

Review

Not peer-reviewed version

Connect-4 AI: A Comprehensive Taxonomy and Critical Review of Methods and Metrics

[Mohammed Alaa Ala'anzy](#)^{*}, Akerke Madiyarova, Aidos Aigeldiyev, [Raiymbek Zhanuzak](#), [Omar Alnaseri](#)

Posted Date: 2 February 2026

doi: 10.20944/preprints202601.2416.v1

Keywords: Connect-4; explainable AI; game-playing artificial intelligence; reinforcement learning; search algorithms




Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Review

Connect-4 AI: A Comprehensive Taxonomy and Critical Review of Methods and Metrics

Mohammed Alaa Ala'anzy^{1,*} , Akerke Madiyarova¹, Aidos Aigeldiyev¹, Raiymbek Zhanuzak^{1,2} and Omar Alnaseri³

¹ Department of Computer Science, SDU University, Kaskelen, Kazakhstan

² Department of Computer Science, Hof University of Applied Sciences, Hof, Germany

³ Department of Electrical Engineering, Baden-Wuerttemberg Cooperative State University, Ravensburg, Germany

* Correspondence: m.alanzy@ieee.org

Abstract

Connect-4, a solved two-player perfect-information game, offers a compact benchmark for artificial intelligence research due to its strategic depth and structural regularities, including board symmetries. This review presents a taxonomy-driven synthesis of Connect-4 AI research, encompassing game-theoretical foundations, classical search algorithms, reinforcement learning methods, explainable AI, and formal verification approaches. Analysis of search-, learning-, and hybrid-based methods reveals three dominant patterns: (i) classical search techniques prioritize determinism and efficiency but face scalability limits; (ii) reinforcement learning and neural approaches improve adaptability at the cost of interpretability and computational resources; and (iii) explainable and formally verified frameworks enhance transparency and reliability while imposing additional performance constraints. Recent advances in Connect-4 AI are driven less by raw performance gains than by strategic integration of efficiency, adaptability, interpretability, and robustness. Structuring the literature through a multidimensional taxonomy clarifies conceptual relationships, highlights underexplored research intersections, and points to emerging trends, including hybrid search-learning systems and explainable game intelligence. Overall, Connect-4 serves as a concise experimental domain for investigating fundamental challenges in game-playing AI, system design, and human-AI interaction.

Keywords: Connect-4; explainable AI; game-playing artificial intelligence; reinforcement learning; search algorithms

1. Introduction

1.1. Connect-4 as a Benchmark Game

Connect-4 is a classic two-player, perfect-information, zero-sum game that has long served as a benchmark for evaluating artificial intelligence (AI) and search algorithms. The game is played on a 6×7 vertical grid, where players alternately drop discs into columns with the objective of forming a sequence of four pieces in a horizontal, vertical, or diagonal direction. Owing to its simple rules and non-trivial strategic depth, Connect-4 has frequently been used to study adversarial decision-making, heuristic search, and algorithmic optimisation [1].

Originally introduced in 1974 by Howard Wexler and Ned Strongin, Connect-4 quickly attracted research interest as a computationally tractable yet strategically rich game. A major milestone was achieved in 1988, when Allis presented a formal, knowledge-based solution to the game, proving that the first player can always force a win under perfect play [2]. Beyond optimal strategies, Connect-4 also became a subject in AI teaching and robotics research, as early studies used it to illustrate fundamental AI concepts and game-playing challenges [3]. These foundational works laid the groundwork for subsequent developments in game search and learning algorithms.

Beyond the specific case of Connect-4, Allis later generalised these ideas in his broader study of search methods in games and artificial intelligence, providing foundational insights into proof-number

search and systematic game-solving strategies [4]. These contributions have had a lasting impact on the design of search algorithms for deterministic, perfect-information games.

Although Connect-4 is theoretically solved, it continues to be widely used as an experimental testbed for AI research. Its moderate state space and deterministic dynamics allow researchers to analyse algorithmic behaviour in a controlled setting while still capturing meaningful strategic complexity. An example configuration of the Connect-4 game board is illustrated in Figure 1, which highlights the structural simplicity underlying the game's strategic richness.

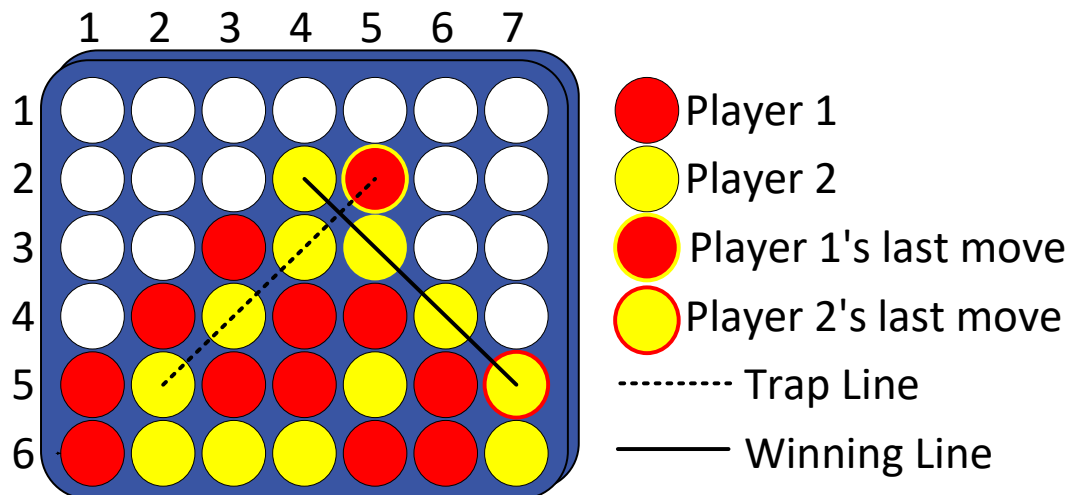


Figure 1. Example of a Connect-4 game configuration.

From a structural perspective, Connect-4 exhibits several forms of symmetry, including horizontal reflection symmetry of board states, equivalence classes of move sequences, and symmetry in player roles under optimal play. These properties reduce the effective state space and have been implicitly or explicitly exploited by classical search algorithms through pruning, transposition tables, and heuristic evaluation. In learning-based approaches, symmetry is often captured implicitly through data augmentation or invariant feature representations. Understanding how different AI paradigms leverage or ignore symmetry provides important insight into their efficiency, generalisation, and interpretability.

1.2. From Classical Search to Learning-Based Approaches

Early research on Connect-4 focused primarily on classical search-based techniques, including Minimax and its optimised variants, which exploit the deterministic and fully observable nature of the game. These approaches enabled exhaustive analysis of game trees and played a central role in establishing optimal strategies. However, their scalability is limited when extended to larger boards, rule variations, or real-time decision-making constraints.

To address these limitations, researchers increasingly adopted simulation-based and learning-based approaches. Monte Carlo Tree Search (MCTS) emerged as a powerful alternative, enabling agents to balance exploration and exploitation through stochastic sampling rather than exhaustive enumeration. Reinforcement learning (RL) further expanded this paradigm by allowing agents to learn effective policies through self-play and experience-driven optimisation. Connect-4 has been extensively used to evaluate such methods, particularly as an intermediate benchmark between small toy problems and large-scale games.

Recent work has demonstrated the effectiveness of hybrid approaches that combine MCTS with reinforcement learning and neural network function approximation. Systems inspired by AlphaZero have achieved strong performance in Connect-4 using end-to-end learning without handcrafted heuristics [5,6]. Additional studies have explored novel learning objectives, optimisation strategies, and

exploration mechanisms tailored to Connect-4 and similar adversarial games [7,8]. These developments confirm that Connect-4 remains a relevant platform for advancing modern AI techniques.

1.3. Explainability and Human–AI Interaction

As AI agents for game playing have grown more sophisticated, concerns regarding transparency and interpretability have become increasingly prominent. Many high-performing models, particularly deep neural networks, operate as black boxes, offering limited insight into the reasoning behind their decisions. In adversarial games such as Connect-4, understanding why an agent selects a specific move is essential for debugging, trust calibration, and educational applications.

Explainable artificial intelligence (XAI) has therefore become an important research direction in game AI. Recent studies have applied XAI techniques to Connect-4 agents in order to explain move selection, strategy formation, and policy evolution, thereby improving transparency and interpretability [9]. Such explanations are critical for both researchers and end users, particularly when AI systems are used as decision-support tools rather than autonomous players.

In parallel, research has examined human–AI interaction in Connect-4, focusing on how users interpret and rely on AI recommendations of varying skill levels. Empirical studies have shown that effective collaboration depends not only on agent performance but also on users' ability to calibrate trust and confidence in AI advice [10]. These findings highlight that the value of game-playing AI lies not only in optimal performance but also in its ability to support meaningful human understanding.

1.4. Emerging Trends and Motivation for This Review

Beyond classical and learning-based methods, recent work has begun to explore the integration of Large Language Models (LLMs) with traditional game AI techniques. Notably, LLM-guided Monte Carlo Tree Search has been shown to achieve strong performance across multiple board games, including Connect-4, by incorporating natural language reasoning into the search process [11]. This emerging direction suggests a broader shift toward hybrid systems that combine symbolic reasoning, learned representations, and search-based planning.

Despite the extensive body of research on Connect-4, existing surveys typically focus on general board games or broader classes of game-playing agents. As a result, the diverse range of methods applied specifically to Connect-4, spanning game theory, search algorithms, reinforcement learning, explainability, and computational optimisation, has not yet been systematically synthesised.

This study addresses this gap by presenting, to the best of the authors' knowledge, the first comprehensive survey and taxonomy dedicated exclusively to Connect-4 AI research. The proposed taxonomy integrates multiple dimensions, including game-theoretical foundations, algorithmic and learning-based approaches, strategic and tactical reasoning, explainability, and computational enhancements. By structuring the literature in this way, the survey provides both a historical perspective and a roadmap for future research.

The main contributions of this work are summarised as follows:

- A structured and focused survey of artificial intelligence research applied to the Connect-4 game, addressing the lack of a consolidated Connect-4-specific review in the existing literature.
- A multidimensional taxonomy that organises existing Connect-4 research into clearly defined categories, including game-theoretical foundations, algorithmic and learning-based approaches, strategic and tactical reasoning, explainability, and computational and formal techniques.
- A taxonomy-driven synthesis of the literature, highlighting how different classes of methods are positioned within the proposed framework and how they address distinct aspects of Connect-4 gameplay.
- An analysis of open challenges and future research directions, derived from gaps identified across the taxonomy, with particular attention to explainability, adaptability, and evaluation practices.

The remainder of this paper is organised as follows. Section 2 describes the methodology used to identify and select relevant studies. Section 3 presents the proposed taxonomy and categorises existing

research accordingly. Section 4 discusses key findings and cross-cutting trends. Section 5 outlines open challenges and promising future directions. Finally, Section 6 concludes the paper.

2. Review Protocol

This section describes the systematic procedure adopted to identify, screen, and analyse relevant literature on Connect-4. The review protocol follows established guidelines for structured literature reviews, ensuring transparency, reproducibility, and comprehensive coverage of the research landscape [12].

2.1. Data Sources

The literature search was conducted across multiple reputable digital libraries to ensure broad coverage of high-quality research. Specifically, IEEE Xplore, Scopus, Web of Science, and SpringerLink were queried. Only peer-reviewed journal articles and conference proceedings were considered, as these sources represent validated and influential contributions to the field [12,13].

2.2. Search Strategy

The literature search employed a comprehensive query incorporating alternative names and broader game descriptors commonly used in the literature:

("Connect-4" OR "Connect Four" OR "four-in-a-row" OR "four in a row") OR ("gravity-based game" OR "connection game" OR "alignment game") AND ("general game playing" OR "board game" OR "perfect-information game")

This search retrieved 293 records. After limiting to English-language publications, 284 records remained. Filtering by publication type yielded 248 peer-reviewed studies, comprising 147 journal articles and 101 conference proceedings. The temporal scope of the review was limited to publications appearing between 2015 and 2026, thereby concentrating the analysis on contemporary advances in artificial intelligence and game-playing research. This procedure yielded a final corpus comprising 153 documents.

