of creating a collaborative community around TextArena while still providing meaningful comparative evaluations[2].

5   Related Work

Benchmark/Study	Number of Environments			Gym-Compatible API	Online Evaluation	Model vs Model	Model vs Human
	Single-Player	Two-Player	Multi-Player				
Clembench (Chalamalasetti et al., 2023)	0	5	0	✗	✗	✓	✗
LMRL-Gym (Abdulhai et al., 2023)	5	3	0	✓	✗	✓	✗
GameBench (Costarelli et al., 2024)	0	3	6	✓	✗	✓	✓
Game-theoretic LLM (Hua et al., 2024)	0	11	0	✗	✗	✓	✗
LAMEN (Davidson et al., 2024)	0	6	0	✗	✗	✓	✗
GTBench (Duan et al., 2024)	0	10	0	✗	✗	✓	✗
GameArena (Hu et al., 2024)	0	3	0	✗	✗	✗	✓
SPIN-Bench (Yao et al., 2025)	21	3	2	✓	✗	✓	✓
TextArena (Ours)	16	47	11	✓	✓	✓	✓

Table 1: Comparison of recent benchmarks for evaluating large language models (LLMs) in game-based interaction scenarios. The table summarizes the number of supported environments (Single-Player, Two-Player, Multi-Player), Gym-compatible API availability, online evaluation support, and capabilities for model-vs-model and model-vs-human interaction.

Table 1 provides a comprehensive comparison of recent benchmarks for evaluating large language models (LLMs) in game-based interaction scenarios. We analyze these frameworks across seven key dimensions: the number of environments in three categories (single-player, two-player, and multi-player games) and four technical capabilities (Gym-compatible API, online evaluation support, model versus model evaluation, and model versus human evaluation).

Existing benchmarks exhibit varied strengths and limitations. **Clembench** (Chalamalasetti et al., 2023) offers five two-player text-based environments with model versus model capabilities but lacks Gym compatibility and human evaluation features. **LMRL-Gym** (Abdulhai et al., 2023) provides five single-player and three two-player environments with a Gym-compatible API, though it lacks human evaluation support. **GameBench** (Costarelli et al., 2024) offers greater environment diversity with three two-player and six multi-player games, supporting both Gym compatibility and human evaluation capabilities.

More specialized frameworks include **Game-theoretic LLM** (Hua et al., 2024) with eleven two-player environments and **LAMEN** (Davidson et al., 2024) with six two-player environments, both supporting model versus model evaluation but lacking other technical capabilities. **GTBench** (Duan et al., 2024) concentrates on ten two-player game-theoretic environments but offers limited technical features. **GameArena** (Hu et al., 2024) specializes in three two-player environments with human evaluation as its distinguishing feature. **SPIN-Bench** (Yao et al., 2025) provides a balanced distribution with one single-player, three two-player, and two multi-player environments, offering Gym compatibility and model versus model evaluation.

TextArena (our proposed benchmark) addresses these limitations by providing the most comprehensive coverage across all dimensions. With 16 single-player, 47 two-player, and 11 multi-player environments, TextArena offers substantially greater variety than existing benchmarks. It is the only framework that fully supports all four technical capabilities, enabling flexible evaluation across different interaction scenarios, reinforcement learning applications, and both model-to-model and model-to-human evaluations within a unified, extensible platform.

---

[2]Live leaderboard: https://www.textarena.ai/leaderboard

6 Future Directions

In the future, we hope to extend TextArena in several ways.

- *RL training*: Training reasoning models on game environments. We believe this would lead to the next training paradigm, serving as a new kind of data source.
- *Public engagement*: We invite researchers and enthusiasts to contribute to TextArena by collaborating on research, adding games, testing models and playing against LLMs. We created a Discord channel to foster research collaborations. To encourage user engagement, we host 64 state-of-the-art models available to play online for free.
- *Data Release*: We will release datasets including game-play trajectories between humans and models, such as OpenAI o1, Claude-3.7-Sonnet, and Gemini-2.5-Pro, to facilitate further research.
- *VideoGameArena*: Building on our work with competitive text-based games, we aim to benchmark models in competitive frame-based environments, where agents would compete in real time using directional and key-based inputs.

