Alphazero on Go

Daniel Hu 5d

29 March 2019

Abstract

Over the past three years, the go world has been revolutionised with the advent of computer programs that can beat top professional level humans at the game of go. In this paper, we look at some of the ways in which they have influenced our thinking and taught us new ideas. I provide an analysis concept by concept of the most important and critical ideas that artificial intelligence (AI) has brought. Watch out for these generic ideas when they appear in your own games and those that you are analysing. This paper is for anyone who plays go but doesn't necessarily understand AI very well, and also for AI researchers who want to understand the impact and progress of AI in go. There are two sections, one on go and one on AI.

DISCLAIMER: The views expressed below are just my personal opinion as a high dan (very strong) amateur go player from watching AI play go at superhuman level. Unfortunately these are not necessarily valid as the AI can't speak for itself (yet). On the other hand, I do believe that learning the **principles** demonstrated by AI (or strong players in general) can raise your go strength and consistency dramatically.

Joseki are standard corner sequences that appear in professional play. Generations of joseki have been overturned by recent AI. Their joseki are more direct and territorial, and they are very active in taking sente. There is a lot that you can learn from watching and studying their play and asking **why** they make one move and not another and comparing them. For example, they are good at invading a structure at the perfect moment, and better at judging the right depth and timing.

Alphazero, developed by Google, is the strongest known AI at go and several other board games, and its design serves as the template for the other top AI such as leelazero[2]. Papers and games can be found at the link in the references section[5]. Alphago and Alphago zero, also developed by Google are related programs and earlier versions. I call them all alphazero for convenience. These are all following a neural network design, modelled on theories in neuroscience. In go, they lead to robots that turn low level number crunching into the emergent phenomena of astonishing intuition and creativity.

1 Go concepts

Efficiency per move Alphazero emphasises efficiency per move very greatly, and nowhere is it more obvious than when it takes gote powerfully, or ambitiously takes sente. The Japanese words gote and sente respectively refer to a result that loses the initiative, and to a result where you can take the initiative to play elsewhere. Taking gote can be treated as spending an extra move locally compared to sente. Therefore you should expect a much better result taking gote than when taking sente. But it is a matter of degree as to what difference you should aim for in real play. In the opening, the difference should be bigger than the endgame. And the difference should go up when fighting over weak groups occurs since every move matters more.

I find that in general alphazero treats the value of initiative more seriously and noticeably than human players have done. When alphazero takes gote in a local sequence, it invariably ends with a great result in territory and thickness locally. And if there is the slightest sense that it doesn't have to respond to a move and the local "temperature" is low enough, then it takes sente, looking for greener pastures. See figure 1. This can be seen by related AI judgements on old joseki in the alphago online teaching tool[1, 3], see figure 2. I suppose this is simply a show of its strength, and there is a lot we can learn from watching how it makes judgement calls of where the ideal balance lies in different spectrums:

- security vs initiative
- defense vs attack or defense with attack: you can try to achieve more things with one move as each move is very valuable