2.3. Study Selection and Eligibility Criteria

Records were screened in a multi-stage process:

- Title screening: 50 records were excluded as clearly irrelevant, leaving 103 records.
- Abstract screening: 37 records were excluded, leaving 67 records.
- Full-text screening: 17 records were excluded as closely related but outside the review scope, resulting in 49 studies included in the final review.

Eligible studies explicitly involved Connect-4 or structurally equivalent games as primary or secondary evaluation platforms and employed computational, artificial intelligence, or learning-based methods. Publications not written in English, non-peer-reviewed materials (e.g., theses, dissertations, book chapters), and duplicate records were excluded.

2.4. Data Extraction and Synthesis

For each selected study, relevant metadata, including title, authors, publication year, venue, and research focus, were extracted. Methodological aspects such as algorithmic approach, learning paradigm, evaluation strategy, and reported outcomes were also analysed.

Due to the heterogeneous nature of experimental setups, datasets, and computational environments across Connect-4 studies, this review adopts a qualitative, taxonomy-driven synthesis rather than a uniform experimental comparison, as direct experimental normalisation would be methodologically inappropriate.

The extracted data were synthesised qualitatively to identify prevailing research trends, methodological advances, and underexplored directions in Connect-4 research. This structured synthesis

forms the basis for the taxonomy presented in Section 3 and supports the discussion of emerging themes and open research challenges in subsequent sections.

3. Taxonomy and Research Directions

In this section, we categorise existing research on Connect-4 into clearly defined areas, offering a structured overview of major contributions to game-solving strategies, AI methodologies, and computational enhancements. This classification allows us to trace how different lines of inquiry, including theoretical models, algorithm design, practical implementation, and optimisation, have shaped the development of intelligent agents for Connect-4. By analysing foundational theories, algorithmic innovations, and formal verification methods, we highlight how diverse approaches contribute to robust, efficient, and explainable AI decision making. The overall taxonomy of research areas and their corresponding sub-themes is illustrated in Figure 2, providing a visual summary of how these contributions interconnect within the broader research landscape. Some AI methods appear across multiple dimensions of the taxonomy. For example, MCTS is discussed in both game-theoretical analysis and algorithmic strategy due to its relevance to strategic reasoning and practical implementation. These overlaps are included deliberately to reflect the diverse roles such methods play. Each sub-theme highlights a specific focus, such as theoretical foundations or applied performance, and we now clarify these distinctions within the relevant sections.

This taxonomy introduces a multidimensional structure that organises the literature across five complementary dimensions: algorithmic approaches, application contexts, explainability, and computational techniques and strategies. Unlike prior reviews that focus on a single axis, such as learning methods or search strategies, this framework provides an integrated view that connects technical foundations with human-AI interaction and system-level considerations. This structure is intended to guide future research by revealing underexplored combinations and clarifying the roles of each method in strategic gameplay.

3.1. Game Theoretical Foundations

Game theory provides the foundation for understanding optimal strategies, computational complexity, and zero-sum dynamics in Connect-4. These principles shape AI decision-making by linking theoretical reasoning with algorithmic implementations such as Minimax and Alpha-Beta Pruning, while also supporting explainability by offering principled insights into agent behaviour. As shown in Figure 3, this sub-theme is organised into three core areas: Optimal Strategies, Complexity and NP-completeness, and Game Theory in Two-Player Zero-Sum Games.

In Connect-4, game-theoretical properties such as determinism, perfect information, and zero-sum structure enable substantial reduction of the state space through search-based algorithms. Classical minimax exploits these properties by pruning branches that dominate, while Alpha-Beta Pruning further reduces the effective branching factor by eliminating moves that cannot influence the outcome. The proof-number search leverages win-loss decomposition to focus exploration on critical subtrees, significantly reducing the number of evaluated positions. These techniques allow Connect-4 agents to achieve strong performance with limited computational resources, and they form the foundation upon which later enhancements, such as heuristic evaluation, parallelisation, and hybrid learning-based methods, are built.

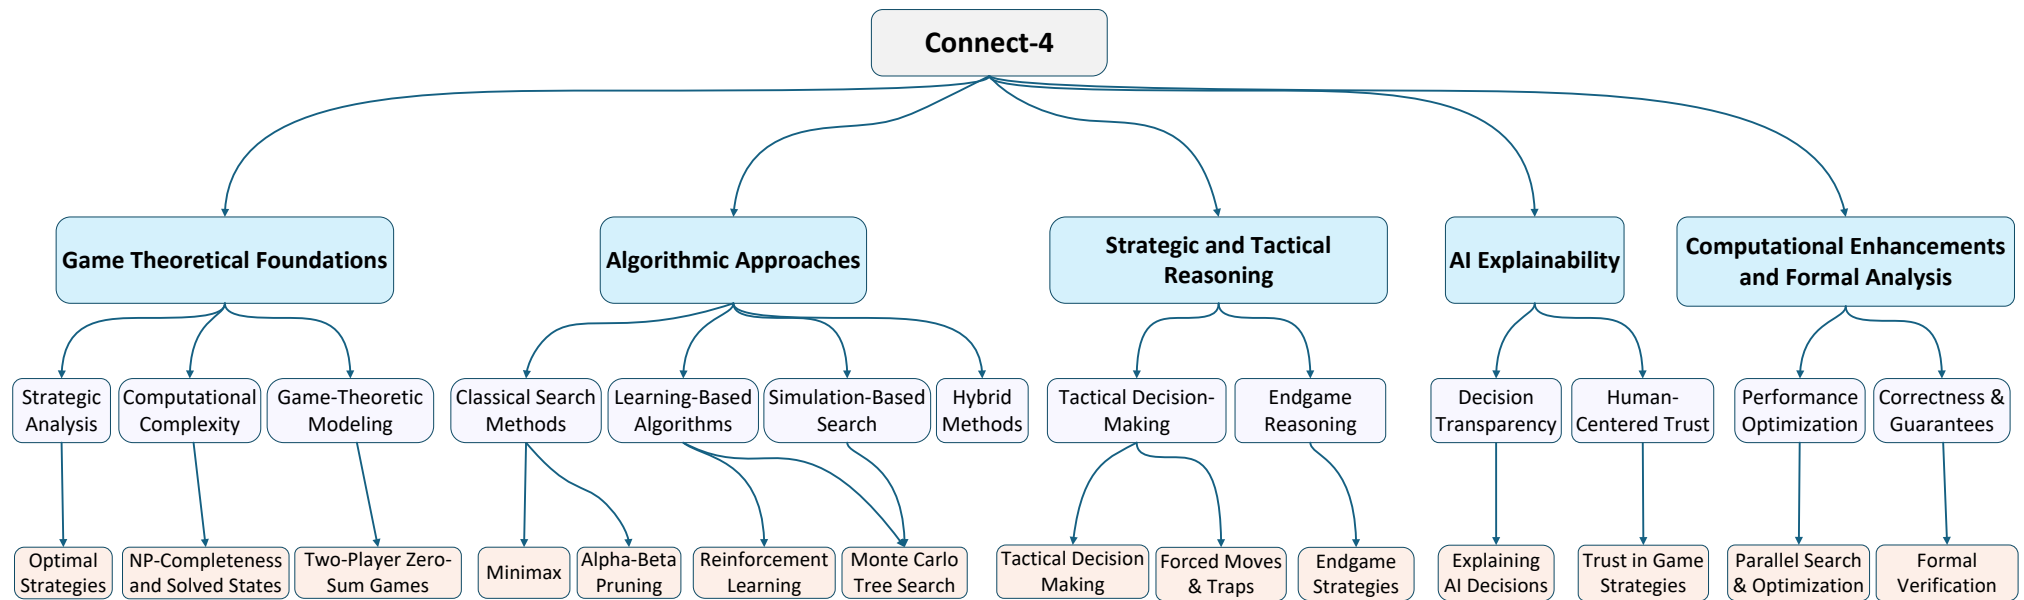


Figure 2. Taxonomy of Connect-4 research illustrating the relationships between game-theoretical foundations, algorithmic approaches, strategic reasoning, explainability, and computational enhancements.

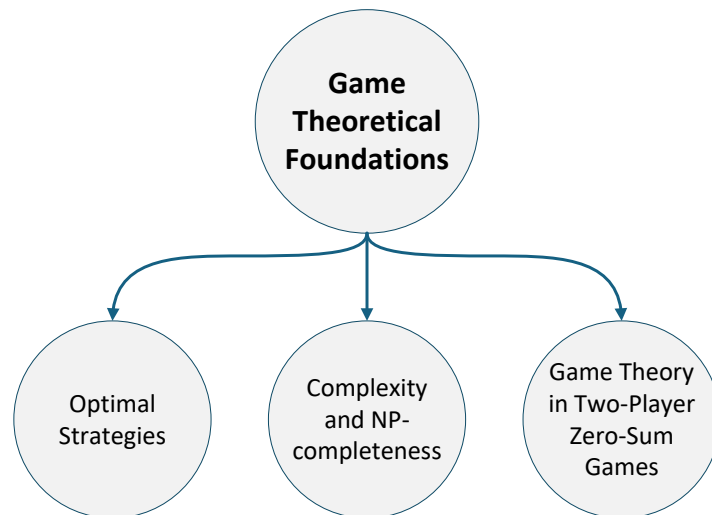


Figure 3. Game Theoretical Foundations sub-theme.

3.1.1. Optimal Strategies

The work in [14] compares MCTS and Minimax (enhanced with Alpha-Beta Pruning and dynamic programming) to evaluate their effectiveness in Connect-4. Optimising Minimax reduces redundant calculations and improves game tree depth, resulting in a 79% win rate against MCTS at a search depth of five moves. While MCTS excels in exploration-heavy environments, Minimax demonstrates superior performance in structured, low-complexity scenarios, though its scalability to larger games remains limited.

A comparative study in [15] introduces an Evolutionary Tournament framework to benchmark Minimax, MCTS, and Reinforcement Learning (RL) approaches. Results show that MCTS achieves the highest win rate, Minimax ranks second, and Q-Learning performs fastest but with the lowest success rate. The study highlights trade-offs between computational speed and strategic depth, while noting that hybrid methods (e.g., MCTS combined with RL) remain underexplored. It also suggests that Deep Q-Networks (DQN) could improve upon the limitations of classical Q-Learning.

In [16], the authors optimise n-tuple network structures using genetic algorithms, achieving a 75% win rate in Othello through dynamic adaptation of network shapes to different game phases. Although this work focuses on Othello, it is included because n-tuple networks and genetic optimization are game-agnostic techniques applicable to grid-based, deterministic, perfect-information games. Given the structural similarities between Othello and Connect-4, such as fixed board representations and discrete action spaces, it suggests that similar adaptive approaches could enhance Connect-4 agents by improving feature representation and evaluation functions.

3.1.2. NP-Completeness and Solved States

The study in [17] examines Shift-Tac-Toe, a Connect-4 variant that introduces row-shifting mechanics. The authors demonstrate that determining winning strategies in this game is NP-complete, using a reduction from the Directed Hamiltonian Path problem. While this variant increases strategic depth, it also introduces potential infinite loops and unpredictability, highlighting the computational challenges of extending traditional Connect-4 mechanics.

In [18], a theoretical model for positional games on partially ordered sets is proposed, reflecting Connect-4's gravity constraint. The authors prove that determining the winner is generally PSPACE-complete but solvable in polynomial time for specific board structures. Although this work is theoretical, it advances understanding of the computational limits of structured two-player games like Connect-4.

