

General Game Playing

MICHAEL GENESERETH, Computer Science Department, Stanford University, USA

General game players are agents able to play strategy games based solely on formal game descriptions supplied at run time. In other words, they do not know the rules until the game starts! Unlike specialized game players, such as Deep Blue, general game players cannot rely on algorithms designed in advance for specific games; they must discover such algorithms themselves. General game playing expertise depends on intelligence on the part of the game player rather than intelligence of the programmer of the game player.

CCS Concepts: • **Software and its engineering** → **Software creation and management; Designing software; Requirements analysis;**

Additional Key Words and Phrases: General Game Playing, Logic Programming, Artificial Intelligence

ACM Reference Format:

Michael Genesereth. 2025. General Game Playing. *J. ACM* 37, 4, Article 111 (August 2025), 11 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 Introduction

Games of strategy, such as chess and checkers, couple intellectual activity with competition. We can exercise and improve our intellectual skills by playing such games. The competition adds excitement and allows us to compare our skills to those of others. The same motivation accounts for interest in Computer Game Playing as a testbed for Artificial Intelligence (AI). The idea is that programs that think better should be able to win more games, and so we can use game playing as an evaluation technique for intelligent systems.

Unfortunately, building programs to play specific games has limited value in this regard. (1) To begin with, specialized game players have a very narrow focus. They can be good at one game but not another. Deep Blue [2] may have beaten the world Chess champion, but it has no clue how to play checkers. (2) A second, more subtle problem with specialized game playing systems is that they do only part of the work. Most of the interesting analysis and design is done in advance by their programmers. The systems themselves might as well be tele-operated.

All is not lost. The idea of game playing can be used to good effect to inspire and evaluate good work in Artificial Intelligence, but it requires moving more of the design work to the computer itself. This can be done by focussing attention on General Game Playing.

General game players are systems that accept descriptions of arbitrary games at runtime and that use such descriptions to play those games effectively without human intervention. In other words, they do not know the rules until the games start.

Unlike specialized game players, such as Deep Blue [blue] and Stockfish, general game players must be able to play different kinds of games. They must be able to play simple games (like Tic Tac Toe) and complex games (like Chess), games with differing numbers of players, cooperative games

Author's Contact Information: Michael Genesereth, genesereth@stanford.edu, Computer Science Department, Stanford University, Stanford, California, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 1557-735X/2025/8-ART111

<https://doi.org/XXXXXXXX.XXXXXXX>

and competitive games, games with or without communication among the players, and so forth. Importantly, they must be able to play games they have never seen before, including games that no-one has seen before. See example in figure 1.



Fig. 1. An unusual version of Chess

Importantly, general game players cannot rely on algorithms designed in advance for specific games as in the case of specialized game players. General game playing expertise depends on intelligence on the part of the game player and not just intelligence of the programmer of the game player.

General Game Playing underscores the importance of knowledge representation, reasoning, and rational decision making in a world increasingly focussed on machine learning and large language models. It is worth noting that, despite important and impressive game-related developments in these latter areas (e.g. [1][25][16][11]), no machine learning program has ever won an official general game playing competition, due to time limits on game play; and it is unlikely that systems based on large language models would perform any better, due to limits on time and space and due to "obfuscation" (wherein words in game descriptions are replaced by nonsense words). By contrast, undergraduates with minimal training in knowledge representation and reasoning methods are able to create GGP programs running on laptops that excel in such competitions and frequently beat humans.

2 Recent History

The idea of *general problem solving* dates from the earliest days of AI. In 1958, Newell and Simon proposed a *General Problem Solver* [19] that worked by applying general, task-independent problem-solving processes to descriptions of domain-specific tasks. The first explicit mention of *general game playing* appeared in a subsequent paper by Jacques Pitrat [21] in 1968. A related notion was described by Barney Pell [20], who further developed the concept in his 1993 doctoral thesis. In the early 2000s, researchers at Stanford enlarged the concept from chess-like games to arbitrary discrete dynamic systems, thus broadening the concept and making it applicable beyond the world of recreational games.

In order to promote work on GGP, in 2005 the Association for the Advancement of Artificial Intelligence established the International General Game Playing Competition [8][7], an annual contest to determine the best automated general game players in the world. The following table shows the winners of the first dozen competitions.

