A Game That Stymies AI

Steven J. Brams, New York University

Abstract

Catch-Up is a simple 2-person sequential game in which one player (A) begins by choosing a number from the set of natural numbers, $\{1, 2, 3, \ldots, n\}$. The other player (B) then chooses one or more numbers whose sum equals or just exceeds A's number. The players then alternate making choices of numbers, without replacement, so that their sums equal or just exceed their opponent's on each round—until all numbers have been chosen—culminating in one player's sum equaling or exceeding its opponent's sum and making it the tied or absolute winner. Unlike Chess or Go, no AI (artificial intelligence) or deep-learning program has been found that consistently beats an opponent in Catch-Up—say, 90% or more of the time—who randomizes its choices of numbers on each round, whereas making random moves in Chess or Go would be disastrous. Has AI met its match in its strongest fields, computation and learning?

1. Introduction

AI (artificial intelligence) has succeeded spectacularly in fields of specialized learning, such as playing Chess or Go, and in finding patterns in visual displays, such as identifying when a human organ is diseased or otherwise abnormal. It has done much less well in situations requiring generalized learning, such as in understanding a text, including one that it may be able to expertly translate from one language into another. Worse, acting human-like by expressing—and, especially, feeling—appropriate emotions seems well beyond AI's capacity.

AI's panoply of tools, such as neural networks, have helped to fuel recent advances, including the design of self-driving cars and trucks. Such advances require enormous computational power as well as smart algorithms to put this power to good use. With uncanny but not perfect accuracy, AI has been able to identify driving situations in which it is safe to proceed, slow down, speed up, turn, stop, or choose myriad variations on these basic actions, though not without incurring some fatal accidents.

It is surprising, therefore, to discover a game, Catch-Up, whose absurdly simple rules and a particular strategy have so far foiled AI's ability to outplay human players, whom it has trounced in much more complex games like Chess and Go. This is especially ironic when the human player makes random choices, which would be disastrous against an AI opponent in Chess or Go. I will suggest reasons why this is the case later, but first consider what constitutes optimal play in Catch-Up.

2. Catch-Up and Optimal Play

Catch-Up has five rules of play:

1. Two players alternately choose numbers, without replacement, from the set of natural numbers, {1, 2, 3, …, *n*}.
2. The first player to choose a number, is P1, and the second player is P2. At the outset, P1 chooses one of the *n* original numbers.
3. Thereafter, P1 and P2 successively choose one or more of the remaining numbers, but each must stop—and turn play over to the other player—when the sum of its choices up to that point equals or just exceeds its opponent's previous sum.
4. The goal of the players is to have a higher sum than an opponent at the end of play—and by as much as possible—or that failing, to have the same sum. If neither of these goals is achievable, a player prefers to lose by as small amount as possible.
5. The game ends when all numbers have been chosen, and one player's sum equals or exceeds its opponent's sum, making it the tied or absolute winner.[1]

I illustrate these rules and optimal play in two simple cases, ignoring the trivial case in which there is only the number 1, which P1 will choose and win:

1. If the numbers are {1, 2}, P1 will choose 2, and P2 cannot catch up and so loses by 2-1.
2. If the numbers are {1, 2, 3}, there are three cases:

    a. If P1 chooses {3}, P2 will choose {1, 2} in either order and guarantee a 3-3 tie.

[1] The *idea* of Catch-Up has been applied in different sports, wherein the player or team that is behind in a game is given the opportunity to catch up by being given such prerogatives as kicking next (penalty kicks in soccer) or serving next (in some racquet sports). See Brams et al. (2018a, 2018b).

b. If P1 chooses {2}, P2 will choose {1, 3}, in that order, and guarantee a 4-2 win, which is better for P2 than choosing only {3} that leads to a 3-3 draw.
c. If P1 chooses {1}, P2 will choose {3} and guarantee a 3-3 tie after P1 subsequently chooses {2} on the next round.

Clearly, P1 can force a draw in the cases 2a and 2c, whereas in case 2b P1 will lose when P2 plays optimally. Thus, P1 should not choose {2} initially but instead should choose {1} or {3} to ensure a draw.

In asking what P1 should do at the outset, one first determines what is optimal for P2 to do subsequently and then determines the consequences for P1. In the {1, 2, 3} example, P1 cannot force a win, as it can in the {1, 2} case, but it can force a draw by choosing {1} or {3} initially.

The thought process underlying optimal play of Catch-Up is *backward induction*, a venerable idea going back more than 100 years that is used in the solution of sequential games like Chess (Zermelo, 2013; for a translation of this paper and a discussion of its results, see Schwalbe and Walker, 2001). The idea is to work from the end of the game back to the beginning, assuming at each stage that the priority of the players is as given by their goals in rule 4 of the rules of play. I assume that each player, looking ahead, anticipates that it and its opponent will make optimal choices when it is each's turn to choose from the set of numbers still available.