3.1.3. Game Theory in Two-Player Zero-Sum Games

The framework proposed in [19], called Adversarial Flow Networks (AFlowNets), extends generative flow models to competitive settings. It enables agents to learn equilibrium strategies via self-play while favouring shorter winning sequences. Applied to Connect-4, AFlowNets outperform AlphaZero in both learning speed and move quality, though the approach may face scalability challenges due to memory-intensive trajectory objectives.

The Probability-Based Proof Number Search (PPNS) algorithm is evaluated in [20] for Connect-4 and Othello. PPNS integrates information from both explored and unexplored nodes, reducing the number of nodes explored by up to 57% compared to traditional Proof Number Search (PNS). While effective for long-term planning, PPNS is prone to floating-point precision issues, which the authors address using precision-rate adjustments. Although part of the evaluation is conducted on Othello, the method is directly transferable to Connect-4, as PPNS is a domain-independent search algorithm designed for deterministic, turn-based, perfect-information games. The inclusion of Othello serves to demonstrate the generality and robustness of the approach rather than reliance on game-specific mechanics.

These studies collectively demonstrate how game-theoretical frameworks inform both strategy design and computational feasibility in Connect-4. Research comparing Minimax, MCTS, and Q-Learning highlights trade-offs between deterministic depth, exploration, and execution speed, while complexity studies clarify the computational challenges inherent to solving or generalising Connect-4 variants. Formal frameworks such as AFlowNets connect equilibrium reasoning with efficient learning, underscoring the role of game theory in advancing competitive AI agents.

3.2. Algorithmic Approaches

AI agents for Connect-4 have been developed using diverse algorithmic approaches, including classical methods such as Minimax and Alpha-Beta Pruning, probabilistic techniques like Monte Carlo Tree Search (MCTS), learning-based strategies grounded in RL, and hybrid models combining multiple paradigms. Each category presents trade-offs between strategic depth, computational efficiency, and adaptability. As shown in Figure 4, these algorithms not only build on theoretical models of optimal play but also contribute to tactical reasoning, computational scalability, and explainability.

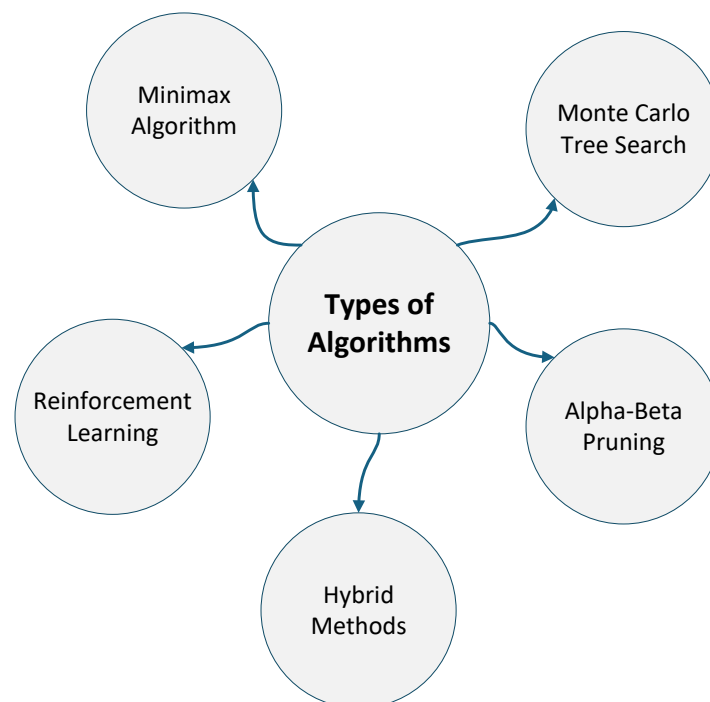


Figure 4. Types of Algorithms sub-theme.

3.2.1. Minimax Algorithm

In [21], Minimax with Alpha-Beta Pruning was evaluated across varying board sizes (7×6 to 13×12) against a greedy agent. The results confirm Minimax's strong performance and scalability for moderate grid sizes, though efficiency declines as the state space expands. Alpha-beta Pruning significantly reduces computational overhead, allowing deeper searches. However, Minimax becomes less effective in large or complex environments, motivating further enhancements through RL techniques and advanced heuristics.

3.2.2. Alpha-Beta Pruning

Alpha-Beta Pruning remains a core enhancement to Minimax. The study in [22] shows its effectiveness for shallow searches but also highlights the superiority of MTD(f) for deeper analyses due to reduced node evaluations. Additionally, [23] demonstrates that tuning exploration noise in AlphaZero's MCTS component can significantly improve win rates, underscoring the importance of careful parameter selection.

Overall, these algorithmic approaches reveal trade-offs between deterministic precision, probabilistic flexibility, and learning-based adaptability. Minimax and Alpha-Beta Pruning excel in structured environments but lack scalability, while MCTS offers robust adaptability, especially when combined with deep learning or hybrid methods. RL introduces generalisation but is sensitive to design parameters, and hybrid models require careful tuning. Selecting the right approach depends on balancing tactical depth, computational resources, and the desired level of interpretability.

3.2.3. Reinforcement Learning

RL methods have shown adaptability through self-play but often require extensive training data and are sensitive to hyperparameters. In [24], temporal difference learning (TDL) with n-tuple networks achieved strong performance, defeating a perfect Minimax agent in favourable scenarios. The approach automatically learns useful board patterns but suffers from slow convergence.

The study in [25] introduces a two-tier multi-agent RL model that combines match-level strategic learning with in-game tactical optimisation, outperforming single-focus agents. Adaptive learning rate strategies, as shown in [26], further accelerate training and improve decision-making by dynamically tuning learning rates.

Addressing the data efficiency challenges in RL, imitation learning approaches offer an alternative by learning from expert demonstrations rather than extensive self-play. [27] introduces the Strategy Representation for Imitation Learning (STRIL) framework for multi-agent games, which tackles the problem of learning from diverse player trajectories that may contain sub-optimal strategies. The framework combines strategy representation learning with data filtering mechanisms to identify dominant trajectories, successfully demonstrating improved performance in competitive environments including Connect-4. As a plugin method, STRIL can be integrated into existing imitation learning algorithms, offering a practical solution to strategy representation challenges in multi-agent scenarios.

Improving the training efficiency of AlphaZero has been addressed through enhanced value target formulations. [28] presents a comprehensive family of AlphaZero value targets, introducing AlphaZero with greedy backups (A0GB) that uses more effective training targets for neural networks. The research establishes a three-dimensional space describing training targets that encompasses the original AlphaZero approach as well as soft-Z and A0C variants from existing literature. Notably, A0GB successfully finds optimal policies in small tabular domains where the original AlphaZero target fails, while also achieving superior performance and faster training than standard AlphaZero on Connect-4 and Breakthrough. The work provides valuable insights into the relationship between tabular learning and neural network-based value function approximation.

3.2.4. Monte Carlo Tree Search

Recent advancements in MCTS, reviewed in [29], highlight improvements to exploration-exploitation balance, computational efficiency, and hybridisation with RL and deep learning. Domain-

specific adaptations and parallelisation have enhanced strategic gameplay, though challenges remain for high branching factors and real-time constraints.

The author in [7] enhances MCTS by incorporating novelty-driven exploration into the selection phase. Instead of relying solely on value estimates, the algorithm combines them with novelty scores to encourage exploration of less-visited or uncertain states. Evaluated on Connect-4 and several other board games, the approach consistently improves MCTS performance across both heuristic-based and neural-guided settings, demonstrating that novelty estimation can help MCTS avoid local optima and achieve more robust search behaviour.

Dual MCTS, proposed in [30], improves AlphaZero by using two search trees and a single deep network, reducing training time by over 32% on Connect-4 compared to AlphaZero. Similarly, [31] introduces dynamic Minimax depth adjustments within MCTS, improving tactical accuracy and decision speed while reducing unnecessary computation.

Parameter randomisation techniques, examined in [32], show that randomising search-control settings per simulation increases adaptability and performance in games with many winning paths, though less effective in narrow-win scenarios.

3.2.5. Hybrid Methods

Hybrid models integrate the precision of Minimax with the probabilistic strengths of MCTS. In [33], three hybrid strategies (MCTS-MR, MCTS-MS, MCTS-MB) are tested, with MCTS-MS proving most effective in tactical scenarios like Connect-4. Similarly, [34] demonstrates that combining Minimax during various MCTS phases improves performance in adversarial settings but increases computational costs.

The author in [5] proposes a lightweight alternative to AlphaZero that significantly reduces computational requirements. Instead of using MCTS during training, the authors train agents using temporal-difference learning with n-tuple networks and apply MCTS only at test time to enhance decision making. This design preserves much of the strategic AlphaZero strength while avoiding expensive self-play training. The approach is evaluated in Othello, Connect-4, and Rubik's Cube, demonstrating that strong performance can be achieved on standard hardware.

Advances such as evolutionary self-learning and tree search optimisation, introduced in [35], reduce AlphaZero's computational overhead without sacrificing strength. Transfer learning approaches [36,37] enable cross-domain knowledge reuse, reducing training time, though negative transfer risks remain.

Hybrid training strategies for RL, explored in [38], balance exploration and exploitation, significantly reducing the training required for perfect play.

Hybrid approaches that combine multiple algorithmic strategies have shown promise in enhancing traditional game-playing methods. [39] introduces Go-Exploit, a novel search control strategy for AlphaZero that addresses limitations in policy improvement. Unlike standard AlphaZero that begins self-play from initial game states, Go-Exploit samples start states from an archive of states of interest, enabling more effective exploration of the game tree and better value function generalization. The approach demonstrates superior sample efficiency compared to standard AlphaZero in Connect-4 and 9×9 Go, showing stronger performance against reference opponents and in head-to-head play. When compared to KataGo, Go-Exploit's search control strategy proves more effective, with further improvements when incorporating KataGo's other innovations.

Evolutionary approaches combined with reinforcement learning offer promising solutions to training stability issues. [40] addresses the sample inefficiency and unstable learning dynamics inherent in self-play reinforcement learning by introducing coevolutionary principles. The approach leverages competitive pressures from evolutionary computation to optimize the training of a population of Rainbow DQN agents in Connect-4. This coevolutionary framework compels agents to continuously innovate and adapt, resulting in a 15% improvement in win percentage over leading self-play algorithms. Additionally, the method demonstrates enhanced training stability with reduced variation in evaluation performance, addressing key challenges in self-play learning environments.

3.3. Strategy and Tactical Play

This subsection examines how strategic and tactical reasoning enhance AI performance in Connect-4. While high-level algorithms establish an agent's decision-making framework, situational awareness and precise tactical execution often determine game outcomes, especially during dynamic or endgame phases. As illustrated in Figure 5, this sub-theme is divided into three areas: Tactical Decision Making, Endgame Strategies, and Forced Moves and Traps. These components emphasise the interplay between local decision accuracy, long-term planning, and opponent modelling, highlighting the need to convert abstract strategies into effective actions while responding to evolving board states.

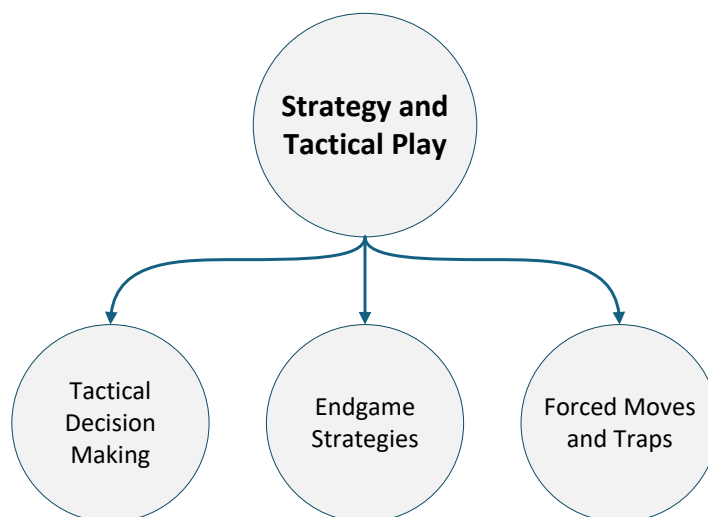


Figure 5. Strategy and Tactical Play sub-theme.