Year	Player	Developer(s)
2005	Cluneplayer	Clune (USA)
2006	Fluxplayer	Schiffel, Thielscher (Germany)
2007	Cadiaplayer	Bjornsson, Finnsson (Iceland)
2008	Cadiaplayer	Bjornsson, Finnsson (Iceland)
2009	Ary	Mehat (France)
2010	Ary	Mehat (France)
2011	TurboTurtle	Schreiber (USA)
2012	CadiaPlayer	Bjornsson, Finnsson (Iceland)
2013	TurboTurtle	Schreiber (USA)
2014	Sancho	Draper, Rose (UK)
2015	Galvanise	Emslie (UK)
2016	WoodStock	Piette (France)

In addition to pitting automated players against each other, the Competition frequently included a man-machine demonstration match between the competition winner and a human player(s). While the human player won the first of these demonstrations, the computer won all of the subsequent competitions. For example, in 2012, CadiaPlayer, in addition to defeating the other automated competitors, defeated the human race, represented by Chris Welty, one of the developers of Watson, the IBM computer that defeated the best players on the American game show Jeopardy! (As a consolation prize, the human was awarded two bottles of Scotch to ease his disappointment at letting down the human race.)



Fig. 2. 2012 Carbon vs Silicon match: The human player taking counsel from another human

The International GGP Competition was suspended after 2016. However, by then, the competition had served its primary purpose - it led to workshops on GGP at multiple international conferences and the publication of numerous research papers on GGP and its applications. Today, regional competitions continue to be run, and there are rumors that the International Competition will resume before long, albeit in slightly different format.

3 Game Playing

Despite the variety of games treated in General Game Playing, all games share a common abstract structure. Each game takes place in an environment with finitely many states, with one distinguished initial state and one or more terminal states. In addition, each game has a fixed, finite number of players; each player has finitely many possible actions in any game state, and each state has an associated goal value for each player.

Given this common structure, we can think of a game as a state graph, like the one shown in figure 3. In this case, we have a game with one player, with eight states (named s_1, \dots, s_8). The arcs in this graph capture the transition function for the game. For example, if the game is in state s_1 and a player does action a , the game will move to state s_2 . if the game is in state s_1 and a player does action b , the game will move to state s_5 . Each game has just one initial state, in this case s_1 . However, there can be any number of terminal states, in this case s_4 and s_8 . The numbers associated

with each state indicate the utilities of those states for the player. Players earn those scores only in terminal states. However, they are provided for all states and in some games indicate incremental progress.




Fig. 3. Game as a state graph

We can extend this model to accommodate games with multiple players with two modifications. (1) We annotate each state with rewards for all players in the game. (The rewards can be the same or different for different players.) (2) For each state, we specify which player is in control, i.e. whose turn it is to play. Turns need not strictly alternate; in some cases, a single player may get several turns in a row. (This model can be extended to games with simultaneous moves, but for the sake of simplicity, we avoid that complexity in what follows.)

Since all of the games that we are considering are finite, it is possible, in principle, to describe such games in the form of state graphs. Unfortunately, such explicit representations are not practical in all cases. Even though the numbers of states and actions are finite, these sets can be extremely large; and the corresponding graphs can be larger still. For example, in chess, there are more than 10^{40} possible states.

In the vast majority of games, states and actions have composite structure that allows us to define a large number of states and actions in terms of a smaller number of more fundamental entities. In Chess, for example, states are not monolithic; they can be conceptualized in terms of pieces, squares, rows and columns and diagonals, and so forth. By exploiting this structure, it is possible to encode games in a form that is more compact than direct representation.

The first step in solving this problem is to conceptualize states as *datasets* (as sets of facts that are true in those states) and to conceptualize actions as operations applied to arguments, as suggested by the graph in figure 4.

