



Mastering Gomoku with AlphaZero: A Study in Advanced AI Game Strategy

Wen Liang

University of California, San Diego
La Jolla, CA 92093, wel245@eng.ucsd.edu

Brian Whiteaker

University of California, San Diego
La Jolla, CA 92093, bwhiteak@ucsd.edu

Hua Shao

University of California, San Diego
La Jolla, CA 92093, h5shao@eng.ucsd.edu

Chao Yu

University of California, San Diego
La Jolla, CA 92093, chy018@eng.ucsd.edu

Inyoung Huh

University of California, San Diego
La Jolla, CA 92093, i1huh@ucsd.edu

Youzhi Liang

Stanford University
Stanford, CA 94305, youzhi@stanford.edu

Abstract

This study delves into the application of the AlphaZero algorithm to Gomoku, a classic board game. Unlike traditional AI methods, AlphaZero learns and strategizes without human input. Our research contrasts AlphaZero's innovative approach with the Monte Carlo tree search technique, highlighting its advanced capabilities in strategic decision-making. The findings reveal AlphaZero's remarkable proficiency in mastering the complexities of Gomoku, marking a significant advancement in artificial intelligence's role in game strategy and decision-making. This paper provides a comprehensive analysis of AlphaZero's learning process and strategic execution in Gomoku, offering insights into the future of AI in strategic gaming.

Introduction

Reinforcement learning (RL) is a pivotal and rapidly advancing domain within contemporary artificial intelligence research. It offers a distinctive framework wherein agents progressively improve their performance not through explicit instruction but through continual interaction with their surroundings [1, 2]. As these agents take actions in a given environment, they are provided feedback in the form of either rewards for desirable actions or penalties for undesirable ones [3]. This trial-and-error mechanism aids the agents in understanding the consequences of their actions and refining their strategies accordingly. The primary goal of RL is to determine an optimal strategy, often termed as a "policy", which instructs the agent on the best possible action to take in any given situation. The optimal policy is one that, when followed, will lead to the maximization of the cumulative rewards over a period, ensuring that the agent's actions result in the most favorable outcomes in its environment. Board games, with their intricate complexities and well-defined reward structures, make a fitting domain for RL, offering a perspective ripe for academic inquiry. MCTS (Monte Carlo tree search) [4] has emerged as a leading algorithm for decision-making in these complex environments. Recently, deep learning catalyzed groundbreaking developments in diverse research domains, from computer vision and natural language processing to state-of-the-art recommender systems [5, 6, 7]. It constructs a search tree by exploring possible future moves, using statistical sampling to evaluate the potential outcomes. It constructs a search tree by exploring possible future moves, using statistical sampling to evaluate the potential outcomes. Bridging the gap between RL and MCTS, the original AlphaGo [8] algorithm showcased a fusion of deep learning and tree search techniques, revolutionizing the game-playing AI landscape. This groundbreaking approach further evolved with the introduction of AlphaZero [9], which uses zero human

knowledge and experience in this game and removes the supervised learning stage. It allows the algorithm to master the game with self-learning only.

Gomoku, often referred to as "Five in a Row," is usually played on a 15x15 grid (though variations can feature larger grids). The game's objective is straightforward yet captivating: two players, typically designated as black and white, take turns placing stones on the board with the aim to align five of their own stones consecutively in a vertical, horizontal, or diagonal line. The game's seemingly simple rules mask a depth of strategy. Early moves tend to be concentrated around the center of the board, providing players with maximal opportunities to expand and form their sequences. As the game progresses, the board transforms into a complex battleground of potential sequences, blocked attempts, and intricate traps. The nature of the game allows for both defensive and offensive tactics. A player might focus on preventing their opponent from completing a sequence, or strategically placing their stones to create multiple potential winning avenues simultaneously. Its straightforward rules combined with its profound complexity make it an ideal candidate for studying artificial intelligence's prowess and potential in mastering classic board games. In recent decades, there have been concerted efforts to solve Gomoku using computational methods. One notable attempt was by Allis [10], who employed proof-number search algorithms to analyze specific game positions and paths, making significant strides in understanding the game's complexities. Another significant contribution came from Chen [11], who utilized pattern recognition and threat-space search techniques to advance AI capabilities in Gomoku, offering a fresh perspective on potential winning strategies.

Driven by the recent monumental strides in board game artificial intelligence, especially the unparalleled triumphs of the AlphaZero algorithm, we were compelled to believe that harnessing this cutting-edge approach for the game of Gomoku was not only feasible but imperative. In embarking on this ambitious journey, our contributions to the Gomoku AI research landscape manifest in two significant dimensions:

1. We generalized the AlphaZero approach for the Gomoku game, achieving impressive results. Initiating from a state of random play, and without any domain knowledge apart from the game rules, our model swiftly learned a winning strategy on a 6x6 table after just a few hours of training on an economical GPU.
2. We embarked on an extensive research endeavor, wherein we juxtaposed the efficacy of our refined AlphaZero methodology against a conventional method that exclusively leverages the Monte Carlo tree search (MCTS). Our aim was to critically assess how these two distinct techniques fare in terms of both efficiency and effectiveness under comparable conditions, to shed light on their relative strengths and potential areas of improvement.