3.3.1. Tactical Decision Making

The study in [41] explores how AI agents can learn evaluation functions by imitating expert game-play trajectories. Using an n-tuple feature-based linear model with preference learning, the authors employ the Veleno game engine to generate high-quality training data. To improve generalisation, the Dagger algorithm refines decision-making based on prior states, while random play diversifies trajectories. Results show that n-tuple models trained on expert data perform well but depend heavily on the variety of observed game states. The study recommends opponent modelling and adaptive learning to further enhance performance.

In [42], the authors propose generating alternative starting positions to adjust game difficulty and discover new variants. A symbolic search with iterative simulation explores large state spaces, creating balanced start states based on parameters like game length, player skill, and win conditions. Results confirm improved engagement and strategic diversity, though some generated positions strongly favour one player, requiring refinement to ensure fairness.

The study in [43] analyses Alpha-Beta Pruning as part of a comparative evaluation of AI algorithms in adversarial settings, using Connect-4 as a baseline. Alpha-Beta significantly reduces the number of explored nodes and improves decision-making efficiency. However, its static nature and computational cost limit scalability in highly dynamic or complex scenarios. Despite these challenges, Alpha-Beta remains valuable where state spaces are predictable and manageable.

3.3.2. Forced Moves and Traps

The authors in [44] examine the role of stochasticity in search-based strategies, particularly in MCTS. Controlled randomness in rollout policies improves exploration and reduces redundancy in node expansions, enhancing computational efficiency. Entropy-based randomness modulation in GPU-accelerated simulations is shown to further optimise resource allocation. However, the study remains largely theoretical, leaving practical parallel implementations for future work.

In [45], a Brain-Computer Interface (BCI) version of Connect-4 is introduced, using code-modulated visual evoked potentials (c-VEPs) for move selection. The system achieves a 94.10% accuracy rate, enabling real-time, gaze-controlled interactions that enhance endgame precision. While effective in controlled settings, challenges such as calibration complexity and performance degradation in multiplayer tasks (due to fatigue or competitive pressure) limit its broader applicability.

These studies underscore the importance of integrating tactical awareness with long-term strategy in Connect-4 AI. Techniques like n-tuple evaluation, adaptive endgame search, and idle-time analysis improve decision accuracy, while randomness-based trap detection enhances robustness in real-time play. Nevertheless, many approaches assume static conditions or idealised environments, highlighting the need for adaptive models that balance reactivity, prediction, and computational efficiency in dynamic gameplay. Table 1 summarises the major techniques used in tactical and strategic adaptation across Connect-4 AI agents.

Table 1. Strategy and Tactical Play techniques in Connect-4.

Ref.	Sub-area	Technique / Focus	Claimed benefit	Key limitation(s)
[41]	TD	Preference learning with n-tuple features (imitation of expert trajectories)	Learns after-state evaluation functions aligned with expert play; improved generalisation via DAgger	Depends on distribution of training trajectories; limited opponent modelling
[42]	TD	Automatic generation of balanced or skill-adjusted start states	Symbolic search + simulation to create positions that calibrate difficulty and increase variety	Some generated starts favour one player; fairness calibration needed
[43]	TD	Alpha-Beta Pruning for adversarial decision-making (cybersecurity testbed)	Prunes large parts of the game tree to improve minimax efficiency	Limited adaptability in dynamic settings; computationally heavy for large state spaces
[46]	EG	MCTS & Alpha-Beta for late-game search	Focuses on critical late stages; structured optimisation of final plies	Deep searches costly in real time; poor transfer to 3D variants
[47]	EG	Dynamic time allocation for MCTS (STOP, BEHIND, UNST, CLOSE)	Allocates time where most needed; improves decision efficiency	Requires careful tuning; misallocation can waste time or miss key moves
[48]	EG	Idle-time analysis (FDS, ADS, OADS)	Uses idle cycles to precompute responses; improves win rates and reaction speed	Assumes deterministic settings; effectiveness in stochastic games untested
[44]	FM&T	Stochasticity measurement and control in search (parallel MCTS)	Moderate randomness reduces redundant exploration and speeds convergence	Mostly theoretical; lacks large-scale GPU implementations
[45]	FM&T	c-VEP BCI-controlled Connect-4	Enables non-traditional interaction with high move-selection accuracy	Accuracy drops in multiplayer; calibration overhead and fatigue issues

TD = Tactical Decision; EG = Endgame; FM&T = Forced Moves & Traps.

3.3.3. Endgame Strategies

The research in [46] investigates endgame optimisation, highlighting the transition from general strategies to precise, calculated moves as the board fills. Both MCTS and Alpha-Beta Pruning are shown to improve decision-making during these critical stages, though Alpha-Beta struggles with unpredictability, and MCTS incurs high computational costs for deep searches. These limitations are further amplified in complex variants like 3D Connect-4.

Time management strategies for MCTS, studied in [47], demonstrate how dynamic allocation (e.g., STOP, BEHIND, UNST, CLOSE) improves performance by prioritising critical moves. The STOP strategy, in particular, reduces unnecessary search without compromising decision quality, though improper parameter tuning can lead to inefficient resource use.

Idle-time analysis strategies, introduced in [48], leverage unused computational time to precompute responses and refine decisions. Techniques like Frequent Decision Strategy (FDS), Analysed Decision Strategy (ADS), and Online Analysis Decision Strategy (OADS) increase win rates and reduce reaction times by up to 69%. However, these methods are best suited to deterministic environments, and their efficacy in probabilistic or complex domains remains uncertain.

3.4. AI Explainability

This subsection explores how AI agents in Connect-4 communicate their decision-making processes to users, focusing on the broader theme of XAI. Although these studies do not focus on improving game-solving performance, they provide critical insights into how Connect-4 AI systems are interpreted, trusted, and integrated into human-centred decision-making settings.

As AI systems become increasingly capable, understanding how and why they make certain decisions becomes critical, especially in strategic games where trust, transparency, and interpretability play a central role. As illustrated in Figure 6, this sub-theme is structured into two key areas: Explaining AI Decisions and Trust in AI for Game Strategies. These categories examine how XAI techniques such as saliency attribution, recursive reasoning, and brain-computer interfaces influence not only gameplay performance but also user trust, collaboration, and psychological engagement. The ability to interpret agent behaviour often relies on principled models of optimality, reflects the internal dynamics of the decision algorithms, and helps illuminate subtle tactical choices made during critical moments of play.

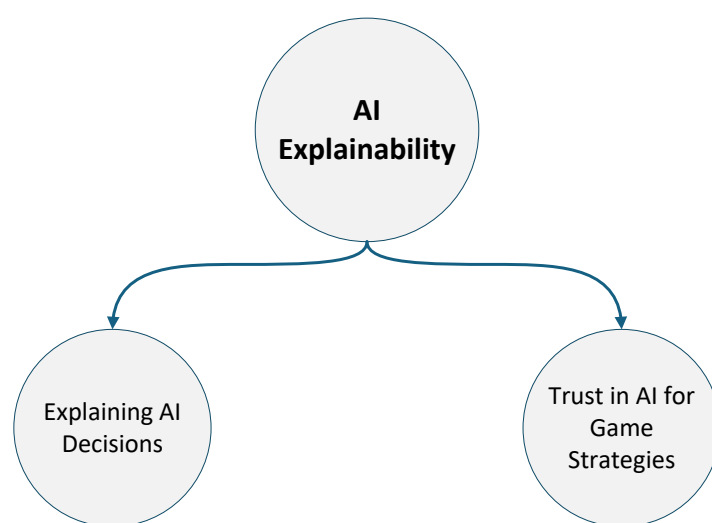


Figure 6. AI Explainability sub-theme.

3.4.1. Explaining AI Decisions

The study in [49] explores how RL agents trained with partial information can enhance XAI evaluation. The authors introduce a setup where agents play Connect-4 while being trained with missing colour information, allowing them to develop saliency-based decision-making strategies. The research leverages characteristic functions from cooperative game theory, training AI to assess the importance of specific board positions. By hiding certain game features during training, the study eliminates off-manifold counterfactuals, which are common limitations in standard XAI methods. The trained agents then compete in a round-robin tournament, where different saliency attribution techniques (such as Shapley values and heuristic-based methods) are tested to determine their effectiveness. The results show that agents trained with partial information can still play optimally while providing more interpretable decision-making processes. Among the tested XAI methods, DeepShap and Guided Backpropagation outperform others in selecting critical board features. However, the study notes that Shapley sampling becomes unreliable when a high percentage of information is missing, suggesting that RL-based XAI could benefit from alternative evaluation techniques in complex environments.

The authors in [50] explore how LLMs handle multi-turn strategic reasoning in games like Connect-4. The authors propose a Recursively Thinking-Ahead (ReTA) agent, designed to enhance long-term decision making by simulating the future moves of the opponent and integrating reward signals into the reasoning process. The research evaluates LLMs in two scenarios: Online Racing, where models directly compete against each other, and Offline Probing, where targeted game-based questions assess their strategic abilities. The study reveals that existing LLMs struggle with multi-turn strategic thinking, often failing to recognise immediate winning moves or optimal board positioning

in Connect-4. By implementing ReTA's recursive prompting mechanism, the study demonstrates that LLMs significantly improve their ability to anticipate future actions, leading to better strategic gameplay. Experimental results show that ReTA outperforms standard reasoning methods such as Chain-of-Thought (CoT), Tree-of-Thought (ToT), and ReAct in Connect-4. However, the study notes that LLMs still lag behind classical game-solving techniques like Minimax, highlighting the need for further advancements in AI-driven strategic reasoning.

In [51] the authors explore how revealing the presence of a human operator behind a robot affects the interaction between the robot and human participants during a Connect-4 game. The research delves into the psychological effects of knowing whether a robot is controlled by a human or is fully autonomous, and how this knowledge influences the human's decision-making. The study finds that when participants were aware that a human was controlling the robot, they were more likely to follow the robot's suggestions than when they believed the robot was autonomous. This behaviour suggests that trust and transparency in decision making play significant roles in how a human engages with the robot, which is crucial in strategic play like Connect-4. Forced moves and traps were also analysed through the interaction, where the robot's suggestions could lead to sub-optimal or forced plays, potentially affecting the player's perception of the robot's intelligence and effectiveness. A notable limitation of this study is that the real-world applicability of these results could vary, as the scenario was heavily controlled, and the perceptions of autonomy of participants may differ outside this experimental context. Furthermore, the limitations of the Wizard of Oz technique in revealing the operator early could lead to biases in gameplay, where participants might distrust the robot's suggestions if they feel manipulated or misled.

3.4.2. Trust in AI for Game Strategies

The study in [52] investigates the impact of different XAI strategies in human-robot interactions, using Connect-4 as a decision-making task. The authors compare classical counterfactual explanations with shared experience-based counterfactuals, which are generated from previous human-robot games. The study evaluates how these explanations affect team performance, robot persuasiveness, and user perception. Results indicate that while both explanation strategies lead to similar team performance, shared experience-based explanations make the robot's suggestions more persuasive, especially for low-performing players. The findings highlight potential risks, as less experienced users tend to follow AI recommendations without critically assessing them. The study concludes that personalised, experience-based XAI can enhance trust and collaboration in human-robot teams. However, the authors caution against over-reliance on AI explanations, suggesting the need for further research into balancing AI guidance with human autonomy in decision-making tasks.