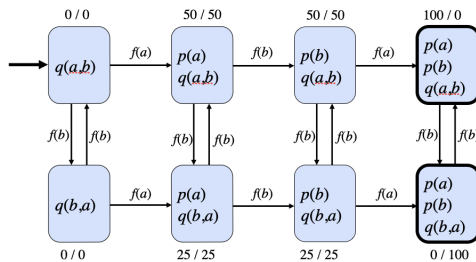


Fig. 4. Game as a "structured" state graph

Given a "structured" state graph of this sort, the second step is to encode game rules in terms of these datasets and actions. In GGP, the most popular game description language is GDL (for Game Description Language) [17]. Rules in GDL are written in Epilog [9] (a dynamic logic programming language based on Prolog). See below for the rules of Tic Tac Toe as written in GDL.

```

role(x)
role(o)

init(cell(1,1,b))
init(cell(1,2,b))
init(cell(1,3,b))
init(cell(2,1,b))
init(cell(2,2,b))
init(cell(2,3,b))
init(cell(3,1,b))
init(cell(3,2,b))
init(cell(3,3,b))
init(control(x))

legal(mark(M,N)) :- cell(M,N,b)

mark(M,N) :: control(R) ==> cell(M,N,R) & ~cell(M,N,b)
mark(M,N) :: control(x) ==> ~control(x) & control(o)
mark(M,N) :: control(o) ==> ~control(o) & control(x)

line(Z) :- row(M,Z)
line(Z) :- column(M,Z)
line(Z) :- diagonal(Z)
row(M,X) :- cell(M,1,X) & cell(M,2,X) & cell(M,3,X)
column(N,X) :- cell(1,N,X) & cell(2,N,X) & cell(3,N,X)
diagonal(X) :- cell(1,1,X) & cell(2,2,X) & cell(3,3,X)
diagonal(X) :- cell(1,3,X) & cell(2,2,X) & cell(3,1,X)

goal(x,100) :- line(x) & ~line(o)
goal(x,50) :- ~line(x) & ~line(o)
goal(x,0) :- ~line(x) & line(o)
goal(o,100) :- ~line(x) & line(o)
goal(o,50) :- ~line(x) & ~line(o)
goal(o,0) :- line(x) & ~line(o)

terminal :- line(x)
terminal :- line(o)
terminal :- ~open

open :- true(cell(M,N,b))

```

There are two players - x and o. In the initial state, all cells are blank (represented here by the symbol b), and the x player has control. It is legal for the player in control to mark a cell if that cell

is blank. (In GDL, symbols that begin with capital letters are variables, while symbols that begin with lower case letters are constants. The :- operator is read as "if" - the expression to its left is true if the expressions that follow it are true.) If the player in control marks a cell, that cell contains the player's mark in the next state and the b is removed. Also, control alternates on each play. (The :: operator relates the action mentioned to its left as a transition rules that applies whenever the conditions before the ==> are true and produces a state in which the conditions to its right are true.) A line is a row of marks, a column of marks, or a diagonal. The x player / o player receives 100 points in any state with a line of x marks / o marks; it receives 50 points in any state with no lines; and it gets 0 points in any state with a line of o marks / x marks. Finally, a game terminates whenever there is a line of player marks or if the game is no longer open, i.e. there are no blank cells.

The main thing to note about this example is that one page of rules fully describes a game of thousands of states. That's a significant saving over the state graph for Tic Tac Toe (which contains more than 5000 states). The compression in more complex games can be even more dramatic. For example, it is possible to describe the rules of Chess in just four pages of rules like the ones above.

Interestingly, in order to prevent programmers from building in specialized capabilities for specific words in game descriptions, it is common in general game playing to *obfuscate* descriptions. All words are consistently replaced by nonsense words. The only exceptions are variables and game-independent constants, such as role, legal, terminal.

4 Game Playing

Having a formal description of a game is one thing; being able to use that description to play the game effectively is something else entirely. The player must be able to compute the initial state of the game. It must be able to compute which moves are legal in every state. It must be able to determine the state resulting from particular moves. It must be able to compute the value of each state for each player. And it must be able to determine whether any given state is terminal.

From a description like this, a general game player can reconstruct a game tree. See the partial game tree shown in figure 5. The player starts with the initial state, computes the legal moves in that state, and for each move deduces the next state. This is done repeatedly to expand the tree until a terminal state is reached on each branch. Given a game-tree, a player can use minimax [22] or an equivalent technique to determine its best possible move; and it can save work by using sophisticated search techniques, such as alpha-beta pruning [14].




Fig. 5. Partial Game Tree for Tic Tac Toe

Unfortunately, it is not always possible to search to the end of the game tree. In Tic-Tac-Toe, there are thousands of states. This is large but manageable. In Chess there are more than 10^{40} states [24]. A state space of this size, being finite, is fully searchable in principle but not in practice. Moreover, the time limit on moves in most games means that players must select actions without knowing with certainty whether they are the best or even good moves to make.

The alternative is to do incomplete search, on each move expanding the game tree as much as possible within the available time and then making a choice based on the estimated values of non-terminal states. In traditional game playing, where the rules are known in advance, the programmer can invent game-specific evaluation functions to help in this regard. For example, in chess, we know that states with higher piece count and greater board control are better than ones with less material or lower control. Unfortunately, it is not possible for a GGP programmer to invent such game-specific rules in advance, as the game's rules are not known until the game begins. The program must evaluate states for itself. Doing this effectively is the key to effective general game playing.

