

PokerBench: Training Large Language Models to Become Professional Poker Players

Richard Zhuang¹, Akshat Gupta^{1*}, Richard Yang¹, Aniket Rahane¹, Zhengyu Li²,
Gopala Anumanchipalli¹

¹University of California Berkeley, 110 Sproul Hall, Berkeley, CA 94720 USA

²Georgia Institute of Technology, 225 North Avenue NW, Atlanta, GA 30332 USA

{richardzhuang0412,akshat.gupta,ryang3881,aniketarahane,gopala}@berkeley.edu, {brian.li}@gatech.edu

Abstract

We introduce POKERBENCH - a benchmark for evaluating the poker-playing abilities of large language models (LLMs). As LLMs excel in traditional NLP tasks, their application to complex, strategic games like poker poses a new challenge. Poker, an incomplete information game, demands a multitude of skills such as mathematics, reasoning, planning, strategy, and a deep understanding of game theory and human psychology. This makes Poker the ideal next frontier for large language models. POKERBENCH consists of a comprehensive compilation of 11,000 most important scenarios, split between pre-flop and post-flop play, developed in collaboration with trained poker players. We evaluate prominent models including GPT-4, ChatGPT 3.5, and various Llama and Gemma series models, finding that all state-of-the-art LLMs underperform in playing optimal poker. However, after fine-tuning, these models show marked improvements. We validate POKERBENCH by having models with different scores compete with each other, demonstrating that higher scores on POKERBENCH lead to higher win rates in actual poker games. Through gameplay between our fine-tuned model and GPT-4, we also identify limitations of simple supervised fine-tuning for learning optimal playing strategy, suggesting the need for more advanced methodologies for effectively training language models to excel in games. POKERBENCH thus presents a unique benchmark for a quick and reliable evaluation of the poker-playing ability of LLMs as well as a comprehensive benchmark to study the progress of LLMs in complex game-playing scenarios.

Introduction

As large language models (LLMs) become exceedingly better at performing traditional natural language processing tasks (Wang et al. 2018, 2019; Radford et al. 2019; Brown et al. 2020), they are now evaluated on more complicated tasks like recalling world knowledge (Hendrycks et al. 2020), reasoning (Talmor et al. 2018) and the ability to mathematics (Cobbe et al. 2021). A natural next evaluation setting for the ever-growing capabilities of these models is “game-playing” - a setting that requires not just the ability to do math and reasoning, but also planning, decision-making, and a deeper understanding of opponent behavior and state

of mind (Hu et al. 2024; Xu et al. 2024; Kosinski 2023). Poker is one such game that requires the above-described complex skill set.

Poker is an example of an incomplete information game (Harsanyi 1995) where a player has complete information about their own holdings, but incomplete information about the holdings of their opponents. The game requires a player to make strategic decisions based on their estimation of the opponent’s holdings by accounting for their actions, style of play, game situation, and possible future outcomes (Gupta 2023). This requires a complex combination of skills including math, reasoning, memory, long-term and short-term planning, and strategy, as well as a deep understanding of game theory and player behavior and psychology. Thus, the development of LLMs in game-playing settings like poker can potentially unlock higher cognitive capabilities in these models.

Existing AI systems for poker, commonly called poker “solvers”, play “game theory optimal” poker and have been shown to have superhuman performance (Brown et al. 2019; Brown and Sandholm 2019, 2018). Yet, these solvers have several limitations. (i) Firstly, poker solvers can take a long time to produce solutions for a spot¹, thus making them unusable for real-time use. (ii) Secondly, poker solvers are only able to calculate solutions for a limited and discrete set of scenarios, since the game tree in poker² can become unmanageably large. The game tree explosion also limits the use of solvers in multi-player settings. (iii) Thirdly, poker solvers do not necessarily provide the most profitable strategies in poker. Solvers are trained to play game theory optimal poker, which means they are trained to be unexploitable. However, being game theory optimal also means that these solvers do not take further advantage when their opponents are playing imperfectly. (iv) Finally, the solutions or strategies presented by solvers are abstract and not always interpretable which makes it hard to group and study these strategies.

With these limitations in mind and the growing cognitive capabilities of large language models, we explore the use of LLMs as poker solvers. Using LLMs for poker immediately allows for overcoming the disadvantages of traditional

*Correspondence to: akshat.gupta@berkeley.edu

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹A “spot” in poker is an intermediate point within a game

²Here we refer to Texas No-Limit Hold’em version of poker. This version of poker is also the focus of our paper.

poker solvers. LLMs can provide solutions for any spot instantly, and the time taken for a solution is not affected by the number of players or the number of possible scenarios (for example - bet sizes) being considered. LLMs also have the potential to take exploitative actions leading to more profitable strategies by taking into account the user’s history of playing using in-context learning (Brown et al. 2020). Finally, since LLMs are text-generation systems, they can also be used to explain their own solutions and strategies. While we do not explore these use cases in this paper, the potential advantages of using LLMs to play poker are exciting.