In [53], the authors examine how personality traits influence human-robot collaboration in a decision-making task using Connect-4. The researchers compare interactions with an explainable and non-explainable AI robot, focusing on how users accept AI suggestions and adapt their strategies over time. The study finds that agreeableness and negative agency significantly impact the likelihood of following the recommendations of a robot, especially when provided with example-based counterfactual explanations. Participants with higher agreeableness were more likely to trust and follow the robot's advice, while those with higher negative agency showed reduced autonomy in decision making. Over multiple interactions, users increasingly aligned their playing style with the robot, suggesting a learning effect. The findings indicate that explainability enhances human-AI collaboration, but also raises concerns about over-reliance on AI in strategic decision making. The study suggests that balancing AI guidance with user autonomy is crucial, particularly in long-term human-robot interactions.

The authors in [54] investigate the role of trust in AI within the context of game strategies, particularly in Connect-4, where AI makes strategic decisions based on the analysis of game states. The study focuses on how transparency in AI's decision-making process influences user trust, especially when players rely on AI for guidance during gameplay. The research shows that players are more likely to accept AI-generated strategies and feel comfortable with AI decisions when the system can explain

its reasoning in an understandable way. This transparency allows users to see why certain moves are suggested, which in turn builds confidence in the system and enhances its perceived reliability. The authors find that AI that communicates its reasoning effectively, such as explaining why a certain move is optimal based on the game state, significantly improves the player's trust in its strategic suggestions. However, the study also addresses some limitations, particularly that while explaining AI's decisions can improve trust, over-explaining can lead to information overload, making the decision-making process more complex for players. Additionally, the research highlights that trust in AI can be highly dependent on the player's prior experiences and expectations of AI behaviour, suggesting that tailoring explanations to the individual user might be necessary to maintain effective interaction.

The study in [55] explores how combining Brain-Computer Interface (BCI) technology with multiplayer video games can affect player trust in AI strategies. The study investigates the application of c-VEPs (code-modulated visual evoked potentials) for controlling game decisions and examines how the transparency of AI decision making influences trust during gameplay. The authors show that when players can see how AI makes decisions based on real-time brain signals, their level of trust increases, as they are able to better understand the reasoning behind the moves of AI. This enhanced transparency makes AI-driven gameplay more acceptable and reliable, particularly in competitive environments where decision-making accuracy is critical. Findings of the study suggest that trust in AI improves when players are able to see the connection between their brain signals and strategic decisions of the game, fostering a deeper understanding of capabilities of the AI. However, the study also highlights some limitations. For instance, while transparency increases trust, it does not always enhance performance in highly dynamic game scenarios. The calibration process for BCI systems is also time-consuming, and player fatigue can affect performance, reducing the overall effectiveness of the system in real-world settings. Additionally, human trust in AI is not always aligned with optimal game strategies, as players may prefer more intuitive decisions over transparent but suboptimal AI moves.

Another study in [56] explores how robot rejection after playing Connect-4 affects trust in AI and self-esteem. The study reveals that, despite participants knowing that the robot is not human, those who were rejected by the robot experienced a significant decrease in self-esteem, reflecting psychological effects similar to interpersonal rejection. This finding raises important questions about how rejection by a robot can reduce trust in AI's decisions, making it harder for users to accept AI-driven strategies. On the other hand, acceptance from the robot did not significantly increase self-esteem, suggesting that trust in AI is not easily gained through positive feedback alone. The paper also discusses the role of explainability in enhancing trust, when users understand the reasoning behind AI decisions, they are more likely to trust the system's actions. However, the study also has some limitations: It was conducted in a controlled environment, and participants were aware that the robot lacked real emotions, which may have affected their response. Moreover, the impact of such robot rejection on broader societal attitudes toward AI was not fully explored, leaving questions about how this dynamic might influence the integration of AI in more complex, real-world applications.

The study in [57] explores the application of data mining techniques in adversarial search for Connect-4, focusing on player movement prediction. The authors simulate multiple Connect-4 games with players of varying skill levels, including human and AI-controlled agents. Three different data mining algorithms, such as Nearest Neighbour (k-NN), Generalised Linear Model (GLM), and Deep Learning, are applied to classify moves and predict which player is likely to make a given move. The research demonstrates that GLM achieves the highest accuracy (92.25%), significantly outperforming Deep Learning (80.50%) and k-NN (54.41%). The findings suggest that pattern-based classification techniques can be highly effective for predicting game moves, with structured players being more predictable than human players. Additionally, the study highlights that introducing randomisation into AI strategies makes prediction significantly harder, reinforcing the role of non-deterministic decision-making in competitive gameplay. A key strength of this research is its comprehensive evaluation of different machine learning approaches for pattern recognition in Connect-4, providing valuable

insights into how AI can anticipate player behaviour. However, one limitation is that the study focuses primarily on move classification, without exploring how such predictive models could be directly integrated into adaptive AI gameplay strategies.

These works show that explainability plays a key role in improving trust, strategic learning, and user engagement in Connect-4. Techniques such as saliency-based evaluation, recursive reasoning, and experience-driven explanations provide meaningful insights into AI decision making, yet their integration into real-time gameplay remains limited by computational demands and scalability issues. While experimental results highlight improvements in interpretability and collaboration, many methods rely on controlled environments or require significant calibration, which may hinder practical deployment. There is also a trade-off between model complexity and transparency, as highly interpretable systems often sacrifice performance. Addressing these challenges requires continued exploration of efficient, adaptable XAI solutions that can balance explainability and competitiveness in dynamic game environments.

Despite the promising results of these studies, integrating explainable AI into real-time Connect-4 agents presents several challenges. Many XAI techniques, such as Shapley values or recursive reasoning, introduce significant computational overhead, which can delay decision-making and reduce responsiveness in fast-paced gameplay scenarios. In practical settings, especially in embedded or resource-constrained environments, maintaining both interpretability and performance is difficult. Moreover, some methods rely on offline analysis or highly controlled setups, which limit their applicability in real-time matches. There is also an inherent tension between model complexity and transparency. Highly interpretable models often do not perform as well as more complex ones, while complex models are typically more difficult to explain. These limitations highlight the need for further research into lightweight and real-time compatible XAI techniques that preserve strategic strength while also enhancing user trust and understanding.

3.5. Computational Enhancements and Formal Analysis

This subsection focuses on advancements aimed at improving the computational efficiency and reliability of AI agents in Connect-4. As illustrated in Figure 7, this sub-theme is divided into two main areas: Parallel Search and Optimisation, and Formal Verification of AI Moves. These components explore how algorithmic performance can be enhanced through parallelisation techniques and how formal methods can be used to ensure the correctness of AI decision-making processes. Together, they highlight the importance of both speed and reliability in developing robust game-playing agents, particularly as search spaces grow and real-time constraints become more prominent. The approaches discussed here support the practical deployment of decision algorithms at scale, reinforce soundness in strategic reasoning, and contribute to the transparency and trustworthiness of agent behaviour across varied gameplay scenarios.

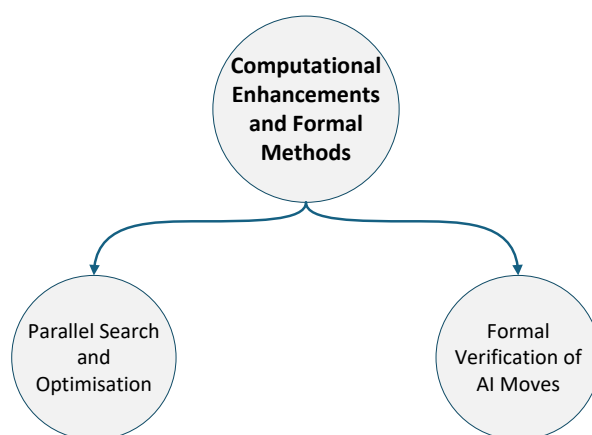


Figure 7. Computational Enhancements and Formal Methods sub-theme.

3.5.1. Parallel Search and Optimisation

The study in [58] explores how game complexity and uncertainty impact AI decision making and strategy formation in Connect-4. The authors analyse probability-based PPNS and single conspiracy number (SCN) as methods for evaluating the difficulty of a given game state. These metrics help quantify how challenging a position is for an AI agent and determine whether a more exploratory or deterministic approach is needed when searching for optimal moves. One of the key insights from this research is that pattern recognition techniques can improve AI adaptability by adjusting search heuristics based on the perceived complexity of a game position. AI agents using PPNS and SCN-based evaluations demonstrate a better ability to detect critical game states that require deeper searches while recognising less impactful positions where computational resources can be conserved. This strategic balance between computation and move importance enhances the efficiency of Connect-4 AI by preventing unnecessary exhaustive searches in trivial board positions. While this study provides valuable theoretical insights, one of its limitations is that it does not propose a specific pattern-recognition model that could be implemented in practical AI frameworks. Instead, it lays the groundwork for future research into adaptive AI agents that can dynamically adjust their search parameters based on complexity evaluation.

Understanding computational scaling relationships is crucial for optimizing AI performance in Connect-4. [59] presents an extensive analysis of performance scaling for AlphaZero agents, establishing power-law relationships between neural network parameter count and playing strength. The study demonstrates that player strength scales as a power law in network parameters when not bottlenecked by compute, and as a power of compute when training optimally sized agents. Notably, the research reveals nearly identical scaling exponents for both Connect-4 and Pentago, suggesting universal scaling principles. The findings show that larger AlphaZero models achieve superior sample efficiency, performing better than smaller models with equivalent training data.

3.5.2. Formal Verification of AI Moves

Authors in [60] explore the formal modelling and verification of a parallelised Minimax algorithm applied to Connect-4. The authors utilise mCRL2, a formal modelling language, to ensure the algorithm's correctness in concurrent execution environments. The research focuses on verifying key properties of the parallel Minimax implementation, including deadlock freedom, liveness, and correctness of state transitions. By distributing computations across multiple threads, the parallel version enhances efficiency while maintaining deterministic decision-making. Labelled Transition Systems (LTSs) are generated to analyse system behaviours, ensuring that no redundant actions occur and that all decisions align with the expected algorithmic framework. The results confirm that parallel Minimax remains consistent under formal verification, proving its reliability for multi-threaded AI in Connect-4. However, scalability challenges remain, particularly when expanding to larger game trees or real-time decision-making environments. Future improvements could focus on further optimising synchronisation mechanisms to reduce computational overhead while maintaining formal correctness.

These studies demonstrate how computational enhancements and formal verification contribute to building faster and more reliable Connect-4 AI agents. Techniques such as PPNS and SCN offer insights into how AI can prioritise search effort based on game state complexity, improving efficiency without sacrificing performance. Formal modelling of parallelised Minimax ensures that algorithmic correctness is preserved even in multi-threaded environments, supporting the safe deployment of concurrent AI systems. While both approaches offer valuable improvements, they remain limited by scalability concerns and lack integration with adaptive real-time agents. Further research is needed to combine performance-oriented enhancements with formal guarantees in a way that supports responsiveness, correctness, and resource-aware optimisation under practical constraints.

Table 2 and Table 3 summarise the literature review, including key studies on Connect-4. It lists the reference number, publication year, study title, and the main algorithms used. This summary highlights the breadth of approaches and findings across the literature, providing an overview of the advancements and trends in the field.

Table 2. Review Summary (Part 1).