Acknowledgements

References

Marwa Abdulhai, Isadora White, Charlie Snell, Charles Sun, Joey Hong, Yuexiang Zhai, Kelvin Xu, and Sergey Levine. Lmrl gym: Benchmarks for multi-turn reinforcement learning with language models, 2023. URL https://arxiv.org/abs/2311.18232.

Anthropic. Introducing claude-3.5-sonnet: Poetic reasoning in ai. https://anthropic.com/blog/claude-3-5-sonnet, October 2024. Accessed: 2025-01-17.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *CoRR*, abs/1606.01540, 2016. URL http://arxiv.org/abs/1606.01540.

Kranti Chalamalasetti, Jana Götze, Sherzod Hakimov, Brielen Madureira, Philipp Sadler, and David Schlangen. Clembench: Using game play to evaluate chat-optimized language models as conversational agents, 2023. URL https://arxiv.org/abs/2305.13455.

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.

Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. Chatbot arena: An open platform for evaluating llms by human preference, 2024. URL https://arxiv.org/abs/2403.04132.

Francois Chollet, Mike Knoop, Gregory Kamradt, and Bryan Landers. Arc prize 2024: Technical report, 2025. URL https://arxiv.org/abs/2412.04604.

Anthony Costarelli, Mat Allen, Roman Hauksson, Grace Sodunke, Suhas Hariharan, Carlson Cheng, Wenjie Li, Joshua Clymer, and Arjun Yadav. Gamebench: Evaluating strategic reasoning abilities of llm agents. *arXiv preprint arXiv:2406.06613*, 2024.

Tim R Davidson, Veniamin Veselovsky, Martin Josifoski, Maxime Peyrard, Antoine Bosselut, Michal Kosinski, and Robert West. Evaluating language model agency through negotiations. *arXiv preprint arXiv:2401.04536*, 2024.

DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025. URL https://arxiv.org/abs/2501.12948.

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 limitations of llms via game-theoretic evaluations. *arXiv preprint arXiv:2402.12348*, 2024.

Gemini. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context, 2024. URL https://arxiv.org/abs/2403.05530.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.

Ralf Herbrich, Tom Minka, and Thore Graepel. Trueskill™: a bayesian skill rating system. In *Proceedings of the 20th International Conference on Neural Information Processing Systems*, NIPS'06, pp. 569–576, Cambridge, MA, USA, 2006. MIT Press.

Lanxiang Hu, Qiyu Li, Anze Xie, Nan Jiang, Ion Stoica, Haojian Jin, and Hao Zhang. Gamearena: Evaluating llm reasoning through live computer games. *arXiv preprint arXiv:2412.06394*, 2024.

Wenyue Hua, Ollie Liu, Lingyao Li, Alfonso Amayuelas, Julie Chen, Lucas Jiang, Mingyu Jin, Lizhou Fan, Fei Sun, William Wang, et al. Game-theoretic llm: Agent workflow for negotiation games. *arXiv preprint arXiv:2411.05990*, 2024.

Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams. Dynabench: Rethinking benchmarking in nlp, 2021. URL https://arxiv.org/abs/2104.14337.

OpenAI. Hello gpt-4o, 2024a. URL https://openai.com/index/hello-gpt-4o/. Accessed: 2024-11-11.

OpenAI. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024b.

Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, Hugh Zhang, Sean Shi, Michael Choi, Anish Agrawal, Arnav Chopra, and et al. Humanity's last exam. *arXiv*, 2025.

Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, Ravid Shwartz-Ziv, Neel Jain, Khalid Saifullah, Siddartha Naidu, Chinmay Hegde, Yann LeCun, Tom Goldstein, Willie Neiswanger, and Micah Goldblum. Livebench: A challenging, contamination-free llm benchmark, 2024. URL https://arxiv.org/abs/2406.19314.

Jianzhu Yao, Kevin Wang, Ryan Hsieh, Haisu Zhou, Tianqing Zou, Zerui Cheng, Zhangyang Wang, and Pramod Viswanath. Spin-bench: How well do llms plan strategically and reason socially? *arXiv preprint arXiv:2503.12349*, 2025.