As we set out to improve the poker-playing abilities of LLMs, we find a big gap in the form of the absence of a reliable dataset to evaluate the poker-playing ability of LLMs. Poker is a game won or lost in the long term over a large sample of hands. For a strategy in poker to be considered profitable or a player to be considered better, we need to evaluate them over a statistically significant sample of poker hands. This poses a significant challenge in improving LLMs at poker as they need to be evaluated at every step of the way. Running game simulations for tens of thousands of hands after every fine-tuning iteration can significantly slow down this process. Thus, we need a dataset that can be used as a quick and accurate indicator of a model’s poker-playing skills. We thus present **POKERBENCH** - a new benchmark for a comprehensive evaluation of the poker-playing abilities of LLMs accompanied by a training dataset. The POKERBENCH benchmark consists of 11,000 spots for evaluating decision-making in poker, covering an exhaustive list of game situations including 1,000 pre-flop scenarios and 10,000 post-flop scenarios. The dataset and code can be found at <https://github.com/pokerllm/pokerbench>

In this paper, we make the following contributions:

- We introduce the POKERBENCH benchmark, meticulously developed in collaboration with trained poker players to encompass a full spectrum of essential poker scenarios. Grounded in the rigorous principles of game theory optimal poker, POKERBENCH evaluates models across a diverse array of important poker spots, encompassing both pre-flop and post-flop play, something that was missing in prior work (Gupta 2023).
- Using POKERBENCH, we perform an extensive evaluation of various state-of-the-art models including GPT-4 (OpenAI 2023), ChatGPT3.5 (OpenAI 2022) the Llama-3 series models (Meta 2024), the Llama-2 series models (Touvron et al. 2023) and Gemma-2B (Team et al. 2024) on POKERBENCH. We find that these models seriously underperform in the game of poker compared to their performance on other benchmarks, with GPT-4 being the best-performing model with an accuracy of 53.55%. We then fine-tune Llama-3-8B, Llama-2-7B, and Gemma-2B models on the accompanying training dataset and find that the performance of the best fine-tuned model improves significantly on fine-tuning.
- To test the usefulness of POKERBENCH as a quick and reliable measure of the poker-playing ability of LLMs, we have various fine-tuned model checkpoints with different scores on POKERBENCH play against each

other. We clearly see that models with a higher score on POKERBENCH are able to beat models performing poorly on a large sample of hands.

Related Work

Creating machines that are able to beat humans in gameplay settings has a long history. The first big success for an AI system at beating humans happened in 1997 when IBM’s Deep Blue (Campbell, Hoane Jr, and Hsu 2002) beat Garry Kasparov, the world number 1 chess player and an all-time great of the game. While this was a significant step, Garry Kasparov recovered from a single-game loss and ended up beating Deep Blue 4-2. More recently, AlphaGo (Silver et al. 2016) beat a Go world champion in 2016 in a comprehensive victory of 4-1. Different from Chess and Go, Poker represents a unique challenge by being an incomplete information game (Harsanyi 1995). While in Chess and Go, all pieces of the game are visible to all players; in poker, all information about the opponent’s holdings is not available to players. Significant progress was made in creating superhuman AI systems at poker between 2017-2019 with various algorithmic advances (Brown et al. 2019; Brown and Sandholm 2018, 2019).

With the increasing cognitive capabilities of LLMs, recent work has started exploring the use of LLMs in gameplay settings. Ruoss et al. (2024) recently explored the possibility of using LLMs to create non-search-based systems that play “grandmaster” level chess using LLMs. Gupta (2023) evaluated the poker-playing capability of GPT-4 and ChatGPT, which represents the first work exploring the possibility of playing poker using LLMs. While they concluded that LLMs were not good poker players, their analysis was limited to the first betting action in poker, also called “raise first in” or RFI spots. Huang et al. (2024) recently fine-tuned LLMs on data collected from PokerStars, a popular poker site, and showed that this improved the poker-playing abilities of LLMs. Yet, most of this development was done in the blind and models could only be evaluated at the end of the fine-tuning process by having them play a large number of hands. POKERBENCH allows for the constant development of LLMs as poker-playing agents and provides a universal benchmark for creating such systems.

