Episode 08: Go Search, Young Game Player!

Show Notes

Games can be thought of as solving one big problem: How do I win? Cognitive psychologists think about problem solving in terms of a search process, how to get from your current state to a goal state. I discuss what this means for how we play games.

Game References

Assassin's Creed, Breath of the Wild, Chess, Dungeons and Dragons, Europa Universalis IV, Go, Just Cause, Legend of the Five Rings, Magic: The Gathering, Othello, Pac Man, Ratchet and Clank, Space Invaders, Tic-tac-toe, Ticket to Ride, Twilight Imperium,

Research References


Newell, A., & Simon, H. A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

Smith, Q. (September, 2017). Review: Legend of the five rings, https://www.shutupandsitdown.com/review-legend-of-the-five-rings/


Transcript

Hello! This is Episode 8 of the Cognitive Gamer podcast. I am your host, Steve Blessing. Apologies for the long time since the last episode, but Hurricane Irma intervened. I'm here in the Tampa, FL area, and while we ended up not being impacted by the hurricane that much, it wasn't clear that was going to be the case. Between prepping for the storm and then de-prepping, it took a toll on all of us. We just had some limbs and other debris in the yard to clean up, and we were without internet for five days. But, we are back up and going, and our hearts ache for those who did not fare as well, either with Hurricane Irma or Maria.

In this episode of Cognitive Gamer we will be talking about a topic that initially got me into cognitive psychology, problem solving. When I was an undergraduate at the University of Illinois, I started thinking I wanted to do something in computer science and artificial intelligence. My plan was to do this five-year program where I would get two bachelor degrees, one in engineering as a computer scientist and then another in liberal arts as a psychologist. I ended up being much more interested in the psychology side of things, and so ended up just pursuing that aspect. How humans solve problems was the topic that really got me interested in cognitive psychology at that point, and propelled me into graduate school.

If you are listening to this podcast, I imagine you already have some understanding of one of the big things we'll be talking about today. In discussing games, we often talk about how complex the game is, and in doing so often mention how many different states the game can be in. Indeed, in an earlier podcast I talked about this, when talking about computer programs that play chess and go. If you remember, chess has like 10 to the 120th power possible board positions, and the game of Go has like 10 to the 360th power possible board positions. This is a common way to talk about game complexity, to think about how many different states a game can be in. And then, to think about playing the game, you need to think about how you navigate between these states in order to get from the initial board position to one of the win states.

Well, cognitive psychologists take this similar notion of game players navigating game states and apply it more generally to how people solve problems. To be clear, cognitive psychologists also have a very general definition of what a problem is. Any time your current state doesn't match your goal state, where you want to be, you have a problem to solve. So, when you get up in the morning and you are in your pajamas, and you want to be well-dressed for the day, that's a problem. You are in algebra class, and on your worksheet is 2x + 3 = 7, and you want to be in a state where you have x = something, that's a problem. You are working on a presentation for work, and you just opened a blank powerpoint file, and want to have an attractive presentation by the end of the day, that's a problem. Finally, maybe you're a policy maker and you want to solve world hunger, that's a problem.

This definition of problem also works in other types of games beyond just chess and go, both board games and video games. You just started a Magic: The Gathering Game? You have a problem to solve, how to best play your cards in order to win. Just started controlling Link on the Great Plateau in Breath of Wild, and trying to figure out what to do to kill Gannon? That's a problem.

Before continuing, now is a good time to talk about some ways that cognitive psychologists describe these very different types of problems. One set of terms that are sometimes used are knowledge lean versus knowledge rich problems. Not surprisingly, this refers to how much knowledge is required in order to navigate through the problem. Simple games and problems that do not require much knowledge in order to solve them are referred to as knowledge lean problems. Some games are advertised as taking only minutes to learn or somesuch. The go-like game Othello I believe is advertised as taking "a minute to learn but a lifetime to master." It's a knowledge-lean game and problem; you just need to know how to capture pieces by laying down your markers. Other games like Twilight Imperium have a very long rules explanation. Coin-op video games like Pac Man and Space Invaders needed to be very knowledge lean so that people felt comfortable plunking in quarters, but home games can require more knowledge of the player. So, we can have knowledge lean and knowledge rich problems, and this acts not as a dichotomy but as a continuum; Pac Man on one side, Europa Universalis IV on the other, and Ratchet and Clank someplace in between.