The approach used in early GGP programs was to develop game-independent heuristics [3] [23] [13] [15], e.g. proximity to a goal state, player mobility, and opponent restriction. Consider mobility. Proponents argue that, all other things being equal, it is better to move to a state that affords the player greater mobility, that is more possible actions. Better than being boxed into a corner. Symmetrically, proponents of mobility argue that it is good to minimize the mobility of one's opponents. All of these heuristics have been shown to be effective in some games. Unfortunately, they are only heuristics. They frequently fail, sometimes with comical consequences.

The final match of GGP-06 is an example. The game was cylinder checkers, i.e. checkers played on a cylinder. In this game, as in ordinary checkers, a player is permitted to move one of his ordinary pieces (pieces that are not kings) one square forward on each turn. In the position shown in figure 6 red is moving from top to bottom and black is moving from bottom to top. If a piece is blocked by an opponent's player, he can "jump" that player if there is an empty square on the other side. Moreover, the player *must* make such a jump if one is available. The objective of the game is to take all or as many of the opponent's pieces as possible while preserving one's own pieces. Here is a snapshot of the game. It is red's turn to play. What should it do? And what do you think it did?




Fig. 6. Position in Final Match of IGGPC 2006

Here's a hint. The player in this case was Cluneplayer [3]; and it had decided, for some reason or other, that limiting the opponent's mobility was a good heuristic. If it were to move the rearmost piece, black would have multiple moves. However, if it were to move the piece in front, black would be forced to capture its piece. In other words, it would have at most one move. Clearly, moving the forward piece minimizes the opponent's mobility, so that is what Cluneplayer did. Actually, the whole match played out this way, with red giving black captures at every opportunity. It was sad to

watch but also a little comical. The moral is that, while non-guaranteed heuristics are sometimes useful, they are not always useful.

5 Monte Carlo Search

The second generation of GGP programs used a different approach to evaluating non-terminal states, notably Monte Carlo Search and Monte Carlo Tree Search.

In Monte Carlo Search (MCS), the player expands the tree a few levels. Then, rather than using a local heuristic to evaluate a state, it makes some probes from that state to the end of the game by selecting random moves for all players. It sums up the total rewards for all such probes and divides by the number of probes to obtain an estimated utility for that state. It can then use these expected utilities in comparing states and selecting actions.




Fig. 7. Monte Carlo Search Tree

Monte Carlo Tree Search [6] improves on this by replacing the random choice by more intelligent selections. The result of introducing such techniques was dramatic. Suddenly, automated general game players began to perform at a high level. Using MCTS, Cadioplayer won the competition 3 times.

The use of statistical techniques changed the character of general game players from curiosities to programs capable of serious play. Virtually every general game playing program today includes some variant of MCS or MCTS. Unfortunately, statistical search alone has weaknesses, notably on games with complex descriptions requiring a lot of computation time and on games where there are few goal states; and, consequently, it must be combined with other methods to be successful.

6 Metagaming

Whereas the GGP techniques described above concentrated on game tree search done at *runtime*, subsequent work was focussed on *offline* processing of game descriptions. Examples include things like factoring games into independent subgames, reformulating game descriptions, and automatic programming. Often it is the case that this sort of processing more than pays for itself. In such cases, players can expend a little time up front and gain a lot in processing time later.

This is really what programmers do when building specialized game players and now we are building game players that can do these things for themselves. And this is what General Game Playing was intended to stress from its very beginnings.

One example of offline analysis is game decomposition, also called *factoring* [4][10]. Consider the game of Hodgepodge pictured in figure 8. Hodgepodge is actually two games glued together.

Here we show chess and othello, but it could be any two games. One move in a joint game of Hodgepodge corresponds to one move in each of the two constituent games. Winning requires winning at least one of the two games while not losing the other.




Fig. 8. Hodgepodge

What makes Hodgepodge interesting is that it is factorable, that is it can be divided into two independent games. Realizing this can have dramatic benefit. To see this, consider the size of the game tree for hodgepodge. Suppose that one game tree has branching a and the other has branching factor b . Then the branching factor of the joint game is a times b , and the size of the fringe of the game tree at level n is $(a*b)^n$. However, the two games are independent. Moving in one subgame does not affect the state of the other subgame. So, the player really should be searching two smaller game trees, one with branching factor a and the other with branching factor b . In this way, at depth n , there would be only $a^n + b^n$ states. This is a huge decrease in the size of the search space. Luckily, in many cases, it is possible to discover such factors in time proportional to the size of the game description.