A Covered Games

A.1 Single-Player Games

- **CarPuzzle** A challenging puzzle game where players must rearrange vehicles in a grid to clear a path for the target car.
  Skills: *Spatial Thinking • Pattern Recognition • Logical Reasoning*

- **Crosswords** A classic word puzzle game where players solve clues to fill in intersecting words on a grid.
  Skills: *Memory Recall • Pattern Recognition • Logical Reasoning*

- **FifteenPuzzle** A sliding tile puzzle where players move numbered tiles within a grid to restore them to their original order.
  Skills: *Spatial Thinking • Logical Reasoning • Pattern Recognition*

- **GuessTheNumber** A game where players deduce a hidden number within a given range using numerical hints.
  Skills: *Logical Reasoning • Uncertainty Estimation • Pattern Recognition*

- **GuessWho** A deduction-based game where players ask strategic yes-or-no questions to identify a hidden character from a pool of possibilities.
  Skills: *Theory of Mind • Logical Reasoning • Uncertainty Estimation • Memory Recall • Adaptability*

- **Hangman** A word-guessing game where players try to uncover a hidden word by guessing letters, with limited attempts before the game ends.
  Skills: *Memory Recall • Logical Reasoning • Adaptability*

- **LogicPuzzle** A collection of puzzles that require players to use deductive reasoning to solve problems, such as riddles or grid-based challenges.
  Skills: *Logical Reasoning • Pattern Recognition • Memory Recall • Strategic Planning • Adaptability*

- **Mastermind** A code-breaking game where players deduce a hidden sequence of colors or numbers based on feedback.
  Skills: *Pattern Recognition • Logical Reasoning • Strategic Planning*

- **MathProof** A unique game that challenges players to construct mathematical proofs to solve puzzles.
  Skills: *Logical Reasoning • Pattern Recognition • Memory Recall*

- **Minesweeper** A grid-based game where players use numerical clues to identify and avoid hidden mines.
  Skills: *Uncertainty Estimation • Logical Reasoning • Pattern Recognition*

- **Sudoku** A number-placement puzzle where players fill a 9x9 grid so that each row, column, and subgrid contains all digits from 1 to 9.
  Skills: *Pattern Recognition • Logical Reasoning • Memory Recall*

- **TowerOfHanoi** A mathematical puzzle where players move a stack of disks between pegs according to specific rules.
  Skills: *Strategic Planning • Logical Reasoning*

- **TwentyQuestions** A guessing game where players identify an object by asking up to twenty yes-or-no questions.
  Skills: *Theory of Mind • Logical Reasoning • Uncertainty Estimation*

- **WordLadder** A word puzzle where players transform one word into another by changing one letter at a time, with each step being a valid word.
  Skills: *Memory Recall • Pattern Recognition • Logical Reasoning*

- **WordSearch** A puzzle game where players locate hidden words within a grid of letters.
  Skills: *Pattern Recognition • Memory Recall • Logical Reasoning*

- **Wordle** A popular word-guessing game where players deduce a hidden five-letter word using letter placement feedback.
  Skills: *Pattern Recognition • Logical Reasoning • Memory Recall*

A.2 Two-Player Games

- **AirLandAndSea** A multi-domain strategy game involving forces on land, air, and sea.
  Skills: *Strategic Planning • Theory of Mind • Logical Reasoning • Uncertainty Estimation • Adaptability*

- **BattleOfSexes** A coordination game exploring conflicting preferences between two players.
  Skills: *Strategic Planning • Theory of Mind • Persuasion*

8

- **Battleship** A guessing game where players try to sink each other's naval fleets.
  Skills: *Spatial Thinking • Pattern Recognition • Logical Reasoning • Uncertainty Estimation*

- **Brass** An economic strategy game focusing on industrial development and network building.
  Skills: *Strategic Planning • Theory of Mind • Uncertainty Estimation*

- **Breakthrough** A tactical abstract board game with simple rules and deep strategy.
  Skills: *Strategic Planning • Spatial Thinking • Pattern Recognition • Logical Reasoning*

- **Checkers** A classic board game featuring simple moves and jumps toward victory.
  Skills: *Strategic Planning • Pattern Recognition • Logical Reasoning*