Poker Preliminaries

In this paper, we study the most popular version of poker, called Texas No-Limit Hold’em (NLH), which has also been the focus of prior work (Gupta 2023; Brown and Sandholm 2019, 2018). The number of players in a Texas NLH game varies from 2-10, with the 6-player game being the most popular. Texas NLH poker is the epitome of decision-making under uncertainty and incomplete information with a near-infinite decision tree. In Texas NLH, each player is given two private cards, also called *hole cards*, that are only known to them, and five community cards that are visible to everyone, peeled in three rounds accompanied by four betting rounds. Each player thus possesses seven cards, out of which two cards are private and only known to the player, and five cards are common to everyone. Out of these seven

cards, a player presents their top 5 cards to make a winning combination. The betting rounds are as follows:

- **Pre-Flop:** This is the first betting round that happens right after the players see their hole cards. At this point, none of the five community cards have been opened.
- **Post-Flop:** After the previous betting round, three community cards are opened at once. This event is called the “flop”. This is followed by a betting round called the post-flop betting round.
- **Post-Turn:** After the post-flop betting round, a fourth community card is opened. The fourth card is called the “turn”, and is followed by a betting round.
- **Post-River:** After the post-turn betting round, the fifth and final community card is opened. The final card is called the “river”, followed by a final betting round.

A game of poker takes place in multiple iterations which involves distributing hole cards, community cards, and the above-mentioned four betting rounds, resulting in a player winning or losing a pot of chips at stake. One such iteration of the above process is commonly referred to as a “*hand*”.

Actions in Poker

There are four basic actions that a player takes while playing poker :

- **Check:** This action means that a player wants to continue playing in a betting round without wagering any chips. This can only be done if no wagers have yet been made previously in that betting round.
- **Bet:** The action of placing a wager in a betting round is called a *bet*. In poker, “*bet*” is specifically referred to a situation when no wager has yet been made in a betting round, then the first wager is called a “*bet*”.
- **Call:** The action of matching a wager made by a player previously is called a “*call*”.
- **Raise:** The action of wagering a larger amount of chips than the previous wager is called “*raise*”.
- **Fold:** The action of choosing not to match a previous wager (bet or raise) and thus giving up claim on the pot is called a “*fold*”. When a player folds, they are no longer part of the current *hand* being played.

The Unit of Measurement in Poker

Poker is a very popular game played at different stakes. While the lowest stakes in most casinos require a buy-in of 100\$, the buy-in can go as high as millions of dollars in high-stakes games around the world. Yet, every person in the game usually starts with the same amount of effective stack size³ according to a normalized unit of measurement.

The normalized unit of measurement in poker is the minimum amount of money a player is allowed to bet, and is called a “**big blind**” (BB). If the minimum bet amount in a game is 2\$, then 1BB = 2\$. If the minimum bet amount in a game is 200\$, then 1BB = 200\$. Poker games are usually classified by the amount required to buy 100 big blinds. So

³“stack size” is the total number of chips a player has

if a game is classified as a 100\$ buy-in game, this means that the minimum bet is 1\$. A standard starting stack size for most poker games and solver calculations is 100BB.

Throughout this paper, we will be describing gameplay using the BB unit. The win rate of a poker player is also defined in terms of big blinds, as the number of big blinds won per hundred hands played.

Game Theory Optimal Poker

“Game Theory Optimal” (GTO) strategy of playing poker refers to the optimal way of playing poker such that a player cannot be exploited by their opponents. A GTO strategy is usually a balanced strategy where the opponent is unable to correlate the actions of a player with their holdings. For a simplistic pedagogical example to illustrate this, if Player A’s strategy is to go all-in with only pocket aces (AA), which are the best starting hands in poker, then Player B can exploit Player A by folding against Player A’s all-in with any two cards that are not pocket aces (AA). As a result, Player A is unable to extract value from their strongest hand. Thus, in GTO play, Player A should go all-in with a wider selection of hands, and also not go all-in with pocket Aces every single time. This is a more optimal strategy in the long run against all kinds of players since it becomes challenging for our opponent to narrow down player A’s exact holding. A widely used technical term for the philosophy behind such a strategy is called playing a “balanced” game.

The POKERBENCH Benchmark

We carefully design the POKERBENCH benchmark to thoroughly evaluate the poker-playing abilities of LLMs with an exhaustive coverage of many types of poker spots. The aim of creating POKERBENCH is to evaluate LLMs at playing game theory optimal poker in a quick and reliable way. POKERBENCH is designed such that the higher a model scores on our benchmark, the better it is at playing optimal poker.