Method

2.1 Value and Policy Network

Deep neural network in AlphaZero [9] and AlphaGo [8] often employs two primary neural networks: the Value Network (V) and the Policy Network (π).

- Value Network (V): This network estimates the value of a given state, i.e., the expected outcome from that state. Formally, for a given state s , $V(s)$ predicts the expected outcome, with values close to +1 indicating favorable outcomes for the player and values close to -1 indicating unfavorable outcomes.

$$V(s) \approx E[r|s] \quad (1)$$

where E is the expectation and r is the eventual reward.

- **Policy Network (π):** This network provides a probability distribution over all possible moves from a given state. For a state s and an action a , $\pi(a|s)$ represents the probability of taking action a as a highly optimized game player.

$$\pi(a|s) = P(a \text{ is the best move} | s) \quad (2)$$

The neural network structure we use is shown in Figure 1.




Figure 1: Value and Policy Networks

2.2 Monte Carlo Tree Search (MCTS)

Monte Carlo Tree Search (MCTS) stands out as a revolutionary algorithm, reshaping decision-making processes in intricate environments through the methodical construction of a search tree. At its core, the algorithm meticulously evaluates prospective game moves, striking a harmonious equilibrium between the dual tenets of exploration (unearthing new moves) and exploitation (leveraging known advantageous moves). The integration of Policy and Value networks into this framework bestows it with unparalleled depth and precision:

- The Policy Network, through its discerning output, serves as the beacon guiding the expansion of the search tree. Instead of branching out indiscriminately, it casts the spotlight on moves radiating promise and potential, ensuring that the exploration process remains strategic and focused.
- On the other hand, the Value Network steps in as an adept evaluator, meticulously scrutinizing leaf nodes within the tree. This network diminishes the traditional reliance on random rollouts for evaluation, infusing the process with a heightened level of precision. This capability not only speeds up the evaluation but also endows it with a more profound insight into the game's dynamics.

In essence, the Policy Network acts as a compass, navigating the vast possibilities in the MCTS landscape and directing it towards potentially rewarding paths. Simultaneously, the Value Network functions as an astute analyst, swiftly gauging the potential outcomes of different game scenarios. Their synergistic interplay ensures that the search process within MCTS remains both streamlined and enriched, leading to decisions that are both efficient and strategically sound.

2.3 Environment and Supervised Learning

Within the realm of reinforcement learning, our agent actively interacts with a specially-designed Gomoku gaming environment, drawing feedback in the form of rewards or penalties based on its moves. As depicted in Figure 2, we meticulously implemented a Gomoku game board that closely mirrors traditional gameplay. Given the computational overhead associated with larger boards, we strategically focused our experimental investigations on boards sized 6×6 , targeting a 4-in-a-row win condition, and 8×8 , targeting the standard 5-in-a-row. To provide a comprehensive representation of each game state, we innovatively devised four distinct binary feature matrices. These matrices encapsulate essential game facets, including the current player's move, the adversary's move, the most recent move, and the initiating player. Notably, these matrices not only serve as a holistic game state representation but also act as pivotal input layers for our deep learning neural network. In terms of game mechanics, we faithfully incorporated Gomoku's conventional victory criteria. Moves are delineated at board intersections, eschewing placement within board squares. Traditionally, the white player initiates the game, with both players alternating their moves in succession until a conclusive game outcome is achieved. The essence of the game revolves around players positioning a stone of their chosen color on an unoccupied intersection. Triumph is heralded by the first player to strategically place five of their stones consecutively, irrespective of orientation - horizontal, vertical, or diagonal. In the rare scenario where the board reaches saturation without either contender achieving the coveted five-in-a-row, the game is ceremoniously declared a stalemate. While Gomoku might superficially seem straightforward, it belies a profound strategic depth, characterized by its myriad winning patterns and tactical nuances.

Gomoku's strategic intricacies are underscored by the delicate balance and importance of specific board configurations, notably the 'threes' and 'fours'. These patterns play pivotal roles in dictating the pace and outcome of a match. When leveraged effectively, they can swiftly shift the advantage to one player, often pushing the opponent into a corner from which recovery becomes arduous.