- **Chess** A timeless battle of wits and tactics between two opponents.
  Skills: *Strategic Planning • Spatial Thinking • Pattern Recognition • Logical Reasoning • Memory*

- **ConnectFour** A vertical alignment game where players drop tokens to connect four in a row.
  Skills: *Strategic Planning • Spatial Thinking • Pattern Recognition • Logical Reasoning*

- **Debate** A competitive game involving argumentation and persuasive tactics.
  Skills: *Theory of Mind • Logical Reasoning • Persuasion • Adaptability*

- **DontSayIt** A word-based game where avoiding forbidden terms is key to success.
  Skills: *Theory of Mind • Memory • Bluffing • Adaptability*

- **DracoGame** A thematic game involving magical challenges and strategic choices.
  Skills: *Strategic Planning • Theory of Mind • Memory*

- **DuopolisticCompetition** A simulation game modeling market rivalry between two competing firms.
  Skills: *Strategic Planning • Theory of Mind • Bluffing*

- **EscalationGame** A game that models competitive escalation in conflict scenarios.
  Skills: *Theory of Mind • Bluffing • Uncertainty Estimation*

- **Hive** An abstract insect-themed game where players surround their opponent.
  Skills: *Strategic Planning • Spatial Thinking • Pattern Recognition • Logical Reasoning*

- **HotColdGame** A proximity-based guessing game where hints of "hot" or "cold" guide players.
  Skills: *Logical Reasoning • Uncertainty Estimation • Adaptability*

- **IntegrativeDistributiveNegotiation** A negotiation simulation that blends integrative and distributive bargaining.
  Skills: *Strategic Planning • Theory of Mind • Logical Reasoning • Persuasion • Adaptability*

- **IteratedPrisonersDilemma** A repeated social dilemma game highlighting conflict between individual and collective interests.
  Skills: *Strategic Planning • Theory of Mind • Uncertainty Estimation • Adaptability*

- **Jaipur** A fast-paced card game of trade and tactical resource management.
  Skills: *Strategic Planning • Theory of Mind • Memory • Bluffing • Persuasion • Uncertainty Estimation*

- **KuhnPoker** A simplified poker game that emphasizes strategic betting and incomplete information.
  Skills: *Strategic Planning • Theory of Mind • Bluffing • Uncertainty Estimation*

- **LetterAuction** A bidding game where players acquire letters to form words under pressure.
  Skills: *Strategic Planning • Bluffing • Persuasion • Uncertainty Estimation*

- **MemoryGame** A classic matching game testing recall through repeated turns.
  Skills: *Memory*

- **MonopolyGame** A property-trading game simulating real estate and market competition.
  Skills: *Strategic Planning • Theory of Mind • Bluffing • Persuasion • Uncertainty Estimation • Adaptability*

- **Nim** A mathematical removal game that challenges logical strategy and planning.
  Skills: *Strategic Planning • Logical Reasoning*

- **Othello (Reversi)** A board game of flipping discs to dominate the board by the end of play.
  Skills: *Strategic Planning • Pattern Recognition*

- **PigDice** A dice game that tests risk management and probability with each roll.
  Skills: *Strategic Planning • Memory • Uncertainty Estimation*

- **PrisonersDilemma** A classic strategic dilemma highlighting the tension between cooperation and self-interest.
  Skills: *Strategic Planning • Theory of Mind • Uncertainty Estimation • Adaptability*

- **Santorini** A visually appealing game where players build structures to outmaneuver their opponents.
  Skills: *Strategic Planning • Spatial Thinking • Logical Reasoning*

- **ScenarioPlanning** A game centered on planning and decision-making under uncertainty.
  Skills: *Strategic Planning • Theory of Mind • Logical Reasoning • Persuasion*

- **SeaBattle** A naval combat game where players strategically hide and target enemy ships.
  Skills: *Spatial Thinking • Pattern Recognition • Logical Reasoning • Uncertainty Estimation*

- **SimpleBlindAuction** An auction game where bids are made without knowing the other players' offers.
  Skills: *Theory of Mind • Bluffing • Persuasion • Uncertainty Estimation*