We define a “spot” to be a combination of hole cards that a player has, the board⁴, and the actions taken in the different betting rounds. Similar to chess, the search space for Texas NLH poker is extremely large. Thus, we build POKERBENCH with two main balancing principles: diversity and simplicity. We want to be able to evaluate an LLM thoroughly by having them play a wide category of scenarios while keeping a reasonable total inference time to enable quick development. The following subsections explain our design choices in condensing the enormous search space by careful filtering and pruning using principles of optimal poker play.

The POKERBENCH benchmark is created based on 6-max player Texas NLH poker. It consists of two separate sets - a pre-flop and a post-flop evaluation dataset. Pre-flop games are usually very different from post-flop play and have a very different distribution of decisions, which is why we decided

⁴“the board” is a term used to refer to the set of community cards in play in a particular hand since they traditionally get displayed on a board

Dataset Type	# Examples	Raise (%)	Fold (%)	Call (%)	Check (%)	Bet (%)
Pre-flop Train Set	60,000	12.7	47.5	39.6	0.3	N/A
Post-flop Train Set	500,000	15.1	25.0	25.0	25.0	9.9
Pre-flop Benchmark	1,000	25.0	25.0	25.0	25.0	N/A
Post-flop Benchmark	10,000	15.5	25.0	25.0	25.0	9.5

Table 1: Summary of POKERBENCH dataset, including dataset size and action distributions for pre-flop and post-flop datasets in both training and benchmark settings.

to separate the two types of scenarios. We have 1,000 evaluation examples in the benchmark for the pre-flop game and 10,000 evaluation examples for the post-flop game. To create this dataset, we use the GTO strategies from GTOWizard for the pre-flop game and WASM-Postflop to solve GTO strategies for the post-flop game.

Pre-flop Action Selection

After each player is dealt with two hole cards, the first betting round can have an exponentially large number of decisions being made. For example, a player can decide to bet a certain amount, and a following player can decide to bet an even larger amount (called a “raise” in poker), and this raising and re-raising can happen with an exploding number of permutations, each with a different bet and raise size. Existing pre-flop GTO strategies exist for all of these scenarios. In POKERBENCH, we only consider scenarios where a maximum of two raises have happened in the pre-flop betting round. This includes scenarios where (i) all players fold, (ii) only one player bets chips and other players either call or fold, and (ii) one player bets chips, a second player raises that bet by a higher wager followed by only calls or folds (also called “3-bet pots”). This covers the majority of the possible pre-flop scenarios that are considered a viable GTO play as most pots do not go beyond a single raise.

Board Selection

A “board” is the list of community cards that show up during a poker game. There are in total 5 community cards, where each card can be any of the 52 cards in the deck. Thus, we have a total of ${}^{52}C_5 \approx 311$ million possible boards that can show up. It is impossible to evaluate our model on that many boards, which is why we group the boards into 11 classes, called **textures**. These board textures cover the most common situations on the flop.

One of the most commonly studied board textures on the flop is what we call “single-broadway-dry”. As a reminder, “flop” is a term used to describe the first three community cards. The term “broadway cards” is used to refer to the cards in the set $\{A, K, Q, J, 10\}$, that is, it refers to the five strongest cards in a suit. The term “single-broadway” refers to the opening of exactly one Broadway card on the flop. The term “dry” is used to describe flops in poker which do not have a lot of possibilities to make different winning hands. For example, a board $\{K_h, 7_d, 2_s\}$, where the subscripts show the suit of the cards⁵, is a typical dry board in poker. The board does not have any repetition of suit,

⁵h : hearts, d: diamonds, s: spades, c: clubs

which means it is less likely to form flushes⁶, and does not have numbers close to each other, which means it is less likely to form straights⁷. The strategy a player should use to play on a $\{K_h, 7_d, 2_s\}$ board is going to be very similar to strategy on similar boards like $\{K_h, 7_d, 4_s\}$, $\{A_h, 8_d, 3_s\}$, $\{Q_h, 8_d, 2_s\}$, and hence all such boards are group together into a single texture called “single-broadway-dry”. All board textures used to create POKERBENCH are shown in the Appendix⁸ (Zhuang et al. 2025).

We randomly sample an equal number of flops from each of the textures to create the dataset of boards on which the game is evaluated. Then to select the turn cards for each scenario, we select cards that would continue to cover the most new ground. For example, if a flop contains a flush draw⁹, some possible turn cards were selected to either complete or not complete the flush draw. This process ensures an equal coverage of the different categories of boards and results in a more informed evaluation compared to random sampling, especially at small sample sizes.