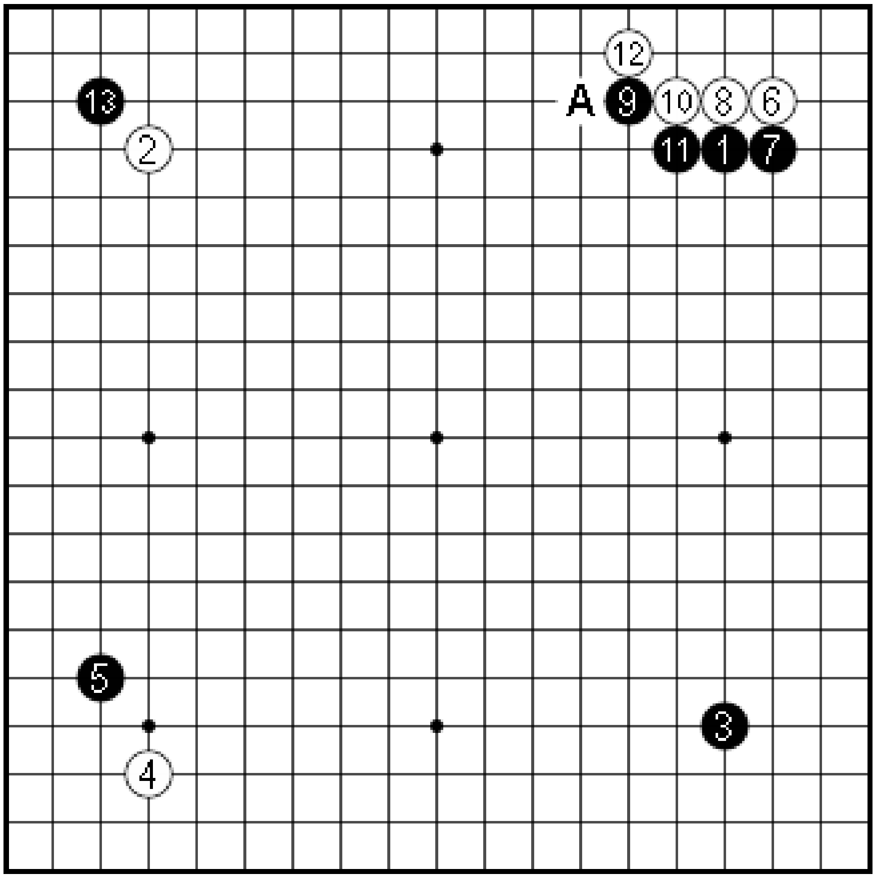

Figure 1: Alphago self-play game 3

Black takes sente from the upper right when human players would normally continue defending the shape by extending at A. It uses the initiative to aggressively invade another corner at move 13. This means it judges that extending wouldn't have satisfactory efficiency.

So white's invasion of the upper right corner may be gote but it is a very high efficiency gote, especially if white has the opportunity to spend another move. All white's stones are working well for territory and eyespace, whereas black's stones aren't making territory (yet), and are too close to thickness and hence weak.

- territory vs influence vs eyespace and thickness: alphago has a very different concept of thickness especially with regard to walls
- reducing vs building vs invading
- light flexibility vs thick commitment

Efficiency is a concept that can be extended through all parts of go, and developed into complicated ideas, irrespective of AI. This is because every move that you have is very valuable as it comes at the cost of conceding your turn for an opportunity for your opponent to make a move to make a plan or counteract your plans. And the difference between a stronger player compared to a weaker player can be defined as someone who spends their moves better. Hence go is reduced to a question of getting good value for money (or good work for energy cost to a physicist). Different moves can work in harmony to build territory or capture enemy stones, and they can also be discordant, such as when you shorten your own liberties in a semeai (capturing race). There is a natural implication of a definition of local **temperature** that refers to the value/opportunity cost of making a move compared to letting your opponent move in terms of the fundamental units of the difference in score (winning margin) after the game ends. The deep mathematical study of such things is called combinatorial game theory, and was developed by John Conway in the 1950s.

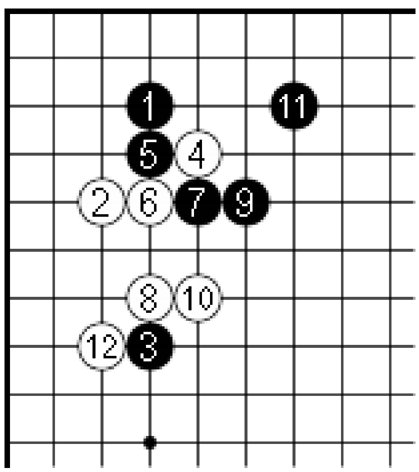

Figure 2: Classic 3-4 point joseki

Move 8 used to be considered the only move as an excellent and vital tesuji to dodge fighting. But it loses 8% in winning chance according to the alphago online teaching tool compared to directly fighting[1, 6]. I suppose alphago views it as slow and inefficient.

This whole joseki up to 12 used to be judged as even for both sides. This is firstly because black has a better position with slightly more territory. The shapes are very similar (imagine moving stone 3 up two spaces and stone 8 diagonally). But as black took the corner first (with move 1), it was considered that this sequence is playing in blacks “sphere of influence” and so black should be expected to get a better result.

The result is now considered much better for black. This is due to an efficiency analysis. We can justify making an efficiency analysis since both sides have settled (the next follow up is worth much less than the big moves that have been played in this sequence). Both sides have spent an equal number of moves here and white was the last player to take gote. There are generally a lot of big points in the opening, so white should expect a better result than even each time white takes gote. But instead, the situation is worse for white, so the result is better for black.

Another way to see this is in the spirit of move 2. Move 2 was played deep and close to the corner in order to pressure (gently attack) it, so white should be able to take advantage of black’s weaknesses or else move 2 isn’t a good move. In fact, AI thinks that move 2 is a good move, hence white must have played something wrong between move 2 and move 12.