Ref	Year	Title	Algorithm	ML	SOA	TA	Var	CA	C/M
[7]	2025	Novelty in Monte Carlo Tree Search	MCTS with Novelty Search	✗	✓	✓	✓	✓	✓
[5]	2022	AlphaZero-Inspired Game Learning: Faster Training by Using MCTS Only at Test Time	MCTS, TD Learning, n-tuple Networks	✓	✓	✓	✓	✓	✓
[14]	2022	Solving connect-4 using Optimized Minimax and Monte Carlo Tree Search	Minimax, Alpha-Beta Pruning, Dynamic Programming, MCTS	✗	✓	✓	✓	✓	✓
[15]	2024	An Evolutionary Framework for Connect-4 as Test-Bed for Comparison of Advanced Minimax, Q-Learning and MCTS	Advanced Minimax, Q-Learning, MCTS	✓	✓	✓	✗	✓	✓
[16]	2025	N-Tuple Network Search in Othello Using Genetic Algorithms	GA, BRKGA	✓	✓	✓	✗	✓	✗
[17]	2024	Row Shifting as a Puzzle Mechanic in Generalized Connect-4	Shift-Tac-Toe, Game Tree Search, NP-Completeness Analysis	✗	✓	✓	✓	✓	✓
[18]	2025	Poset Positional Games	-	✗	✓	✓	✓	✓	✓
[19]	2023	Expected Flow Networks in Stochastic Environments and Two-player Zero-sum Games	EFlowNet, AFlowNet	✓	✓	✓	✗	✓	✓
[20]	2020	Characterizing the Nature of Probability-Based Proof Number Search: A Case Study in the Othello and Connect-4 Games	PPNS, PNS, MCPNS	✗	✓	✓	✓	✓	✓
[21]	2023	Evaluation of the Use of Minimax Search in Connect-4	Minimax, Alpha-Beta Pruning	✗	✓	✓	✓	✓	✗
[29]	2023	Monte Carlo Tree Search- a review of recent modifications	MCTS	✓	✓	✓	✓	✓	✓
[30]	2021	Dual Monte Carlo Tree Search	AlphaZero, MPV-MCTS, Dual MCTS	✓	✗	✓	✗	✓	✓
[31]	2023	MCTS with Dynamic Depth Minimax	Minimax, MCTS	✗	✓	✓	✓	✓	✗
[32]	2021	Analysis of the Impact of Randomization of Search-Control Parameters in Monte-Carlo Tree Search	MCTS	✓	✓	✓	✓	✓	✗
[24]	2012	Reinforcement Learning with N-tuples on the Game Connect-4	TDL	✓	✗	✓	✗	✓	✓
[25]	2024	Multi-agent Dual Level Reinforcement Learning of Strategy and Tactics in Competitive Games	Multi-agent RL, Minimax	✓	✓	✓	✓	✓	✗
[26]	2014	Online Adaptable Learning Rates for the Game Connect-4	TDL, IDBD, TCL	✓	✓	✓	✓	✓	✓
[27]	2025	Learning Strategy Representation for Imitation Learning in Multi-Agent Games	STRIL Framework	✓	✗	✗	✗	✓	✓
[28]	2022	Value targets in off-policy AlphaZero: a new greedy backup	A0GB	✓	✗	✓	✓	✓	✓
[33]	2015	MCTS-Minimax Hybrids	MCTS, Alpha-Beta Pruning, Minimax	✗	✓	✓	✓	✓	✗
[34]	2013	Monte-Carlo Tree Search and Minimax Hybrids	MCTS, Minimax Hybrid	✗	✓	✓	✗	✓	✓
[35]	2021	Improvements to Increase the Efficiency of the AlphaZero Algorithm- A case study in the game connect-4	MCTS, AlphaZero	✓	✓	✓	✓	✓	✗
[36]	2023	Transfer Learning in Monte Carlo Tree Search	MCTS	✓	✗	✓	✗	✓	✗
[37]	2024	Heuristic Knowledge Transfer for General Game Playing	MCTS, AlphaZero	✓	✓	✓	✓	✓	✗

Note: SOA = Search and Optimization Algorithms, TA = Theory Analysis, Var = Variants, CA = Comparative Analysis, C/M = Code or Math.

Table 3. Review Summary (Part 2).

Ref	Year	Title	Algorithm	ML	SOA	TA	Var	CA	C/M
[38]	2022	Hybrid Training Strategies: Improving Performance of Temporal Difference Learning in Board Games	TDL, Hybrid Training	✓	✗	✓	✗	✓	✓
[39]	2023	Targeted Search Control in AlphaZero for Effective Policy Improvement	Go-Exploit	✓	✓	✗	✗	✓	✓
[40]	2020	Coevolutionary Deep Reinforcement Learning	Coevolutionary RL	✓	✓	✗	✗	✓	✓
[22]	2017	The Analysis of Alpha Beta Pruning and MTD(f) Algorithm to Determine the Best Algorithm to be Implemented at Connect Four Prototype	Alpha Beta Pruning, MTD(f)	✗	✓	✓	✗	✓	✓
[23]	2021	Hyperparameter Choice as Search Bias in AlphaZero	MCTS, AlphaZero	✓	✓	✓	✓	✓	✓
[41]	2015	On imitating Connect-4 game trajectories using an approximate n-tuple evaluation function	n-tuple evaluation, Preference Learning	✓	✓	✓	✓	✓	✗
[42]	2015	Automatic Generation of Alternative Starting Positions for Simple Traditional Board Games	Symbolic Methods, Iterative Simulation, Minimax	✓	✓	✓	✓	✓	✗
[43]	2023	Adversarial agent-learning for cybersecurity: a comparison of algorithms	A2C, AlphaZero, MCTS, ES, CEM, DRL	✓	✓	✓	✗	✓	✓
[46]	2018	Connect-4: A Novel Paradigm to Elicit Positive and Negative Insight and Search Problem Solving	✓	✗	✗	✓	✗	✓	✗
[47]	2016	Time Management for Monte Carlo Tree Search	MCTS	✗	✓	✓	✗	✓	✗
[48]	2018	Empirical Evaluation of Idle-Time Analysis Driven Improved Decision Making by Always-On Agents	AlphaZero, FDS, ADS, OADS	✓	✓	✓	✓	✓	✓
[44]	2024	Measuring Randomness in Tabletop Games	MCTS	✗	✓	✓	✗	✓	✗
[45]	2023	Assessing the Potential of Brain-Computer Interface Multiplayer Video Games using c-VEPs: A Pilot Study	c-VEP	✓	✗	✓	✗	✗	✗
[49]	2022	Training Characteristic Functions with Reinforcement Learning: XAI-methods play Connect-4	AlphaZero MCTS	✓	✓	✓	✓	✓	✗
[50]	2024	ReTA- Recursively Thinking Ahead to Improve the Strategic Reasoning of Large Language Models	Minimax, ToT, CoT, ReTA	✓	✓	✓	✓	✓	✓
[51]	2022	Questioning Wizard of Oz: Effects of Revealing the Wizard behind the Robot	MCTS	✗	✓	✓	✗	✓	✓
[52]	2023	Explainable Machina how social-implicit XAI affects complex human-robot teaming tasks	MCTS, Alpha-Beta Pruning, DNN	✓	✓	✓	✓	✓	✗
[53]	2023	Natural Born Explainees how users personality traits shape the human-robot interaction with explainable robots	AlphaZero, MCTS	✓	✗	✓	✓	✓	✗
[54]	2013	Brain-computer interface controlled gaming: Evaluation of usability by severely motor restricted end-users	SMR-BCI	✗	✓	✓	✗	✓	✗
[55]	2023	Combining brain-computer interfaces and multiplayer video games: an application based on C-VEPs	c-VEP	✓	✗	✓	✗	✗	✗
[56]	2018	The bionic blues: Robot rejection lowers self-esteem	-	✗	✗	✓	✗	✓	✗

Note: Table 3 continues from Table 2 and follows the same column structure and abbreviations.

Table 4. Review Summary (Part 3).

Ref	Year	Title	Algorithm	ML	SOA	TA	Var	CA	C/M
[57]	2017	Data Mining in Adversarial Search - Players Movement Prediction in Connect-4 Games	Random, Minimax, Mini-MaxHori	✓	✓	✓	✗	✓	✓
[58]	2021	Computing Games: Bridging the Gap Between Search and Entertainment	PPNS, PNS, MCPNS, MCTS	✓	✓	✓	✗	✓	✗
[59]	2022	Scaling Laws for a Multi-Agent Reinforcement Learning Model	AlphaZero	✓	✓	✓	✗	✓	✓
[60]	2024	Formal Verification of Multi-Thread Minimax Behavior Using mCRL2 in the Connect-4	LTS, Minimax	✗	✓	✓	✓	✓	✓

Note: Table 4 continues from Table 3 and follows the same column structure and abbreviations.

Table 5 summarises the multi-dimensional roles of major AI techniques in Connect-4 research. Each technique is associated with a primary thematic focus, while also contributing to secondary areas such as tactical execution, explainability, and computational optimisation. This cross-taxonomic mapping illustrates the interconnected nature of AI strategy development and highlights how individual methods span multiple functional domains. It provides a consolidated comparison of the main methodological directions and serves as a central reference for the subsequent discussion.

Table 5. Cross-taxonomic roles of AI techniques in Connect-4 research, highlighting the evaluation metrics used across different methodological categories.

AI Technique	Primary Role (Main Sec.)	Secondary Roles (Other Sec.)
Minimax with Alpha-Beta Pruning	3.2.1	3.1, 3.3.1, 3.3.3, 3.2.2, 3.5
Monte Carlo Tree Search (MCTS)	3.2.4	3.1.1, 3.3.3, 3.3.2, 3.2.5, 3.3.1
Reinforcement Learning (TDL, MARL, Q-Learning)	3.2.3	3.1.1, 3.4.1, 3.2.5
Hybrid Models (e.g., MCTS-Minimax, RL-MCTS)	3.2.5	3.2.4, 3.2.3, 3.3
Proof Number Search (PPNS, PNS, MCPNS)	3.1.3	3.5, 3.5.1
Formal Verification (e.g., mCRL2, LTS)	3.5.2	3.2.2, 3.1.3
n-Tuple Networks	3.2.3	3.3.1, 3.1.1
XAI Methods (e.g., DeepShap, ReTA, Saliency, c-VEP)	3.4	3.4.1, 3.4.2, 3.3

Table 5 provides a central synthesis of the proposed taxonomy by explicitly mapping how major AI techniques contribute across multiple research dimensions.

4. Discussion

This discussion is organised around the core taxonomy dimensions identified in this review—computational efficiency, adaptability, interpretability, and reliability—allowing evidence from different methodological strands to be synthesised within a unified analytical framework.

Rather than reiterating individual experimental results, this section adopts a critical and integrative perspective on the Connect-4 literature. It examines how methodological choices reflect fundamental design trade-offs, explaining why particular approaches succeed under specific constraints and where their structural limitations arise. In this sense, the discussion focuses less on incremental performance gains and more on the underlying principles that shape algorithmic behaviour.

Overall, the analysis indicates that research on Connect-4 AI has evolved along several partially overlapping trajectories, including game-theoretical analysis, algorithmic optimisation, learning-based methods, explainability-oriented approaches, and formal verification. These trajectories do not converge toward a single optimal solution; instead, they reflect differing priorities in AI system design. Persistent trade-offs emerge between performance, computational efficiency, adaptability, and interpretability, underscoring the value of Connect-4 as a compact yet expressive experimental domain.

From a game-theoretical perspective, Connect-4 continues to offer a well-structured environment for analysing optimal play, complexity, and zero-sum dynamics. Although the game has been formally solved, theoretical constructs such as proof-number search, equilibrium reasoning, and complexity analysis remain influential. These frameworks shape how algorithms prioritise search, evaluate positions, and manage uncertainty. However, their reliance on fixed assumptions and static environments limits flexibility when learning, stochasticity, or rule variations are introduced.

Algorithmic developments further illustrate these tensions. Classical search-based methods, including Minimax and Monte Carlo Tree Search, benefit from determinism and predictable behaviour, making them well suited to real-time decision-making and controlled settings. Advances such as pruning strategies, parallelisation, and adaptive depth control have substantially improved their efficiency. Nonetheless, these methods depend heavily on explicit search and handcrafted heuristics, which constrains their ability to generalise beyond predefined configurations.

In contrast, reinforcement learning approaches emphasise adaptability and autonomous strategy discovery through self-play. These methods often demonstrate strong generalisation and the emergence of non-trivial strategies, but at the cost of substantial computational resources during training and evaluation. As a result, their applicability in resource-constrained or time-critical scenarios remains limited.