Selecting Hole Cards

GTO play dictates to have nondeterministic strategies since if a deterministic strategy exists in a player’s game, it becomes exploitable. For example, if a player always goes all-in when they have aces (AA), then all other players are likely to fold, ending up in an unprofitable play. Thus, a balanced, game theory optimal way of playing any spot in poker is to have at least two actions chosen with certain probabilities at the time of play. For selecting hole cards, we choose spots where there is a clear dominant strategy as the best action. For example with a board like $\{A_h, K_h, 8_d, 3_s, 10_c\}$, facing a small bet from the opponent, private cards that have a clear dominant strategy are hands like Q_d, J_s of raising, whereas hands like $10_d, 9_d$ have a more mixed strategy of calling or folding. We decided to filter hole cards by selecting action lines that choose one dominant action with greater than 50% probability.

Dataset Summary

The POKERBENCH benchmark consists of 1k evaluation spots for pre-flop and 10k evaluation spots for post-flop play. Along with the evaluation benchmark, we also release

⁶when all five cards have the same suit

⁷a consecutive sequence of 5 cards

⁸The Appendix for this paper can be found in the arxiv version of this paper.

⁹A “flush draw” exists when two cards of the same suit are opened on the board, thus making the possibility of a flush if another card of the same suit is dealt

a training set containing 60k pre-flop spots and 500k post-flop spots. The distribution of actions for the POKERBENCH benchmark and the training set can be seen in Table 1.

For the pre-flop training set, we preserve the original distribution of actions that a player can take because the search space there is relatively small, and so an exhaustive strategy that covers most possible scenarios is learnable. For the post-flop training set, we experiment with a few different sampling strategies and find that resampling the action distribution using a balanced strategy results in the best fine-tuning performance. Specifically, a balanced sample helps ensure coverage of the search space. It also keeps the model from learning a naive strategy that decides to “fold” most of the time, a possible situation if we preserve the original distribution where the action of folding predominates.

However, for the POKERBENCH benchmark, for both pre-flop and post-flop settings, we choose a balanced sampling strategy. Precisely, there is an equal percentage of samples each with the correct decision labels being fold, call, check, or bet/raise. We cannot preserve the highly imbalanced distribution for evaluation because a naive strategy of folding all hands will lead to an accuracy of close to 90%. Because of this, we downsample the fold action significantly to remove hands that do not provide much signal. For example, all starting hands like {72}, {83}, {93}, {84} should be folded, but not all of these hands need to be kept in the evaluation set to get the signal that the model is folding such low-value hands. Whereas hands like {AK}, {KQ}, {AA}, {JJ} should be played before the flop, it is important to analyze the model on a larger set of such hands, as these hands are also more profitable.

Experiments

After carefully curating POKERBENCH, we move on to evaluating the poker-playing ability of LLMs.

Evaluation Metrics We use the following two metrics to evaluate the poker skills of LLMs by evaluating their responses on POKERBENCH:

- **Action Accuracy (AA):** Action accuracy measures if LLMs can take game theory optimal actions (fold, raise, bet, etc.) for a given spot.
- **Exact Match Accuracy (EM):** Actions like *bet* and *raise* are followed by the amount of the bet and raise. Hence, exact match accuracy also considers if the wager amount is game theory optimal.

Are Modern LLMs Good at Playing Optimal Poker?

We evaluate GPT-4 (OpenAI 2023), ChatGPT 3.5 (Ouyang et al. 2022), Llama-3 models (8B, 70B) (Meta 2024) and Llama-2 70B (Touvron et al. 2023) on POKERBENCH. We take inspiration from the evaluation protocol for the popular MMLU dataset (Hendrycks et al. 2020) and evaluate these models in a few-shot setting. While MMLU is a multiple-choice question-answering dataset, POKERBENCH is not. Thus we require the model to generate the action and the exact bet amount for a given spot. For few-shot examples, we




Figure 1: Training dynamics of Llama-3-8B on POKERBENCH training dataset.

select one example randomly for each possible action from the training dataset and add it to the context. Thus our few-shot setting contains 5 examples in context, one for each action as shown in Table 1. An example prompt can be found in the Appendix (Zhuang et al. 2025). For generating text, we set temperature = 0.1 and top-p = 0.95 to generate the most probable answer to get statistically stable results. We use the OpenAI API for evaluating OpenAI models and TogetherAI API for evaluating models from the Llama series.