Applying computer-assisted backward induction to all sets of natural numbers from $n = 2$ to 20, Isaksen et al. (2015) found that when the sum of all numbers is odd, as it is when $n = 2$ (1+2 = 3), optimal play results in a win by either P1 or P2 by a difference in their sums of 1 (2–1 = 1). But when the sum of the numbers is even, as it is when $n = 3$ (1+2+3 = 6), optimal play results in a draw (1+2 = 3).

When the sum of the numbers is odd, one might think that P1, by going first, might always be able to force a win, but this is not the case. Excluding the cases of $n = 1$ and 2 that I just discussed, P2 can force a win when $n = 9$, 10, 14, or 18, whereas P1 can force a win when $n = 5$, 6, 13, or 17. In all other cases ($n = 3$, 4, 7, 8, 11, 12, 15, 16, 19, 20), optimal play produces a draw.

Thus, unlike Chess, in which there is a general agreement that optimal play leads to a draw, or possibly a win by white (less likely, a win by black), backward-induction calculations for Catch-Up show that neither P1 nor P2 wins more often than the other between $n = 3$ and 20 (four times each). Optimal play in the other 10 cases between $n = 3$ and 20 leads to a draw.

3. Randomization Vs. AI

Isaksen et al. (2015) compared the performance of a player who randomizes its choices to three other strategies (more on these shortly). To illustrate what randomization means, assume n = 5, so the numbers at the start are {1, 2, 3, 4, 5}. Randomization works as follows:

P1's first choice

P1 chooses each of the 5 numbers with probability 1/5.

P2's first choice

For purposes of illustration, assume that P1 chooses {3}. Then there are eight possible subsets of the remaining numbers whose sum equals or exceeds 3:

{4} / {5} / {1, 2} / {1, 4} / {1, 5} / {2, 1} / {2, 4} / {2, 5}

Randomizing these choices means that P2 chooses each with probability 1/8.

In six of these cases, the subsets comprise two numbers, wherein the first number—either 1 or 2—is less than 3, so a second number, which is the second number in each subset, is needed to make the sum for P2 equal to or greater than 3. After P2 chooses one of the eight subsets at random, either two or three numbers remain.

P1's second choice

P1 will choose again, and at random, a subset of the remaining numbers such that, when they are added to P1's present score of 3, equals or exceeds P2's score. Depending on the numbers that P1 chooses when it makes a second choice, P2 may or may not have a second choice that equals or exceeds P1's last total.

In summary, when a player randomizes, which I call the strategy RD, it chooses with equal probability any of the subsets of available numbers that equal or exceed an opponent's last total. Isaksen et al. (2015) tested this strategy against three other strategies when n = 10—none using AI—in which an opponent, at each of its turns, (1) maximizes its score, increasing its lead by as much as possible; (2) minimizes its score, decreasing its lead by as much as possible; or (3) uses the maximum number of numbers possible, reducing the numbers available to its opponent by as much as possible.

How do these strategies do against RD? RD won the following percentages of 100,000 games against each of these three kinds of opponents: (1) 41.7%, (2) 60.9%, and (3) 46.5%, suggesting that its performance was middling. The most decisive win score from matching each strategy, including RD, against every other is that strategy (2) beat strategy (1) in 79.7% of games.[2]

4. Problems of AI in Catch-Up

Can an AI player do better against RD than any of the non-AI strategies? It would seem yes if AI can learn to make optimal choices by using backward induction.

The rub is that backward induction may not be learnable by AI, because it's a human discovery based on reasoning backwards from a last move. While this move can be determined if the players make optimal choices at every stage, this is not what RD does. Because RD's choices are random and therefore unpredictable, AI cannot backward induct from a known last move of the game.

To be sure, after every (unpredictable) move of RD, AI can consider all possible moves from that point onward and determine what moves are optimal for it to choose for every contingency that can arise. These optimal strategies may change radically after each random move of RD, but this is not the fundamental problem. More fundamental is whether it is possible, from playing against itself hundreds of thousands or millions of times, that any pattern resembling what backward induction prescribes can be learned.

It is not evident that this is possible when AI's opponent is RD, because there is no rhyme or reason to RD's choices.[3] Indeed, any regularity that AI detects in RD's play will be a chimera.

---

[2] Because the sum of the first 10 natural numbers is odd (55), we know that either P1 or P2 can force a win, based on backward induction, which turns out to be P2 by 28 to 27. This cannot readily be shown by hand; but in the simpler case of the first 5 natural numbers, it is not difficult to show that P1 can force a win by choosing 3 initially. Whichever of the eight possible strategies that P2 can choose to equal or exceed 3, P1 can force a win. For example, if P2 chooses {2, 5}, P1 will choose {1, 4} and win by 8 to 7; similarly, if P2 chooses {1, 2} and ties P1, P1 will then choose {5} and win by the same score when P2 must choose 4.