Diving into the nuances, the 'four' configuration is a fascinating tactical alignment where four stones of identical color stand in unison, beckoning a potential game-winning fifth stone in the subsequent move. The looming threat of this alignment is palpable, sending clear signals of an impending victory. Recognizing this, an opponent is thrust into a defensive mode, compelled to respond instantly. Failure to address this sequence by obstructing the alignment invariably results in a loss, testament to its lethal efficacy.

Equally compelling, yet distinct in its strategic implications, is the 'fork' configuration. In this maneuver, a player crafts a masterstroke with a single move, spawning two formidable attack sequences in tandem. The duality of the threat is what sets the 'fork' apart: it presents a dual quandary that the opponent must navigate. The challenge is steep; it's nearly impossible to stymie both threats concurrently. Thus, successfully engineering a fork is often tantamount to clinching the game, rendering the adversary powerless in the face of this double-edged assault. To truly appreciate the visual elegance and tactical profundity of these configurations, one can refer to Figure 3 and Figure 4. These illustrations vividly capture the essence of 'threes', 'fours', and the enigmatic 'fork', underscoring their pivotal roles in the beautiful complexity of Gomoku.




Figure 2: Gomoku game board from our implementation




Figure 3: 'Four' winning pattern

Result

Our experimentation painted an optimistic picture when applying the AlphaZero methodology to the Gomoku game. Significantly, our rendition not only succeeded but boasted an impeccable 100% victory rate as the initiating player during self-play assessments. Moreover, as the succeeding player, the algorithm manifested a keen aptitude for defense, coupled with a proactive stance towards




Figure 4: 'Fork' winning pattern

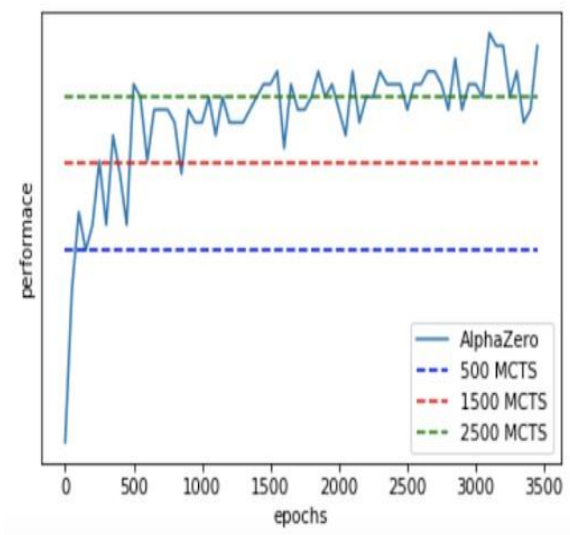


Figure 5: Compare AlphaZero with MCTS

identifying and capitalizing on counterattack chances. A detailed exemplification of this nuanced behavior is cataloged in Appendix I.

In our research, we undertook an in-depth comparative study, juxtaposing the performance of the

AlphaZero approach against the traditional Monte Carlo tree search (MCTS) method. To furnish a comprehensive perspective, we analyzed a spectrum of iterations, spanning from 500 to 2500 in number. The ensuing patterns and performance distinctions are graphically represented in Figure 5. The empirical data unambiguously accentuates AlphaZero's superior efficacy, as it continually eclipses the performance benchmarks set by the MCTS method.



Conclusion

We meticulously crafted a simulation environment tailored for the game of Gomoku and, within this context, developed a specialized agent. By adopting and integrating the AlphaZero methodology into our platform, we were able to achieve not only functional outcomes but also results that exceeded our initial expectations. These findings underscore the profound efficacy of the AlphaZero technique in mastering intricate board games, such as Gomoku, demonstrating its versatility and robustness in diverse gaming scenarios.

References

- [1] Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285, 1996.
- [2] Yuxi Li. Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274*, 2017.
- [3] Marco A Wiering and Martijn Van Otterlo. Reinforcement learning. *Adaptation, learning, and optimization*, 12(3):729, 2012.
- [4] Cameron B. Browne, Edward Powley, Daniel Whitehouse, Simon M. Lucas, Peter I. Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1–43, 2012.
- [5] Wen Liang, Youzhi Liang, and Jianguo Jia. Miamix: Enhancing image classification through a multi-stage augmented mixed sample data augmentation method, 2023.
- [6] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023.
- [7] Wen Liang, Zeng Fan, Youzhi Liang, and Jianguo Jia. Cross-attribute matrix factorization model with shared user embedding, 2023.
- [8] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, January 2016.
- [9] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of go without human knowledge. *Nature*, 550:354–, October 2017.
- [10] L. Victor Allis. Searching for solutions in games and artificial intelligence. 1994.
- [11] Ken Chen and Zhixing Chen. Static analysis of life and death in the game of go. *Information Sciences*, 121(1):113–134, 1999.