The results for evaluation on POKERBENCH can be found in Table 2. We use chat or instruct models for this evaluation as applicable¹⁰. GPT-4 outperforms all other models both in pre-flop and post-flop play. The second best performing model is Llama-3-70B with its performance on the benchmark being significantly lower than GPT-4. A surprising thing to note is that ChatGPT 3.5 performs comparably to a significantly smaller Llama-3-8B model. Llama-3 outperforms Llama-2 in generating correct bet/raise amounts (higher EM score) while this lead diminishes when generating optimal actions (similar AA score). While we also tried to evaluate the smaller models from the Llama-2 series, they are unable to follow poker instructions, which is why we do not report their accuracy scores. The benchmark results show that all modern LLMs significantly lack in their ability to play game theory optimal poker and that there is a lot of room for improvement.

Fine-Tuning LLMs into Better Poker Players

In the previous section, we saw that state-of-the-art LLMs are not good at playing poker and significantly underperform on POKERBENCH compared to other tasks they are evaluated on (Hendrycks et al. 2020; Talmor et al. 2018; Wang et al. 2018, 2019). A possible reason for this could be the complexity of the game which requires a multitude of skills coming together. To improve the poker-playing ability of LLMs, we fine-tune the model on a subset of the accompanying training set released with this paper. The training set consists of 30k pre-flop spots and 80k post-flop spots. We

¹⁰ Available open-sourced base models performed sub-optimally compared to chat/instruct models which is why we do not report the scores of the base models.

EVALUATION TYPE	MODEL	Overall Accuracy		Post-Flop Accuracy		Pre-Flop Accuracy	
		EM \uparrow	AA \uparrow	EM \uparrow	AA \uparrow	EM \uparrow	AA \uparrow
Pre-Trained Models (Few-Shot)	LLAMA-3 (8B)	26.02	40.03	14.96	31.25	37.77	49.30
	LLAMA-2 (70B)	36.48	48.30	32.95	41.11	40.20	55.90
	LLAMA-3 (70B)	39.16	49.78	34.30	45.40	44.30	54.40
	CHATGPT 3.5	29.96	39.69	18.75	34.19	41.80	45.50
	GPT-4	53.55	65.54	52.18	62.69	55.00	66.50
Fine-Tuned Models (Zero-Shot)	GEMMA (2B)	51.84	62.74	41.57	52.94	62.70	73.10
	LLAMA-2 (7B)	78.11	79.91	76.52	79.55	79.80	80.30
	LLAMA-3 (8B)	78.26	80.64	76.52	79.07	80.10	82.30

Table 2: Performance of various pre-trained and fine-tuned LLMs on POKERBENCH.

have a much larger amount of post-flop spots in the training set since there are more possible permutations of post-flop spots compared to the pre-flop game. We choose this subset to balance the number of examples in the pre-flop and post-flop stages.

We fine-tune three models on the POKERBENCH training dataset - Llama-3-8B, Llama-2-7B, and Gemma-2B. We fine-tune the model for 5000 optimization steps with a batch size of 128 and a learning rate of 1e-6, which is lower than the pre-training learning rate for most of these models¹¹. The model goes through one epoch of the dataset approximately every 900 gradient steps. The performance of the fine-tuned models on POKERBENCH can be seen in Table 2. We see that the Llama-3-8B model improves considerably on the POKERBENCH benchmark, outperforming GPT-4. The performance for Llama-3-8B and Llama-2-7b are quite similar with Llama-3-8B performing slightly better, whereas the Gemma-2B model falls behind on the benchmarks. Thus going forward, we pick Llama-3-8B for further investigation. Figure 1 shows the loss curves for Llama-3-8B plotted with the accuracy of different checkpoints during the fine-tuning process. The loss curves for Llama-2-7B and Gemma-2B are presented in the Appendix (Zhuang et al. 2025) and have very similar training dynamics.

Heads-Up Games between Fine-tuned Checkpoints Finally, we test if the scores on POKERBENCH translate into actual performance. To do so, have different models with different scores play an actual poker game against each other. To do so, we pick three different Llama-3-8B model checkpoints with different accuracy scores on POKERBENCH. We simulate a heads-up¹² game between the three players. The details of the selected players are shown in Table 3. The checkpoint name refers to the number of gradient updates that have happened in the fine-tuning process at the time of selection. For example, the Llama-3-8B 1600 has gone through 1600 gradient update steps on the training dataset. Going forward, we will refer to the players by their checkpoint names as player 800, player 1600, and player 5000.

The heads-up games are played with two modifications -

¹¹ A higher learning rate resulted in unstable fine-tuning

¹² A “heads-up” match in poker is a game where two players play against each other.

CHECKPOINT NAME	EM	AA
Llama-3-8B 800	58.62	61.32
Llama-3-8B 1600	65.96	69.60
Llama-3-8B 5000	78.26	80.64

Table 3: Selected checkpoints for heads-up games.