Attack It is said that it is difficult to attack in go, as it takes four moves to capture a single enemy stone, and there is normally always room to escape, or at least a way to sacrifice and exchange. And there is the principle that attacking moves often become scattered disconnected groups that don't make territory, and hence become worthless and hence raise the efficiency of the enemy stones. Hence most forcing moves are bad moves that are worth less than they cost. And hence most stronger players tend to prefer to take territory (or big points) directly or threaten and pressure enemy stones from afar rather than directly attack and harass. And in my opinion alphazero goes even further down this route, reiterating all these points and more.

In a sense, moves that are the opposite of forcing moves, those that give the opponent a free choice of many equally valuable moves are often very good moves. At the least I think there is the general conservative principle that such moves are rarely the worst move, even if they aren't best. Remember the opponent always has many options to choose from, and these can be split into different areas of the board and different "directions" of play to choose from. Different moves you make change the value of different areas, whether to increase or decrease. Normally the opponent only has a few good moves to choose from, so if you successfully balance out the board to give them more options, that makes it difficult for your opponent and hence often good for you.

Alphazero tends to like attacking early on only when it is simple and a direct way to take profit. So forcing moves are only played if they are rapidly exchangeable into hard cash. In my opinion if this is a guiding principle of AI play, then it may miss the more subtle or ambiguous or sharp and complicated ideas in attack and defense that come with long term planning. But it is exactly these ideas that tend to lie on a knife edge and the success or failure of these ideas win or lose a human player a game. But I think the deep neural networks remain much better than humans at even this long term strategic thinking. And nowhere is this clearer than in alphazero's performance in chess, where it has superhuman play on zugzwang and light and dark square control, for example highly valuing a space advantage (like a boa constrictor) and also the combination of the two bishops. And yet they still must always have blind spots as they too are finite, and even they might be worse at some concepts. For example, AI are famously bad at reading long logical sequences, and this is particularly notable in the case of playing out ladder variations.

Areas, thickness, direction and shape, and hence sente Follow on from the above, I think there is a general human mindset for how we find good moves logically. First we ask what the right area of the board to play on is, then what groups are relatively thick (so we can apply the proverb play away from thickness). Next is what good direction you should play for your plan (building, reducing, invading or killing). Finally we consider in the local shape, what stones make a good relation and connection to other stones. But we don't know if AI have similar notions hidden in all their numbers. Or even if there are any such notions in an "ultimate" theory of go. Such notions of plans, logic, comparison, classification and substitution of moves may only be emergent properties from the rules, assuming they are even justified.

If anything, since computers are less emotional and feeling-based with much more calculation, perhaps they are ironically more logical, even when logic hasn't been coded in. Their numbers use a value network as a heuristic for finding good moves. Even if a single function heuristic cannot possibly capture the complexity and depth of go, it is a good guide to help them find good moves. And then they recursively search for good moves for the opponent (to find the consequences and risk of each move) and then select the best move in that context. Given sufficient time, it is a system that converges to finding the "best move" and perfect play. But the hope is that it may also converge to a more meaningful and useful heuristic of winning probability for real world play.

Perhaps we can get better at finding good moves by putting us in the shoes of alphazero and seeing what it sees and what it misses. Firstly, we can consider what the opponent's best few moves are if we pass. From this, we can try to play moves that make the opponent's set of good moves less good for them, or our own follow ups better. Perhaps this is an explanation for why alphazero seems to like to keep the initiative. It is probably because in this context, a sente sequence is likely to be one of the better sequences on the board simply because they stop the opponent from playing what they want to play. It isn't clear to me whether perfect play will have as many sente sequences that seem aji-keshi. Sente sequences may settle the position, so the only thing we have to be careful about is if these moves remove our other options, and hence leave better moves for the opponent that weren't originally good moves. After all just one move can refute a clever idea.

Let's go back to area, direction and shape. These can be considered different levels or hierarchies of go theory or strategy. There are generally only a few big areas of the board, and these depends on the spacing and "shape" between stones and groups. Choosing the right area of the board, and not spending gote in a small area, tends to mean you will be less criticised by a teacher. I think human go teachers tend to view playing good direction as a key indicator of strength. Good direction tends to be mutual, that is, good for both your opponent and you, so that is why putting yourself in the opponent's shoes is a useful tool. At other times, you have a free choice of direction, and your choice indicates what your plan is and should be, and what your style is, whether territory or influence based, thick attacking or thin shinogi and so on. On the other hand, playing good shape tends to be considered a mix of reading and experience. There are principles such as "your opponent's vital point is your own", but these are just as often false as true. Good shape recurs often so recognising similar positions helps train your shape sense. This is a reason why neural networks work in go and why learning is possible: no two positions may be exactly the same, but many similar ideas remain good or bad. On the other hand, what seems good or bad shape no longer seems so as you get stronger, and hence deeper reading is obviously critical.