Appendix I

This section showcases a set of self-play match results on a 6×6 board, where a sequence of four consecutive pieces leads to a win. In the presented board, the 'x' pieces represent the first player, while the 'o' pieces symbolize the second player. The darker shades on the board correspond to positions where the policy network predicted a stronger advantage.

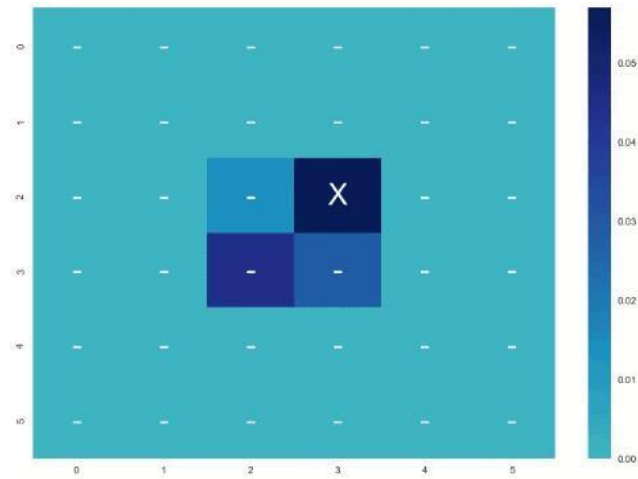


Figure 6: Example step 1

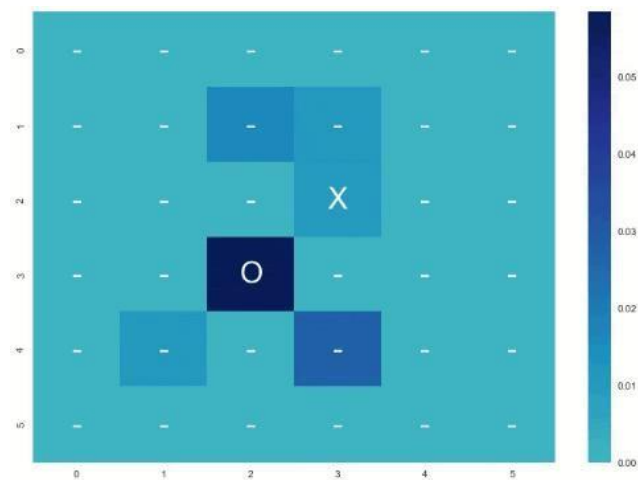


Figure 7: Example step 2

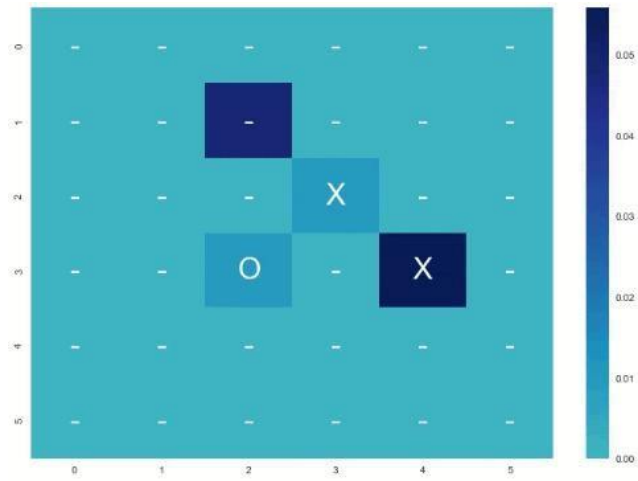


Figure 8: Example step 3




Figure 9: Example step 4

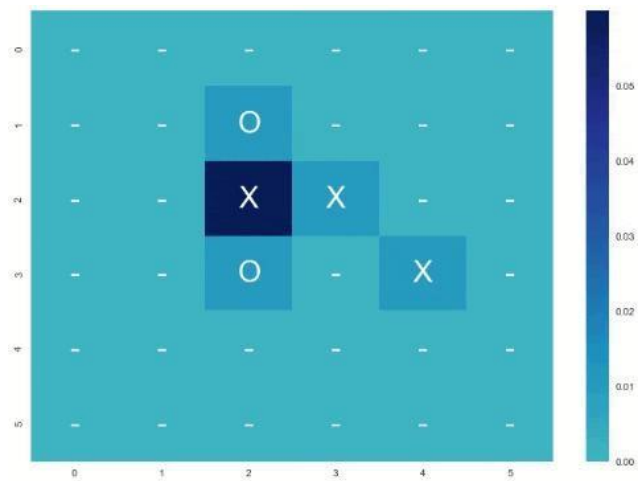


Figure 10: Example step 5




Figure 11: Example step 6




Figure 12: Example step 7




Figure 13: Example step 8




Figure 14: Example step 9




Figure 15: Example step 10




Figure 16: Example step 11




Figure 17: Example step 12




Figure 18: Example step 13