GAME	WINNER	WIN RATE
5000 vs 1600	5000	24.79 bb/100
5000 vs 800	5000	50.88 bb/100
1600 vs 800	1600	34.99 bb/100

Table 4: Results of heads-up games against players 5000, 1600 and 800.

the stacks of the players are always reset to 100BB to match the GTO strategies¹³, and the seating positions of the players are randomly assigned before each hand. To create statistically significant results, we have each player play against the other for 50k hands. The win rate is calculated with a metric called big blinds per hundred hands, or bb/100, defined as the number of big blinds won per hundred hands played. As a reminder, a big blind is the basic unit of measurement in poker and represents the minimum bet that can be made. As a reference, under a statistically significant sample size, a win rate of about 4.0 bb/100 is considered very good in poker. A win rate in the range of 5.0-9.0 bb/100 is considered exceptional and is a sign of a clearly dominant player. At this win rate, players are suggested to move at a higher level of stake.

The results for the heads-up tournament can be seen in Table 4. We see that player 5000, who also has the highest score on POKERBENCH, is a significantly superior player compared to the other players. Player 5000 wins its games against players 1600 and 800, while the same is true for player 1600 in its game against player 800. These experiments show that a higher score on POKERBENCH actually translates to a higher win rate. These experiments also present the value of POKERBENCH as an evaluation tool during development. Instead of having models play against each other for 50k hands for statistical significance, we can

¹³GTO strategies are calculated at 100BB stacks

GAME STAGE	WINNER	WIN RATE
Overall	GPT-4	22.20 bb/100
Pre-flop	5000	17.80 bb/100
Flop	GPT-4	34.97 bb/100
Turn	GPT-4	24.04 bb/100
River	5000	18.02 bb/100

Table 5: Results of the 1,000 heads-up games between Player 5000 and GPT-4. As the sample size is limited, the absolute value of win-rate is less meaningful. Instead, we provide a playing style analysis to better understand the models’ performances.

evaluate the models on a much smaller set of examples chosen from POKERBENCH.

Fine-tuned Model vs. GPT-4 Next, we perform a heads-up game between the best fine-tuned model checkpoint and the best pre-trained model, namely GPT-4, using a same setting. Due to inference cost constraints, the two models played for 1,000 hands only. The result for this tournament can be seen in Table 5.

We notice an unorthodox phenomena, that despite achieving higher test accuracy, our fine-tuned Llama model was outplayed by GPT-4. To investigate this, we separate the win-rate by stage of the game (pre-flop, flop, turn, and river), and conduct an in-depth analysis of the game log. Overall, we identify spots where Llama is effectively learning from the training set and winning by playing more GTO than GPT-4. For instance, as shown in Appendix (Zhuang et al. 2025), Llama is winning in the pre-flop stage by raising a wider range of hands (which is more optimal) and squeezing GPT-4 out of the pot.

However, one reason of GPT-4’s winning is that it adopts a sub-optimal strategy called “donking”. In a nutshell, it is a losing move in the long run, hence not frequently seen in GTO strategies (more details can be found in the Appendix (Zhuang et al. 2025)). But our dataset assumes GTO play from both players. Consequently, this move is unfamiliar to the fine-tuned model, constituting a big part of Llama’s losses as GPT-4 wins 100% of the hands where it donks.

Also, neither player is playing at an optimal action frequency: GPT-4 is being over-aggressive while Llama is over-passive. A good example of this can be illustrated by a rock-paper-scissors game. Imagine player A chooses rock 100% of time, while player B chooses scissors 50% of time and rock 50% of time. We can see that even though player B has a strategy closer to the optimal solution¹⁴, it will never win against player A’s less optimal strategy. Similarly, while GPT-4 has a strategy farther from GTO (as assessed by our dataset), it is still able to win against Llama by coincidentally executing a strategy that happens to exploit Llama’s suboptimal decisions. A more detailed analysis can be found in the Appendix (Zhuang et al. 2025).

Discussion There is a fine line between winning against an imperfect strategy versus minimizing losses against a GTO

¹⁴The optimal strategy in rock-paper-scissors is playing each option 33% times.

strategy. The two methods of evaluation used in this paper are evaluating two different things. The gameplay is evaluating how well a strategy performs against an imperfect strategy, while our dataset is evaluating the how well a strategy perform against a GTO strategy. One of the advantage about training the model to approach GTO strategy is that being GTO is a sufficient condition for winning against any imperfect strategy. But when a non-perfect GTO strategy is learned (for instance in the case of our fine-tuned model since it is not perfectly GTO), the relationship between the two could become complex. Yet, we argue that succeeding on our dataset is a necessary condition for minimizing loss against a GTO strategy because our dataset is a compilation of a diverse set of samples that any near GTO strategy must agree on (and non-GTO strategy is guaranteed to lose in the long run). We also believe that this finding is also a limitation of the simple supervised fine-tuning in training LLMs to effectively succeed in poker. We would like to use this example to encourage further research in improving the adaptability of language models in game environments.