Hybrid models attempt to reconcile the strengths of search-based and learning-based approaches by combining deterministic planning with learned evaluation functions. While such methods frequently improve tactical precision and empirical performance, closer examination of the literature suggests that hybridisation often redistributes complexity rather than eliminating it. Computational savings in one stage of decision making may be offset by increased overhead elsewhere, indicating that hybrid approaches represent pragmatic compromises rather than definitive solutions.

Strategic and tactical studies introduce an additional dimension by focusing on how high-level decision frameworks translate into effective local actions. Techniques targeting endgame optimisation, forced move detection, and time management demonstrate that situational awareness can be as critical as global strategy. However, these approaches typically assume stable and deterministic conditions, limiting their robustness in more dynamic or uncertain environments.

The growing emphasis on explainable AI reflects a broader shift in research priorities from pure performance toward transparency and human–AI interaction. Explainability techniques enhance interpretability and user trust, particularly in interactive or educational contexts. Nevertheless, the surveyed literature reveals that explainability is often treated as an auxiliary layer applied post hoc, rather than as a core design principle. Many XAI methods rely on simplified or retrospective analyses, which may reduce their effectiveness in real-time settings and introduce additional computational overhead.

Formal verification offers a complementary perspective by addressing correctness and reliability. Verification techniques provide strong guarantees for search-based systems, ensuring consistency and preventing unintended behaviour. However, their applicability is largely restricted to static or well-defined models. As AI agents increasingly incorporate learning and adaptation, integrating formal guarantees with dynamic decision-making remains a significant open challenge.

Taken together, these findings suggest that the primary contribution of Connect-4 research lies not in achieving marginal performance improvements, but in serving as a controlled testbed for investigating fundamental AI design trade-offs. The game exposes how different methodological choices prioritise efficiency, adaptability, interpretability, or correctness, often at the expense of others. Consequently, Connect-4 offers insights that extend beyond games, informing the design of real-world decision-making systems.

To support this synthesis, Table 6 summarises the evaluation dimensions and metrics employed across the surveyed literature, ensuring complete coverage of the references discussed in Section 3. The table highlights how different metrics align with distinct methodological priorities and research objectives.

Table 6. Taxonomy of evaluation dimensions and metrics in Connect-4 research, with complete coverage of surveyed references (Section 3).

Evaluation dimension	Metric / focus	References in this survey
Game outcome performance	Win / Draw / Loss rates	[7,14,15,21,30,31,33–35,38–40,46]
Search efficiency	Nodes expanded, pruning effectiveness	[17,20,22,33,34,43,58,60]
Computational cost	Time per move, runtime, scalability	[14,15,46–48,59,60]
Learning efficiency	Episodes to convergence, sample efficiency	[24,25,27,28,30,40,59]
Learning stability	Reward variance, training robustness	[7,30,37,40]
Generalisation	Performance vs unseen agents / domains	[36,37,43,55,57]
Strategic / tactical quality	Forced moves, traps, endgame precision	[41–46]
Formal correctness	Solved status, proofs, complexity	[14,17,18,20,60]
Explainability	Saliency, reasoning transparency	[49,50,52,53]
Human–AI interaction & trust	Trust, persuasion, usability	[45,51,54–56]
Non-traditional interaction	BCI, alternative interfaces	[45,55]
Theoretical learning frameworks	Equilibrium learning, flow models	[19,32]

The temporal evolution of research interests in Connect-4 AI from 2015 to January 2026 is illustrated in Figure 8. Using the taxonomy defined in Section 3, each surveyed paper was assigned to a dominant methodological category, and annual publication counts were aggregated.

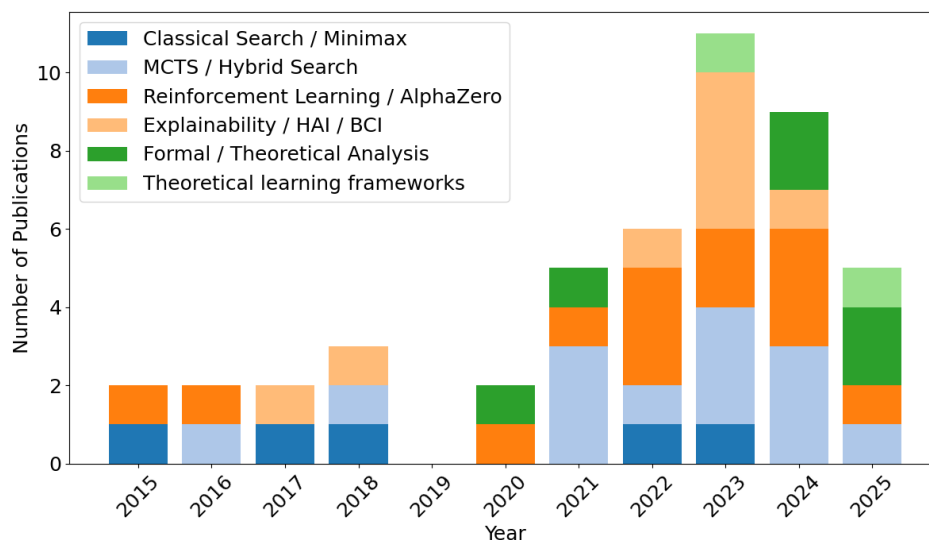


Figure 8. Timeline of research topics in Connect-4 AI (2015–2026).

Early work in this period is dominated by classical search-based approaches, such as Minimax and Alpha-Beta pruning, typically evaluated using win/loss outcomes and search efficiency metrics. From approximately 2017 onward, Monte Carlo Tree Search and hybrid methods become increasingly prominent, reflecting growing interest in anytime algorithms and scalability.

After 2018, a marked rise in reinforcement learning and AlphaZero-inspired approaches is observed, accompanied by the adoption of learning-centric metrics such as convergence speed and training stability. In the most recent period (2021–2025), the literature increasingly emphasises explainability, human-AI interaction, and formal verification, indicating a maturation of the field beyond raw performance optimisation.

These taxonomy-driven trends directly motivate the open challenges and future research directions discussed in the following section.

5. Open Challenges and Future Research Directions

Despite substantial progress across theoretical, algorithmic, and learning-based approaches, several structural challenges remain unresolved in Connect-4 AI research. A central issue is the lack of approaches that can simultaneously optimise computational efficiency, adaptability, interpretability, and reliability. Existing systems typically achieve strong performance along one or two of these dimensions while sacrificing others, indicating that the challenge is not incremental but systemic.

One persistent gap concerns real-time explainability. While many explainable AI techniques provide valuable post-hoc insights, they are rarely designed to operate under the strict time and resource constraints required for real-time gameplay. As a result, explainability is often decoupled from decision-making rather than integrated into it. Developing lightweight explanation mechanisms that function alongside search or learning processes remains an open research direction.

Within learning-based approaches, recent work has begun to explore alternative neural architectures, such as Kolmogorov-Arnold Networks (KANs), as complements or substitutes for conventional deep neural networks. These models align naturally with reinforcement-learning and hybrid search-learning paradigms, offering improved functional interpretability and potentially more transparent decision representations while remaining compatible with established game-playing frameworks. Integrating such architectures into Connect-4 agents represents a promising direction for addressing the dual challenges of performance and interpretability in real-time settings.

Generalisation beyond controlled settings also remains limited. Most Connect-4 agents are evaluated under fixed board sizes, predefined rule sets, or narrowly defined opponent models. Although variants of the game are sometimes introduced, systematic evaluation across diverse configurations is rare. This limits understanding of how well learned strategies transfer across environments and highlights the need for benchmarks that explicitly assess adaptability and robustness rather than isolated performance.

Another challenge lies in the inconsistency of evaluation practices across the literature. Studies frequently rely on different metrics, opponent strengths, computational budgets, and experimental protocols, making direct comparison difficult. The absence of standardised evaluation frameworks hinders cumulative progress and obscures trade-offs between competing approaches. Establishing common benchmarks and reporting standards would significantly strengthen empirical comparability and reproducibility.

Addressing these challenges requires a shift toward integrative research strategies that prioritise joint optimisation, real-time compatibility, and systematic evaluation. Progress in these areas would not only advance Connect-4 AI but also contribute to broader efforts in developing decision-making systems that balance performance, transparency, and adaptability in complex environments.

6. Conclusion

Connect-4 offers a structured yet tactically rich environment that remains a valuable testbed for artificial intelligence research. Despite being computationally solved, the game continues to support innovation in strategic reasoning, learning algorithms, and decision-making systems. This work presents the first dedicated taxonomy of AI research on Connect-4, categorising the literature across five integrated dimensions: game-theoretical foundations, algorithmic approaches, strategic and tactical reasoning, explainability, and computational verification. By critically analysing a broad range of studies, the review highlights how different methods address distinct gameplay challenges, from efficient deterministic search to adaptive and interpretable learning. While classical approaches such as Minimax and Alpha-Beta Pruning retain relevance in constrained scenarios, recent trends point toward hybrid, learning-based, and explainable agents that offer greater adaptability. These developments, however, introduce challenges in scalability, transparency, and integration. The findings of this review underscore the importance of advancing AI agents that are not only effective but also interpretable and formally reliable.

This work offers a structured synthesis of current research and provides a roadmap for future inquiry, supporting the development of more capable and trustworthy AI systems within Connect-4 and beyond.

Supplementary Materials: The supplementary materials are available online alongside this manuscript.

Author Contributions: Conceptualization, M.A.A.; Methodology, M.A.A.; Formal Analysis, M.A.A.; Investigation, M.A.A., A.M., A.A., R.Z.; Resources, M.A.A.; Data Curation, M.A.A.; Writing—Original Draft Preparation, A.M., A.A., R.Z.; Writing—Review and Editing, M.A.A.; Visualization, M.A.A.; Supervision, M.A.A., O.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding. The APC was funded by the authors.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article as it is a survey and review of previously published literature.