[3] RD's strategy is somewhat akin to the use of optimal mixed (random) strategies in 2-person, zero-sum games, in which the value of the game is the same regardless of what strategy an opponent chooses.

But for argument's sake, assume that AI is able to infer that RD is making random choices. Then AI's optimal strategy would be to stop trying to learn and instead use backward induction.

Even if this option were allowed, however, applying backward induction may not be feasible. Because the number of possible choices of a player in Catch-Up increases exponentially with *n*, no present or foreseeable computer will be able to find optimal strategies in Catch-Up for, say, *n* = 100, just as Chess has proved impregnable to this determination.

This is not true of Checkers. It was shown, using multiple computers making calculations over almost two decades, that optimal strategies always produce a draw in this game (Schaeffer et al., 2007).

Unlike Catch-Up, however, a player who randomizes its moves in Chess will lead to its quick demise (i.e., checkmate) against even an amateur player, not to mention one that uses an AI program. This has not proved true in Catch-Up, suggesting that RD vs. AI poses a real challenge to AI for three reasons: (1) AI may not be able to surmise that its opponent is RD; (2) what pattern in its play it does infer may be chimerical; and (3) even if AI is able to infer that RD is making random choices and that backward induction would be optimal, it may not be feasible, except toward the end of a game when most numbers have already been chosen.

But this would not be an AI strategy, which has proved so effective in Chess and Go (Tomasev et al., 2020; McGrath et al., 2021). It would rather be to use an idea from over 100 years ago, backward induction, which has nothing to do with learning either from experience or observing the results of playing a game many, many times, as has been done for Chess and Go. Inferring what moves are optimal at each position of the players, based on this experience or observation, is impossible if RD gives no clue of what it will do next.

5. Conclusion

Catch-Up—and perhaps other behavior that reflects random processes in the natural world or applies them in the play of certain games—may be beyond the reach of AI. What distinguishes Catch-Up from Chess and Go, wherein optimal play depends on making "good" moves in response to the nonrandom choices of an opponent, is that RD's choices provide no basis for AI's making good moves.

The only good moves that RD affords AI is to consider all possible moves at every stage and choose the best available, whatever RD does, attributing no intelligence to RD. But this is no more than a brute-force backward-induction calculation, obviating AI's ability to discern patterns in an opponent's play and exploit this information.

The fact that a player who is behind its opponent on any round can use its turn to catch up in Catch-Up means that no player can ever fall too far behind, as can happen in Chess or Go. The fact that Catch-Up keeps competition close, even if backward induction indicates at an earlier point that one player has an unassailable advantage, also distinguishes Catch-Up from Chess or Go.

But it is the desultoriness of RD that prevents AI from exploiting its choices, except via the (non-intelligent) use of backward induction. By contrast, random choices in Chess or Go are eminently exploitable, singling out Catch-Up as a nontrivial strategic game that, via RD, can stymie AI.

I know of no other competitive games in which random choices are the Achilles heel of an AI program meant to exploit a human or nonhuman opponent. This is not meant to deprecate AI or deep learning as a means for a player to defeat an opponent who is not able to think so deeply. Rather, it is to show that thoughtless random choices, like optimal mixed strategies in game theory, can make it impossible for AI to gain an upper hand in certain games.

References

Brams, S.J.; Ismail, M.S. Making the Rules of Sports Fairer. *SIAM Review* **2018a**, *60*, 181-202.

Brams, S.J., Ismail, M.S.; Kilgour, D.M.; Stromquist, W. A Rule That Makes Service Sports More Competitive. *American Mathematical Monthly* **2018b**, *125*, 771-796.

Isaksen, A.; Brams, S.J.; Ismail, M.; Nealen, A. A Game in Which the Lead Alternates. *Game and Puzzle Design* **2015**, *1*, 38-49.

McGrath, T.; Kapishnikov, A.; Tomasev, N.; Pearce, A.; Hassabis, D.; Kim, B.; Paquet, U.; Kramnik, V. Acquisition of Chess Knowledge in AlphaZero. Unpublished paper, **2021**. https://arxiv.org/pdf/2111.09259.pdf

Schaeffer, J.; Burch, N.; Bjornsson, Y.B.; Müller, M.; Lake, R.; Lu, P.; Sutphen, S. Checkers is Solved. *Science* **2007**, *317*, 1518-1522.

Schwalbe, U.; Walker, P. Zermelo and the Early History of Game Theory. *Games and Economic Behavior* **2001**, *34*, 123-137.

Tomasev, N.; Paquet, U.; Hassabis, D.; Kramnik, V. Assessing Game Balance with AlphaZero: Exploring Alternative Rule Sets in Chess. Unpublished paper, **2020** https://arxiv.org/pdf/2009.04374.pdf

Zermelo, E. Uber eine Anwendung de Mengenlehre auf die Theorie des Schachspiels." *Proceedings of the Fifth International Congress of Mathematicians* **1913**, *2*, 501-504.