Conclusion

In this paper, we present POKERBENCH - a comprehensive benchmark that evaluates optimal poker-playing ability of LLMs. We evaluate multiple state-of-the-art language models and show that current LLMs fail significantly in playing optimal poker. We also fine-tune Llama-3-8B among other models on the accompanying training dataset and show that resultant models can outperform much larger models. We also show that the scores on POKERBENCH actually translate to superior poker-playing skills by evaluating models of different scores through game simulations over 50k hands. Thus, POKERBENCH represents a quick and reliable measure of the optimal poker-playing ability of large language models as well as a comprehensive benchmark to study the progress of LLMs in this domain. This study not only explores the potential of LLMs in strategic game-playing but also presents a benchmarks to evaluate higher-level cognitive capabilities of LLMs in complex game-playing situations.

References

- Brown, N.; Lerer, A.; Gross, S.; and Sandholm, T. 2019. Deep counterfactual regret minimization. In *International conference on machine learning*, 793–802. PMLR.
- Brown, N.; and Sandholm, T. 2018. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374): 418–424.
- Brown, N.; and Sandholm, T. 2019. Superhuman AI for multiplayer poker. *Science*, 365(6456): 885–890.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Campbell, M.; Hoane Jr, A. J.; and Hsu, F.-h. 2002. Deep blue. *Artificial intelligence*, 134(1-2): 57–83.

- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Gupta, A. 2023. Are ChatGPT and GPT-4 Good Poker Players?—A Pre-Flop Analysis. *arXiv preprint arXiv:2308.12466*.
- Harsanyi, J. C. 1995. Games with incomplete information. *The American Economic Review*, 85(3): 291–303.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Hu, S.; Huang, T.; Ilhan, F.; Tekin, S.; Liu, G.; Kompella, R.; and Liu, L. 2024. A survey on large language model-based game agents. *arXiv preprint arXiv:2404.02039*.
- Huang, C.; Cao, Y.; Wen, Y.; Zhou, T.; and Zhang, Y. 2024. PokerGPT: An End-to-End Lightweight Solver for Multi-Player Texas Hold'em via Large Language Model. *arXiv preprint arXiv:2401.06781*.
- Kosinski, M. 2023. Evaluating Large Language Models in Theory of Mind Tasks. *arXiv e-prints*, arXiv–2302.
- Meta. 2024. Llama 3 - <https://ai.meta.com/blog/meta-llama-3/>.
- OpenAI. 2022. ChatGPT - <https://openai.com/blog/chatgpt#OpenAI>.
- OpenAI. 2023. GPT-4 Technical Report - <https://cdn.openai.com/papers/gpt-4.pdf>.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35: 27730–27744.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Ruoss, A.; Delétang, G.; Medapati, S.; Grau-Moya, J.; Wenliang, L. K.; Catt, E.; Reid, J.; and Genewein, T. 2024. Grandmaster-level chess without search. *arXiv preprint arXiv:2402.04494*.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587): 484–489.
- Talmor, A.; Herzig, J.; Lourie, N.; and Berant, J. 2018. Commonsenseqa: A question answering challenge targeting commonsense knowledge. *arXiv preprint arXiv:1811.00937*.
- Team, G.; Mesnard, T.; Hardin, C.; Dadashi, R.; Bhupatiraju, S.; Pathak, S.; Sifre, L.; Rivière, M.; Kale, M. S.; Love, J.; et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models, 2023. URL <https://arxiv.org/abs/2307.09288>.
- Wang, A.; Pruksachatkun, Y.; Nangia, N.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. 2019. SuperGlue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*, 32.
- Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Xu, X.; Wang, Y.; Xu, C.; Ding, Z.; Jiang, J.; Ding, Z.; and Karlsson, B. F. 2024. A Survey on Game Playing Agents and Large Models: Methods, Applications, and Challenges. *arXiv preprint arXiv:2403.10249*.
- Zhuang, R.; Gupta, A.; Yang, R.; Rahane, A.; Li, Z.; and Anumanchipalli, G. 2025. PokerBench: Training Large Language Models to become Professional Poker Players. *arXiv preprint arXiv:2501.08328*.