Acknowledgments: The authors would like to thank the anonymous reviewers for their constructive comments, which helped improve the clarity and quality of this manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Steele, R.; Larremore, D.B. Misère Connect Four is Solved. *ICGA Journal* **2025**, *47*, 118–129.
2. Allis, L.V. A knowledge-based approach of connect-four. *J. Int. Comput. Games Assoc.* **1988**, *11*, 165.
3. Stuart, B.L. Connect 4 as a problem in artificial intelligence and robotics. *ACM SIGCSE Bulletin* **1994**, *26*, 41–46.
4. Allis, L.V. Searching for solutions in games and artificial intelligence. *Ph.D. Thesis, University of Limburg* **1994**.
5. Scheiermann, J.; Konen, W. AlphaZero-inspired game learning: Faster training by using MCTS only at test time. *IEEE Transactions on Games* **2022**, *15*, 637–647.
6. Wu, T.R.; Guei, H.; Peng, P.C.; Huang, P.W.; Wei, T.H.; Shih, C.C.; Tsai, Y.J. Minizero: Comparative analysis of alphazero and muzero on go, othello, and atari games. *IEEE Transactions on Games* **2024**.
7. Baier, H.; Kaisers, M. Novelty in Monte Carlo Tree Search. *IEEE Transactions on Games* **2025**.
8. Zhu, Y.; Cui, G.; Liu, A.; Jia, Q.S.; Guan, X.; Zhai, Q.; Guo, Q.; Guo, X. A Reinforcement Learning Embedded Surrogate Lagrangian Relaxation Method for Fast Solving Unit Commitment Problems. *IEEE Transactions on Power Systems* **2025**.
9. Hassija, V.; Chamola, V.; Mahapatra, A.; Singal, A.; Goel, D.; Huang, K.; Scardapane, S.; Spinelli, I.; Mahmud, M.; Hussain, A. Interpreting black-box models: a review on explainable artificial intelligence. *Cognitive Computation* **2024**, *16*, 45–74.
10. Dunning, R.E.; Fischhoff, B.; Davis, A.L. When do humans heed AI agents' advice? When should they? *Human Factors* **2024**, *66*, 1914–1927.
11. Schultz, J.; Adamek, J.; Jusup, M.; Lanctot, M.; Kaisers, M.; Perrin, S.; Hennes, D.; Shar, J.; Lewis, C.; Ruoss, A.; et al. Mastering board games by external and internal planning with language models. *arXiv preprint arXiv:2412.12119* **2024**.
12. Ahmad, Z.; Jehangiri, A.I.; Ala'anzy, M.A.; Othman, M.; Latip, R.; Zaman, S.K.U.; Umar, A.I. Scientific workflows management and scheduling in cloud computing: taxonomy, prospects, and challenges. *IEEE Access* **2021**, *9*, 53491–53508.
13. Ala'anzy, M.; Othman, M. Load balancing and server consolidation in cloud computing environments: a meta-study. *IEEE Access* **2019**, *7*, 141868–141887.
14. Sheoran, K.; Dhand, G.; Dabas, M.; Dahiya, N.; Pushparaj, P. Solving connect 4 using optimized minimax and monte carlo tree search. *Mili Publications. Advances and Applications in Mathematical Sciences* **2022**, *21*, 3303–3313.
15. Taylor, H.; Stella, L. An Evolutionary Framework for Connect-4 as Test-Bed for Comparison of Advanced Minimax, Q-Learning and MCTS. *arXiv preprint arXiv:2405.16595* **2024**.
16. Kuramitsu, H.; Suzuki, K.; Matsuzawa, T. N-Tuple Network Search in Othello Using Genetic Algorithms. *Games* **2025**, *16*, 5.
17. Cutsinger, J.; Wylie, T. Row Shifting as a Puzzle Mechanic in Generalized Connect Four. In Proceedings of the 2024 IEEE Conference on Games (CoG). IEEE, 2024, pp. 1–4.
18. Bagan, G.; Duchêne, E.; Galliot, F.; Gledel, V.; Mikalački, M.; Ojjid, N.; Parreau, A.; Stojaković, M. Poset positional games. *Discrete Mathematics* **2025**, *348*, 114455.
19. Jiralerspong, M.; Sun, B.; Vucetic, D.; Zhang, T.; Bengio, Y.; Gidel, G.; Malkin, N. Expected flow networks in stochastic environments and two-player zero-sum games. *arXiv preprint arXiv:2310.02779* **2023**.
20. Primanita, A.; Khalid, M.N.A.; Iida, H. Characterizing the nature of probability-based proof number search: A case study in the othello and connect four games. *Information* **2020**, *11*, 264.
21. Touré, A.W. Evaluation of the Use of Minimax Search in Connect-4—How Does the Minimax Search Algorithm Perform in Connect-4 with Increasing Grid Sizes? *Applied Mathematics* **2023**, *14*, 419–427.
22. Tommy, L.; Hardjianto, M.; Agani, N. The analysis of alpha beta pruning and MTD (f) algorithm to determine the best algorithm to be implemented at connect four prototype. In Proceedings of the IOP Conference Series: Materials Science and Engineering. IOP Publishing, 2017, Vol. 190, p. 012044.
23. Weiner, E.M.; Montañez, G.D.; Trujillo, A.; Molavi, A. Hyperparameter Choice as Search Bias in AlphaZero. In Proceedings of the 2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC). IEEE, 2021, pp. 2389–2394.
24. Thill, M.; Koch, P.; Konen, W. Reinforcement learning with n-tuples on the game Connect-4. In Proceedings of the Parallel Problem Solving from Nature-PPSN XII: 12th International Conference, Taormina, Italy, September 1-5, 2012, Proceedings, Part I 12. Springer, 2012, pp. 184–194.

25. Yuan, C.; Al Forhad, M.A.; Bansal, R.; Sidorova, A.; Albert, M.V. Multi-agent Dual Level Reinforcement Learning of Strategy and Tactics in Competitive Games. *Results in Control and Optimization* **2024**, *16*, 100471.
26. Bagheri, S.; Thill, M.; Koch, P.; Konen, W. Online adaptable learning rates for the game Connect-4. *IEEE Transactions on Computational Intelligence and AI in Games* **2014**, *8*, 33–42.
27. Lei, S.; Lee, K.; Li, L.; Park, J. Learning Strategy Representation for Imitation Learning in Multi-Agent Games. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2025, Vol. 39, pp. 18163–18171.
28. Willemsen, D.; Baier, H.; Kaisers, M. Value targets in off-policy AlphaZero: a new greedy backup. *Neural Computing and Applications* **2022**, *34*, 1801–1814.
29. Świechowski, M.; Godlewski, K.; Sawicki, B.; Mańdziuk, J. Monte Carlo tree search: A review of recent modifications and applications. *Artificial Intelligence Review* **2023**, *56*, 2497–2562.
30. Kadam, P.; Xu, R.; Lieberherr, K. Dual Monte Carlo Tree Search. *arXiv preprint arXiv:2103.11517* **2021**.
31. Ji, J.; Thielscher, M. MCTS with Dynamic Depth Minimax. In *Advances in Computer Games*; Springer, 2023; pp. 63–75.
32. Sironi, C.F.; Winands, M.H. Analysis of the Impact of Randomization of Search-Control Parameters in Monte-Carlo Tree Search. *Journal of Artificial Intelligence Research* **2021**, *72*, 717–757.
33. Baier, H.; Winands, M.H. MCTS-minimax hybrids. *IEEE Transactions on Computational Intelligence and AI in Games* **2014**, *7*, 167–179.
34. Baier, H.; Winands, M.H. Monte-carlo tree search and minimax hybrids. In Proceedings of the 2013 IEEE Conference on Computational Intelligence in Games (CIG). IEEE, 2013, pp. 1–8.
35. Clausen, C.; Reichhuber, S.; Thomsen, I.; Tomforde, S. Improvements to Increase the Efficiency of the AlphaZero Algorithm: A Case Study in the Game 'Connect 4'. In Proceedings of the ICAART (2), 2021, pp. 803–811.
36. Cleaver, N.; Neshatian, K. Transfer Learning in Monte Carlo Tree Search. In Proceedings of the 2023 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE). IEEE, 2023, pp. 1–7.
37. Jung, J.D.; Hoey, J. Heuristic Knowledge Transfer for General Game Playing. In Proceedings of the 2024 IEEE Conference on Games (CoG). IEEE, 2024, pp. 1–8.
38. Fernández-Conde, J.; Cuenca-Jiménez, P.; Cañas, J.M. Hybrid Training Strategies: Improving Performance of Temporal Difference Learning in Board Games. *Applied Sciences* **2022**, *12*, 2854.
39. Trudeau, A.; Bowling, M. Targeted Search Control in AlphaZero for Effective Policy Improvement. In Proceedings of the Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, 2023, pp. 842–850.
40. Cotton, D.; Traish, J.; Chaczko, Z. Coevolutionary deep reinforcement learning. In Proceedings of the 2020 IEEE Symposium Series on Computational Intelligence (SSCI). IEEE, 2020, pp. 2600–2607.
41. Runarsson, T.P.; Lucas, S.M. On imitating Connect-4 game trajectories using an approximate n-tuple evaluation function. In Proceedings of the 2015 IEEE Conference on Computational Intelligence and Games (CIG). IEEE, 2015, pp. 208–213.
42. Ahmed, U.; Chatterjee, K.; Gulwani, S. Automatic generation of alternative starting positions for simple traditional board games. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2015, Vol. 29, pp. 1–10.
43. Shashkov, A.; Hemberg, E.; Tulla, M.; O'Reilly, U.M. Adversarial agent-learning for cybersecurity: a comparison of algorithms. *The Knowledge Engineering Review* **2023**, *38*, e3.
44. Goodman, J.; Perez-Liebana, D.; Lucas, S. Measuring Randomness in Tabletop Games. In Proceedings of the 2024 IEEE Conference on Games (CoG). IEEE, 2024, pp. 1–8.
45. Moreno-Calderón, S.; Martínez-Cagigal, V.; Santamaría-Vázquez, E.; Pérez-Velasco, S.; Marcos-Martínez, D.; Hornero, R. Assessing the Potential of Brain-Computer Interface Multiplayer Video Games using c-VEPs: A Pilot Study. In Proceedings of the 2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC). IEEE, 2023, pp. 1–4.
46. Hill, G.; Kemp, S.M. Connect 4: A novel paradigm to elicit positive and negative insight and search problem solving. *Frontiers in psychology* **2018**, *9*, 1755.
47. Baier, H.; Winands, M.H. Time management for Monte Carlo tree search. *IEEE transactions on computational intelligence and AI in games* **2015**, *8*, 301–314.
48. Garapati, S.; Karlapalem, K. Empirical evaluation of idle-time analysis driven improved decision making by always-on agents. In Proceedings of the 2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC). IEEE, 2018, Vol. 2, pp. 141–146.

49. Wäldchen, S.; Pokutta, S.; Huber, F. Training characteristic functions with reinforcement learning: XAI-methods play connect four. In Proceedings of the International Conference on Machine Learning. PMLR, 2022, pp. 22457–22474.
50. Duan, J.; Wang, S.; Diffenderfer, J.; Sun, L.; Chen, T.; Kailkhura, B.; Xu, K. ReTA: Recursively Thinking Ahead to Improve the Strategic Reasoning of Large Language Models. In Proceedings of the Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), 2024, pp. 2232–2246.
51. Nasir, J.; Oppliger, P.; Bruno, B.; Dillenbourg, P. Questioning wizard of oz: effects of revealing the wizard behind the robot. In Proceedings of the 2022 31st IEEE international conference on robot and human interactive communication (RO-MAN). IEEE, 2022, pp. 1385–1392.
52. Matarese, M.; Cocchella, F.; Rea, F.; Sciutti, A. Ex (plainable) machina: how social-implicit xai affects complex human-robot teaming tasks. In Proceedings of the 2023 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2023, pp. 11986–11993.
53. Matarese, M.; Cocchella, F.; Rea, F.; Sciutti, A. Natural Born Explainees: how users' personality traits shape the human-robot interaction with explainable robots. In Proceedings of the 2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN). IEEE, 2023, pp. 1786–1793.
54. Holz, E.M.; Höhne, J.; Staiger-Sälzer, P.; Tangermann, M.; Kübler, A. Brain-computer interface controlled gaming: Evaluation of usability by severely motor restricted end-users. *Artificial intelligence in medicine* **2013**, *59*, 111–120.
55. Moreno-Calderón, S.; Martínez-Cagigal, V.; Santamaría-Vázquez, E.; Pérez-Velasco, S.; Marcos-Martínez, D.; Hornero, R. Combining brain-computer interfaces and multiplayer video games: An application based on c-VEPs. *Frontiers in Human Neuroscience* **2023**, *17*, 1227727.
56. Nash, K.; Lea, J.M.; Davies, T.; Yogeewaran, K. The bionic blues: Robot rejection lowers self-esteem. *Computers in human behavior* **2018**, *78*, 59–63.
57. Ribeiro, A.C.; Rios, L.M.; Gomes, R.M.; Faria, B.M.; Reis, L.P. Data mining in adversarial search—players movement prediction in connect 4 games. In Proceedings of the 2017 12th Iberian Conference on Information Systems and Technologies (CISTI). IEEE, 2017, pp. 1–6.
58. Primanita, A.; Khalid, M.N.A.; Iida, H. Computing games: Bridging the gap between search and entertainment. *IEEE Access* **2021**, *9*, 72087–72102.
59. Neumann, O.; Gros, C. Scaling laws for a multi-agent reinforcement learning model. *arXiv preprint arXiv:2210.00849* **2022**.
60. Escobar, D.; Insuasti, J. Formal Verification of Multi-Thread Minimax Behavior Using mCRL2 in the Connect 4. *Mathematics* **2025**, *13*, 96.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.