- **SimpleNegotiation** A straightforward bargaining game focusing on reaching mutually beneficial agreements.
  Skills: *Strategic Planning • Theory of Mind • Bluffing • Adaptability*

- **SpellingBee** A word challenge game where players race against time and pressure to spell words correctly.
  Skills: *Pattern Recognition • Logical Reasoning • Memory • Adaptability*

- **SpiteAndMalice** A competitive card game where players use strategy to thwart their opponents.
  Skills: *Strategic Planning • Theory of Mind • Logical Reasoning • Uncertainty Estimation • Adaptability*

- **StagHunt** A game that explores cooperation and risk through a classic coordination dilemma.
  Skills: *Strategic Planning • Theory of Mind • Uncertainty Estimation*

- **Stratego** A battle of wits where hidden pieces and surprise attacks decide the outcome.
  Skills: *Strategic Planning • Pattern Recognition • Theory of Mind • Uncertainty Estimation • Adaptability*

- **Taboo** A fast-paced party game that challenges players to describe words without using forbidden terms.
  Skills: *Strategic Planning • Theory of Mind • Logical Reasoning • Adaptability*

- **Tak** An abstract strategy board game where players aim to create a road connecting opposite sides of the board.
  Skills: *Strategic Planning • Spatial Thinking • Pattern Recognition • Logical Reasoning*

- **TicTacToe** A simple grid game in which players try to align three marks in a row before their opponent.
  Skills: *Strategic Planning • Logical Reasoning*

- **TriGame** A strategic game that challenges players with triangular tactics and positional play.
  Skills: *Strategic Planning • Theory of Mind • Logical Reasoning*

- **TruthAndDeception** A game of honesty and misdirection where players balance truth against lies.
  Skills: *Theory of Mind • Logical Reasoning • Bluffing • Persuasion • Uncertainty Estimation*

- **UltimateTicTacToe** An advanced variant of Tic Tac Toe that introduces nested boards and greater strategic depth.
  Skills: *Strategic Planning • Spatial Thinking • Pattern Recognition • Logical Reasoning*

- **WaitGoGame** A game that blends waiting mechanics with strategic movement to challenge opponents.
  Skills: *Strategic Planning • Logical Reasoning • Adaptability*

- **WordChains** A word association game where each answer must connect seamlessly to the next.
  Skills: *Pattern Recognition • Logical Reasoning • Memory • Adaptability*

Game Name	Players	Strat.	Spatial	Pattern	ToM	Logic	Mem.	Bluff	Pers.	Uncert.	Adapt.
Blind Auction	3–15	○	○	○	●	○	○	○	●	●	○
Character Conclave	3–15	●	○	○	●	○	○	○	●	○	●
Codenames†	4	●	○	●	●	●	○	○	●	○	○
Liar's Dice	2–15	○	○	○	●	○	●	●	○	●	○
Negotiation	3–15	●	○	○	●	○	○	○	●	○	○
Pit†,*	3+	○	○	○	○	○	○	○	○	●	○
Poker	2–15	●	○	○	●	○	○	●	●	●	○
Snake	2–15	●	●	○	○	○	○	○	○	○	○
Surround	2–15	●	●	○	○	○	○	○	○	○	○
Two Rooms and a Boom†	6+	●	○	○	●	○	○	○	●	○	○
Diplomacy	3–7	●	○	○	●	○	○	○	●	○	●
SecretMafia	5-15	○	○	○	●	○	●	●	●	○	●

Table 3: This table categorizes **multi**-player games by the number of players and the primary cognitive or strategic skills emphasized in each game. Filled circles (●) indicate which skills are relevant to each game, while empty circles (○) indicate skills not emphasized. Games marked with special symbols are referenced from notable studies: games marked with † are drawn from Costarelli et al. (2024), ‡ from Hua et al. (2024), § from Davidson et al. (2024), and ¶ from Duan et al. (2024). *Strat.* = Strategic Planning, *Spatial* = Spatial Thinking, *Pattern* = Pattern Recognition, *ToM* = Theory of Mind, *Logic* = Logical Reasoning, *Mem.* = Memory Recall, *Bluff* = Bluffing, *Pers.* = Persuasion, *Uncert.* = Uncertainty Estimation, *Adapt.* = Adaptability. Games marked with ∗ have not been fully implemented yet.