The table below illustrates the computational benefits of factoring on three highly factorable games - Multiple Buttons and Lights, Multiple Switches, and Multiple Tic Tac Toe. The first column in the table identifies the game; the second column indicates the depth to which the tree is searched; the third column lists the number of milliseconds to search the game tree to that depth without factoring; and the last column shows the number of milliseconds after factoring.

Game	Depth	Unfactored Cost	Factored Cost
Multiple Buttons and Lights	4	16,700	6
Multiple Switches	5	22M	210
Multiple Tic Tac Toe	3	88,000	150

Factoring is just one example of game reformulation. There are many others. For example, it is sometimes possible to find symmetries in games that cut down on search space. In some games, there are bottlenecks that allow for a different type of factoring. Consider, for example, a game made up of one or more subgames in which it is necessary to win one game before moving on to a second game. In such a case, there is no need to search to a terminal state in the overall game; it is sufficient to limit search to termination in the current subgame. These examples are extreme cases, but there are many simpler everyday examples of finding structure of this sort that can help in curtailing search.

The trick in metagaming is to analyze and/or reformulate a game without expanding the entire game tree. The interesting thing about general game playing is this - that sometimes the cost of analysis is proportional to the size of the description rather than the size of the game tree, as in the example mentioned above. In such cases, players can expend a little time and gain a lot in search savings.

Of course, factoring games and finding transformations like these require a lot of algorithmic expertise. What we really need is Donald Knuth in a box. Or maybe Corman and Leiserson. And, of

course, since we are talking about games, we should have expertise at game tree search. So we had better cram Jonathan Schaeffer into our box as well.

An important point here is that GGP is not just about game tree search; to an even greater extent, it is about game descriptions and their use, and ultimately about automatic programming.

7 Conclusion

General game playing is an interesting application in its own right. It is intellectually engaging and more than a little fun. But it is much more than that. (1) It serves as a laboratory for practical applications, in business and law, science and engineering. In fact, some of the games used in competitions are drawn from such areas. (2) More fundamentally, it provides a theoretical framework for defining rationality in a way that takes into account problem representation, incompleteness of information, and resource bounds. The upshot is that it raises questions about the nature of intelligence and serves as a laboratory in which to evaluate competing approaches to intelligence.

General Game Playing is a setting within which AI is the essential technology. It concentrates attention on the notion of runnable specifications. Building systems of this sort dates from the early years of AI.

It was in 1958 that John McCarthy invented the concept of the "advice taker" [18]. The idea was simple. He wanted a machine that he could program by description. He would describe the intended environment and the desired goal, and the machine would use that information in determining its behavior. There would be no programming in the traditional sense. McCarthy presented his concept in a paper that has become a classic in the field of AI.

The main advantage we expect the advice taker to have is that its behavior will be improvable merely by making statements to it, telling it about its environment and what is wanted from it. To make these statements will require little, if any, knowledge of the program or the previous knowledge of the advice taker.

An ambitious goal! But that was a time of high hopes and grand ambitions. The idea caught the imaginations of numerous subsequent researchers – notably Bob Kowalski, the high priest of logic programming, and Ed Feigenbaum, the inventor of knowledge engineering. In a paper written in 1974, Feigenbaum gave his most forceful statement of McCarthy's ideal [5].

The potential use of computers by people to accomplish tasks can be "one-dimensionalized" into a spectrum representing the nature of the instruction that must be given the computer to do its job. Call it the what-to-how spectrum. At one extreme of the spectrum, the user supplies his intelligence to instruct the machine with precision exactly how to do his job step-by-step. ... At the other end of the spectrum is the user with his real problem. ... He aspires to communicate what he wants done ... without having to lay out in detail all necessary subgoals for adequate performance.

Some have argued that the way to achieve intelligent behavior is through specialization. That may work so long as the assumptions one makes in building such systems are true. For general intelligence, however, general intellectual capabilities are needed, and such systems should be capable of performing well in a wide variety of tasks. To paraphrase the words of Robert Heinlein [12].