Another useful distinction to be made between problems is how well-defined the problem is. For some problems it's very clear what exactly the problem is, how to make moves between problem states, and it's clear once you have found a solution. These are well-defined problems. Backgammon and tic-tac-toe would be in this camp. Other problems are more ill-defined. You mostly know what the problem is, maybe you know how to move between problem states, maybe, and hopefully you know when the problem is solved, but perhaps not. A role-playing game like Dungeons and Dragons might be the great example here. Like knowledge lean versus knowledge rich, this is a continuum as well, though most games tend to be more well-defined than not. Using the non-game examples from before, solving an algebra problem is pretty well-defined, coming up with your work presentation much more ill-defined, and perhaps figuring out what to wear is someplace in between. If tic-tac-toe is well-defined, and role-playing games are ill defined, then maybe an open-world game like Assassins Creed or Just Cause is someplace in between.

We have knowledge lean vs. knowledge rich problems, and we have well-defined problems and ill-defined problems. Any problem, given a cognitive psychologist's general definition of what a problem is falls someplace on this two-dimensional axis. Regardless of where it falls on those axes though, we can now consider another great way to think about problem solving. Again, cognitive psychologists are using a lot of how game designers may think about playing games. In all of these games and problems that we've talked about so far, you have the initial state and the goal state, and then you also have all the states in between; the intermediate states. There are ways to transition between all of these states; the moves of the game, if you will. All of these states and transitions taken together, a cognitive psychologist would refer to as a problem space. Some problem spaces are on the smaller side, like tic tac toe, and some are quite large, like chess and go. The key insight is that any problem can be represented by its problem space, even ill-defined problems. Sure it's hard to conceptualize, and might be essentially infinite, but even something like a Dungeons and Dragons game could be represented by a problem space, the player's initial state, where they could possibly end up, and all the states in between. Not only does this notion of problem space work for problems that are games, but also for all the problems I mentioned previously, like deciding what to wear. Again, depending on the size of your wardrobe, that might be a very large space, but you could imagine the start state, you in your pajamas, and the end states, you wearing something appropriate to leave your house, and all the intermediate states of you being in various states of undress.

So, now we have this notion of problem spaces. What does this gain us, psychologically speaking? It actually gains us something very powerful, and again, if you are used to thinking about it in terms of games, maybe it seems almost too simple or not novel. But, if you have this very general notion of what a problem is, and you think about problems in terms of problem spaces, then thinking about solving problems simplifies to thinking about how we search problem spaces. That is, problem solving is simply search in a problem space—what traversal of states will get you from start state to end state, and *that*, problem solving as search, becomes a much more tractable problem to think about and study.

There are two main ways to consider how a problem solver, or game player, may think about getting from the initial state to the goal state. One is by using an algorithm, another is by using a heuristic. I imagine you already know how humans tend to navigate problem spaces, by use of heuristics, but for the sake of argument, let's consider algorithms as well. Indeed, game players may tend to look at things more algorithmically, by their nature. Algorithms are nice, because they ensure problem solving success. They guarantee success because they ensure that the problem solver visits each state in the problem space. And, if each state is visited, one of them will be a solution state. Yay! But, humans tend not to be so methodical in solving problems, particularly those they encounter in everyday life. Imagine solving what to wear for the day by using an algorithm. You would be trying on all the different combinations of your clothing pieces until you found an appropriate one for the day. Again, though, we probably are more algorithmic in our approach to games, trying to consider our moves more methodically than we would figuring out what to wear for the day or how to do our presentation for work.

Using an heuristic to search the problem space in order to solve a problem or play a game tends to be a more efficient approach, and the one that humans use most often. They are more efficient because heuristics attempt to hone in on a problem solution as quickly as possible, finding a relatively short path from start state to end state, perhaps by simply not considering certain parts of the problem space. Because they probably don't consider all parts of the problem space, they are fallible. Use of a heuristic does not guarantee a solution; indeed, it doesn't even guarantee the most efficient solution. On average, though, using heuristics usually offers a good route for finding an adequate solution to a problem, or figuring out the best move to make next in a game. In that way, it's like our discussion on satisficing from a couple of podcasts ago.

There are a number of different heuristics that people use to navigate problem spaces. To be clear, people of course don't have to be consciously aware they are using a particular heuristic, but by observing how they solve a problem, you can identify a number of prototypical approaches. I'll mention two heuristics quickly, and then go into a bit more depth for two more heuristics.

One heuristic that I'm sure we've all succumbed to is trial and error, or guess and checking: Pick something at random, something that looks good and promising, and see if it works. Cognitively that's not all that exciting, but it does have its uses when you have no other idea how to proceed. Another process that I would classify as a heuristic, and is very interesting from a cognitive point of view is analogy. Indeed, that will be its own future podcast, so I won't say much about it now. But, trying to solve a current problem by making reference back to a previously solved problem, game, or puzzle is a powerful way to solve problems, and has a heuristic kind of feel to it. If you pick a good past problem to analogize with, it can really help you navigate the problem space of your current problem. Again though, that's a topic for another podcast.