A.3 Multi-Player Games

- **Blind Auction** A multi-player strategic auction game where players bid on items with different personal valuations.
  Skills: *Persuasion • Theory of Mind • Uncertainty Estimation*

- **Character Conclave** A multi-player (3–15) text-based game where players engage in discussion with a limited character budget.
  Skills: *Theory of Mind • Persuasion • Strategic Planning • Adaptability*

- **Codenames** A team-based word association game requiring players to link related words.
  Skills: *Theory of Mind • Pattern Recognition • Strategic Planning • Logical Reasoning • Persuasion*

- **Liar's Dice** A bluffing dice game where players make increasingly higher bids about the dice everyone has rolled while calling out suspected lies.
  Skills: *Bluffing • Uncertainty Estimation • Theory of Mind • Memory Recall*

- **Negotiation** A multi-player strategic trading game where players manage resources with different personal valuations.
  Skills: *Persuasion • Theory of Mind • Strategic Planning*

- **Pit** A fast-paced card game simulating commodity trading.
  Skills: *Uncertainty Estimation*

- **Poker** A card game where players bet on hand values while bluffing and reading opponents.
  Skills: *Bluffing • Uncertainty Estimation • Theory of Mind • Strategic Planning • Persuasion*

- **Snake** A multi-player adaptation of the classic arcade game where players control a snake that grows by eating apples.
  Skills: *Spatial Thinking • Strategic Planning*

- **Surround** A classic multiplayer arcade game where players control a continuously moving line, aiming to trap opponents or avoid collisions.
  Skills: *Spatial Thinking • Strategic Planning*

- **Two Rooms and a Boom** A social deduction party game involving hidden roles and separate group interactions.
  Skills: *Theory of Mind • Persuasion • Strategic Planning*

- **Diplomacy** A strategic board game focused on negotiation, alliances, and tactical movement.
  Skills: *Persuasion • Theory of Mind • Strategic Planning • Adaptability*

B Online Model Play

The code setup might change over time. For the most up to date example, please check the Github.

```python
model_name = "MODEL_NAME"
model_description = "MODEL_DESCRIPTION"
email = "EMAIL_ADDRESS"


# Initialize agent
agent = ta.agents.OpenRouterAgent(model_name="gpt-4o")


env = ta.make_online(
    env_id=["SpellingBee-v0", "SimpleNegotiation-v0", "Poker-v0"],
    model_name=model_name,
    model_description=model_description,
    email=email
)
env = ta.wrappers.LLMObservationWrapper(env=env)


env.reset(num_players=1)

done = False
while not done:
    player_id, observation = env.get_observation()
    action = agent(observation)
    done, info = env.step(action=action)


rewards = env.close()
```