*A human being should be able to change a diaper, plan an invasion, butcher a hog, conn a ship, design a building, write a sonnet, balance accounts, build a wall, set a bone, comfort the dying, take orders, give orders, cooperate, act alone, solve equations, analyze a new problem, pitch manure, program a computer, cook a tasty meal, fight efficiently, die gallantly. **Specialization is for insects.***

References

- [1] B. Bannerjee, G. Kuhlmann, and P. Stone. 2006. Value Function Transfer for General Game Playing. In *Proceedings of the ICML Workshop on Structural Knowledge Transfer for Machine Learning*.
- [2] M. Campbell, A. J. Hoare, and F.-H. Hsu. 2002. Deep Blue. *Artificial Intelligence* 134, 1-2 (2002), 57–83.
- [3] J. Clune. 2007. Heuristic evaluation functions for General Game Playing. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*. AAAI, 1134–1139.
- [4] E. Cox, E. Schkufza, R. Madsen, and M. Genesereth. 2009. Factoring General Games Using Propositional Automata. In *IJCAI-09 Workshop on General Game Playing (GIGA-09)*. IJCAI.
- [5] E. Feigenbaum. 2007. *ACM Turing award lectures*. Association for Computing Machinery, Chapter How the 'what' becomes the 'how'.
- [6] H. Finnsson and Y. Bjornsson. 2008. Simulation-Based Approach to General Game Playing. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI-08)*. AAAI, 259–264.
- [7] M. Genesereth and Y. Bjornsson. 2012. The International General Game Playing Competition. *AI Magazine* 34, 2.
- [8] M. Genesereth, N. Love, and B. Pell. 2005. General Game Playing: Overview of the AAAI Competition. In *AI Magazine*, Vol. 26. AAAI, 62–72.
- [9] M. R. Genesereth and V. K. Chaudhri. 2020. *Introduction to Logic Programming*. Springer.
- [10] M. Gunther, S. Schiffel, and M. Thielscher. 2009. Factoring General Games. In *IJCAI-09 Workshop on General Game Playing (GIGA-09)*. IJCAI.
- [11] M. Hausknecht, P. Khandelwal, R. Miikkulainen, and P. Stone. 2012. HyperNEAT-GGP: a HyperNEAT-Based Atari General Game Player. In *GECCO 12: Proceedings of the 14th annual conference on Genetic and evolutionary computation*.
- [12] Robert A. Heinlein. 1988. *Time Enough for Love*. Ace Books.
- [13] M. Kirci, N.R. Sturtevant, and J. Schaeffer. 2011. A GGP Feature Learning Algorithm. *Kunstliche Intelligenz* 25, 1 (2011), 35–41.
- [14] D. E. Knuth and R. W. Moore. 1975. An analysis of alpha-beta pruning. *Artificial Intelligence* 6, 4 (1975), 293–326.
- [15] G. Kuhlmann and P. Stone. 2006. Automatic Heuristic Construction in a Complete General Game Player. In *Proceedings of the 21st AAAI Conference on Artificial Intelligence (AAAI-06)*. AAAI, 1457–1462.
- [16] J. Levine, C. B. Congdon, M. Ebner, G. Kendall, S. Lucas, R. Miikkulainen, T. Schaul, and T. Thompson. 2013. General Video Game Playing. In *Dagstuhl Seminar on Artificial and Computational Intelligence in Games*. 1–7.
- [17] N. Love, T. Hinrichs, and M. Genesereth. 2006. Game Description Language Specification.
- [18] J. McCarthy. 1958. Programs with common sense. In *Symposium on Mechanization of Thought Processes*. National Physical Laboratory.
- [19] A. Newell, J.C. Shaw, and H.A. Simon. 1959. Report on a general problem-solving program. In *Proceedings of the International Conference on Information Processing*. 256–264.
- [20] B. Pell. 1993. *Strategy Generation and Evaluation for Meta-Game Playing*. Trinity College.
- [21] J. Pitrat. 1968. Realization of a general game-playing program. *IFIP Congress 2* (1968), 1570–1574.
- [22] S. J. Russell and P. Norvig. 2021. *Artificial Intelligence: A Modern Approach*. Pearson.
- [23] S. Schiffel and M. Thielscher. 2007. Fluxplayer: A Successful General Game Player. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI-07)*. AAAI, 1191–1196.
- [24] C. E. Shannon. 1950. Realization of a general game-playing program. *Philos. Mag.* 41, 314 (1950), 256–275.
- [25] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, and T. Graepel. 2018. A General Reinforcement Learning Algorithm That Masters Chess, Shogi, and Go Through Self-Play. *Science* 362, 6419 (2018), 1140–1144.