The two problem solving heuristics that I would like to spend a bit more time on are hill-climbing and means-ends analysis. First hill-climbing. You see people do hill climbing a lot in playing games and solving problems. To follow the hill climbing heuristic, all you have to do is to always choose the step that will get you closer to the goal. This sounds like it should always work, but it doesn't. To solve some problems, you sometimes have to take a step or two backwards in order to continue forwards. To use a computer science-y term, you have reached a local maxima, in which you haven't reached a solution, but any move to another state will take you further away from an end state. But in general, this is a great strategy, and works for many problems, puzzles, and games. For example, you're playing Ticket to Ride. In general, you want to always be laying down trains that complete your given paths. As long as you do that, you are getting closer to winning the game Of course, another player may be laying down more trains for higher scoring routes, but for you personally, you are pursuing a hill climbing strategy as you pursue that strategy. You are always progressing that scoring marker higher and higher, hopefully to victory.

For some games, though, perhaps particularly strategy war games, sometimes the best move is to seemingly lose ground in order to score yet more points later on. I was recently reading a review of Legend of the Five Rings by Quintin Smith on Shut Up and Sit Down. In the review he recounted a play he did where he, quote, “played a sequence of cards that sacrificed half of my warriors to inflict still-greater losses on my foe.” End quote. The hill climbing heuristic doesn’t account for that type of play, where you take a step backwards to move 2 or more steps forward. A more sophisticated strategy, means-ends analysis, can account for this type of behavior that people will on occasion do. In means-ends analysis, the problem solver considers their current position and their desired position. They select a move that reduces that difference the most. If they can immediately do that move, great. Otherwise, the problem solver now must set a subgoal to solve this mini-problem, for how to apply this larger operation. It is in solving these subgoals that players may need to take a figurative step backward in order to implement the operation that reduces the difference the most, moving them forward in the problem solving process.

I use an example in the classroom to illustrate the difference between hill climbing and means ends analysis. I tell the students I have to get to the airport to go to a conference; that’s my problem, I’m currently here, but I need to be at the airport gate in order to catch my flight. I have all the students point in the direction of the airport. I then implement a hill climbing strategy, where I literally only take steps that take me closer to the airport; I cannot take steps away from the aiport. After walking over a desk or two, I end up nose against a wall, unable to take any more steps towards the airport. Hillclimbing might not be the best approach here. We then consider means ends analysis. It’s a few miles between campus and the airport. I ask them what would reduce this difference the most. Most of them say a car, and that’s great. I then ask, can Iimplement that strategy straight away, standing here in the classroom? No, my car is not here. So, I now need to solve this subgoal, of how to get to my car, and that’s what is going to allow me to take a step back, to so speak, in order to work on my main problem of getting to the aiport. And, just like I’ll have to satisfy this subgoal at the outset, there will be subgoals for me to solve at the airport, because I can’t drive my car directly to the gate. Some combination of hillclimbing and means ends analysis gives us a couple of great, all-purpose approaches to solving a variety of different problems, puzzles, and of course, playing games.

I would like to leave you with two thoughts. First, is actually a puzzle that you may already be familiar with, but it shows this tension between hillclimbing and means ends analysis. It's the orcs and hobbits problem, also known as the missionaries and cannibals puzzle. Imagine there are 3 orcs and 3 hobbits travelling together, for some reason or another. They come across a river they have to cross. They have a boat that can hold 2 creatures at a time. If at any point in their transition from one side of the river to the other, there are more orcs than hobbits, the orcs will of course eat the hobbits. How can you move all 6 creatures from the one side of the river to the other? If you are not familiar with this puzzle, it can be a little challenging, because at one point you have to violate hillclimbing in order to solve it. The second thought to leave you with follows from this observation from the orc and hobbit problem. What makes a game, puzzle, or problem hard is our ability to navigate its problem space. That can be related to the size of the problem space, but that’s not the whole issue. The hobbit and orc problem space is relatively small, only 16 states, but many people find it challenging. Our ability to apply our problem solving heuristics also a major determiner to problem and game playing difficulty, and that’s affected by a large number of factors, like how well- or ill-defined the problem is. Also, the context the problem is placed in can greatly affect problem difficulty, but that will be a topic for another podcast.

I hope this discussion of problem spaces has given you a deeper appreciation of how you go about solving problems and playing games. Cognitive scientists have really taken this notion of a problem space, and have used it to deeply figure out how people approach those processes. It has led to some real insights into how the mind works, like memory and subgoaling. Next time, we have a real treat here on Cognitive Gamer, as I will interview someone about the role games can play in childhood development, and how certain cognitive milestones can be measured by the playing of games. As always, I welcome any comments or questions you may have, so please email me, steve@cognitivegamer.com and also visit my website, cognitivegamer.com. Also, you can like me on Facebook, Cognitive Gamer, or follow me on Twitter, @cognitive_gamer. Until next time, remember to think about what you play, and have fun doing it.