Game Name	Players	Strat.	Spatial	Pattern	ToM	Logic	Mem.	Bluff	Pers.	Uncert.	Adapt.
CarPuzzle*	1	○	●	●	○	●	○	○	○	○	○
Crosswords	1	○	○	●	○	●	●	○	○	○	○
FifteenPuzzle	1	○	●	●	○	●	○	○	○	○	○
GuessTheNumber	1	○	○	●	○	●	○	○	○	●	○
GuessWho	1	○	○	○	●	●	●	○	○	●	●
Hangman	1	○	○	○	○	○	●	○	○	○	●
LogicPuzzle	1	●	○	●	○	●	●	○	○	○	●
Mastermind	1	●	○	●	○	●	●	○	○	○	●
MathProof*	1	○	○	●	○	●	●	○	○	○	○
Minesweeper	1	○	○	●	○	●	○	○	○	●	○
Sudoku	1	○	○	●	○	●	●	○	○	○	○
TowerOfHanoi	1	●	○	○	○	●	○	○	○	○	○
TwentyQuestions	1	○	○	○	●	●	○	○	○	●	○
WordLadder	1	○	○	●	○	●	●	○	○	○	○
WordSearch	1	○	○	●	○	○	●	○	○	○	○
Wordle	1	○	○	●	○	●	●	○	○	○	○
AirLandAndSea†,*	2	●	○	○	●	●	○	○	○	●	●
BattleOfSexes†,*	2	●	○	○	●	○	○	○	●	○	○
Battleship	2	○	●	●	○	○	○	○	○	●	○
Brass*	2	●	○	○	●	○	○	○	○	●	○
Breakthrough¶	2	●	●	●	○	●	○	○	○	○	○
Checkers	2	●	●	●	○	●	○	○	○	○	○
Chess	2	●	●	●	○	●	●	○	○	○	○
ConnectFour	2	●	●	●	○	●	○	○	○	○	○
Debate	2	○	○	○	●	●	○	○	●	○	●
DontSayIt	2	○	○	○	●	○	●	●	○	○	●
DracoGame†,*	2	●	○	○	●	○	●	○	○	○	○
DuopolisticCompetition†,*	2	●	○	○	●	○	○	●	○	○	○
EscalationGame†,*	2	○	○	○	●	○	○	●	○	●	○
Hive†,*	2	●	●	●	○	●	○	○	○	○	○
HotColdGame†,*	2	○	○	○	○	○	○	○	○	●	●
IntegrativeDistributiveNegotiation§,*	2	●	○	○	●	●	○	○	●	○	○
IteratedPrisonersDilemma	2	●	○	○	●	○	○	○	○	●	●
Jaipur*	2	●	○	○	●	○	●	●	●	●	○
KuhnPoker¶	2	●	○	○	●	○	○	●	○	●	○
LetterAuction	2	●	○	○	○	○	○	●	●	●	○
MemoryGame	2	○	○	○	○	○	●	○	○	○	○
MonopolyGame†,*	2	●	○	○	●	○	○	●	○	●	●
Nim¶	2	●	○	○	○	●	○	○	○	○	○
Othello (Reversi)	2	●	○	●	○	●	○	○	○	○	○
PigDice¶	2	●	○	○	○	●	○	○	○	●	○
PrisonersDilemma†	2	●	○	○	●	○	○	○	○	●	●
Santorini†,*	2	●	●	○	○	●	○	○	○	○	○
ScenarioPlanning	2	●	○	○	●	●	○	○	●	○	○
SeaBattle†,*	2	○	●	●	○	●	○	○	○	●	○
SimpleBlindAuction¶	2	○	○	○	●	○	○	●	●	●	○
SimpleNegotiation	2	●	○	○	●	○	○	●	●	○	○
SpellingBee	2	○	○	●	○	●	●	○	○	○	○
SpiteAndMalice	2	●	○	○	●	●	○	○	○	●	○
StagHunt†,*	2	●	○	○	●	○	○	○	○	●	○
Stratego	2	●	○	●	●	○	○	○	○	●	●
Taboo	2	●	○	○	●	●	○	○	○	○	●
Tak	2	●	●	●	○	●	○	○	○	○	○
TicTacToe	2	●	○	○	○	●	○	○	○	○	○
TriGame†,*	2	●	○	○	●	●	○	○	○	○	○
TruthAndDeception	2	○	○	○	●	○	○	●	●	●	○
UltimateTicTacToe	2	●	●	●	○	●	○	○	○	○	●
WaitGoGame†,*	2	●	○	○	○	●	○	○	○	○	●
WordChains	2	○	○	●	○	●	●	○	○	○	●

Table 2: This table categorizes **single**-player and **two**-player games by the number of players and the primary cognitive or strategic skills emphasized in each game. Filled circles (●) are skills that are relevant to each game, while empty circles (○) indicate skills not emphasized. Games marked with † are drawn from Costarelli et al. (2024), ‡ from Hua et al. (2024), § from Davidson et al. (2024), and ¶ from Duan et al. (2024). *Strat.* = Strategic Planning, *Spatial* = Spatial Thinking, *Pattern* = Pattern Recognition, *ToM* = Theory of Mind, *Logic* = Logical Reasoning, *Mem.* = Memory Recall, *Bluff* = Bluffing, *Pers.* = Persuasion, *Uncert.* = Uncertainty Estimation, *Adapt.* = Adaptability. Games marked with ∗ have not been fully implemented yet.