This seems to be a very logical system, but it is hard to tell if such a system of understanding and logic is reflected in the reality of go, or if they are more a feature of how our minds work. But I suppose this is a recurring question in the divide between the sciences and the humanities.

Ideals and perfection Alphazero tends to like sequences that end with it playing a perfect move, that is one that has perfect efficiency. This is a bit like the proverb that says you should delay invading a moyo until the last moment. This proverb probably works because any group shares a lowest common denominator of two eyespaces and hence breaks a lot of territory. Hence the difference between two alive groups in terms of territory tends to be less than the difference between an alive group and no group. Another example of one stone having perfect efficiency might be when the opponent has two walls, but you have sente to play in the middle with just enough space to settle (i.e. live or get something of equal value), and reduce the value of the walls dramatically. This is because they can't be used to efficiently attack any group (by definition) and can't be used to make territory when you have a settled group in the way.

In my opinion, there is a lot of evidence of AI playing local sequences that end with perfection, and I have a strong suspicion that alphazero can judge this perfection well and actively looks for it when reading out possible sequences to play. (as alphazero is given the data of what the last few (generally seven) moves were and hence can "see" what the plan was) See figure 3.

Perhaps further analysis requires that we break up what exactly we mean by "sequence" and how we define the start and end. So for now, I will just suggest a definition: an end of a "sequence" is when the natural flow and pressure ends and tenuki becomes most attractive, or the shape is settled. Or it is the conclusion to what looks like "plan" in hindsight, and the next few moves will move onto a different strategic/tactical fight.

Thickness of walls and difficult play in the centre Normally when we build a wall we think that it is strong and thick and we are meant to use it to drive the opponent towards as an anvil for a hammer or else turn it into territory. After all building a wall of influence towards the centre tends to be at the cost of giving the opponent solid territory towards the edges and corners. But alphazero has very different views. One is changing what we recognise as thickness and placing a higher value to eyespace and solidity of connections as components of thickness. And another is changing how we play with thickness and increasing the importance of the proverb "don't play near thickness".

See figures 4 and 5 for what counts as thickness. The 3-3 invasion of the 4-4 point used to be considered to give the opponent too much thickness. But now it is considered a form of attack. After what is attack except for cutting off the options for an enemy group, and the invasion undercuts and takes away the 4-4 stone's access to eyespace in the corner.

We used to try to quickly "use" a wall and find it unacceptable if a wall that was made in exchange for territory ended up getting reduced to nothing. But alphazero tends to criticise that driving groups towards walls can sacrifice the strength of the wall for uncertain profit, and trying to make territory off a wall tends to be too hard for too little profit. So instead it treats walls more flexibly as also having value being a **source of initiative** and often leaves it up to the opponent to decide what happens to the wall and is more willing to let the opponent reduce it **at the cost of one move or more**. This means that it exchanges the value of the wall more slowly, perhaps initially in return for a better result elsewhere on the board, and often may still be able to use it to attack **when the timing is right** (and it is worth it), and still make territory from the wall (and not much less than trying to make it directly by oneself). To us, this can seem like it tries to get everything and succeeds, but clearly it just has a better understanding of what is valuable and what to exchange value for.

Reduce loose stones and make tight stones overconcentrated This saying can be considered a variation of old proverbs, and there are many examples of this idea in alphazero's play. It helps guide our play and guide how we analyse and read positions, as they help train us to look where the ideal and efficient follow ups in a sequence may be. The concepts behind the wording are quite subtle and difficult, so I will just explain them and then focus on a few concrete examples. Loose stones have many interpretations including high efficiency stones, such as those which try to achieve multiple things, or those on the 4th line. It is valuable to play near such stones in order to counteract them and exploit their weaknesses.

Tight stones might mean those close to other stones, or close to the edge, such as those on the 3rd line. It is less important to play near such stones as they are already more solid and you have to be flexible to take what you can get when playing near such stones. You have opportunity to profit in either territory, influence, initiativeor attack/defence, depending on how the opponent responds. Hence it is often a good idea to play light probing moves near such stones to see what the opponent commits to. You can start by threatening to surround or take away territory or strength from such stones, and you might

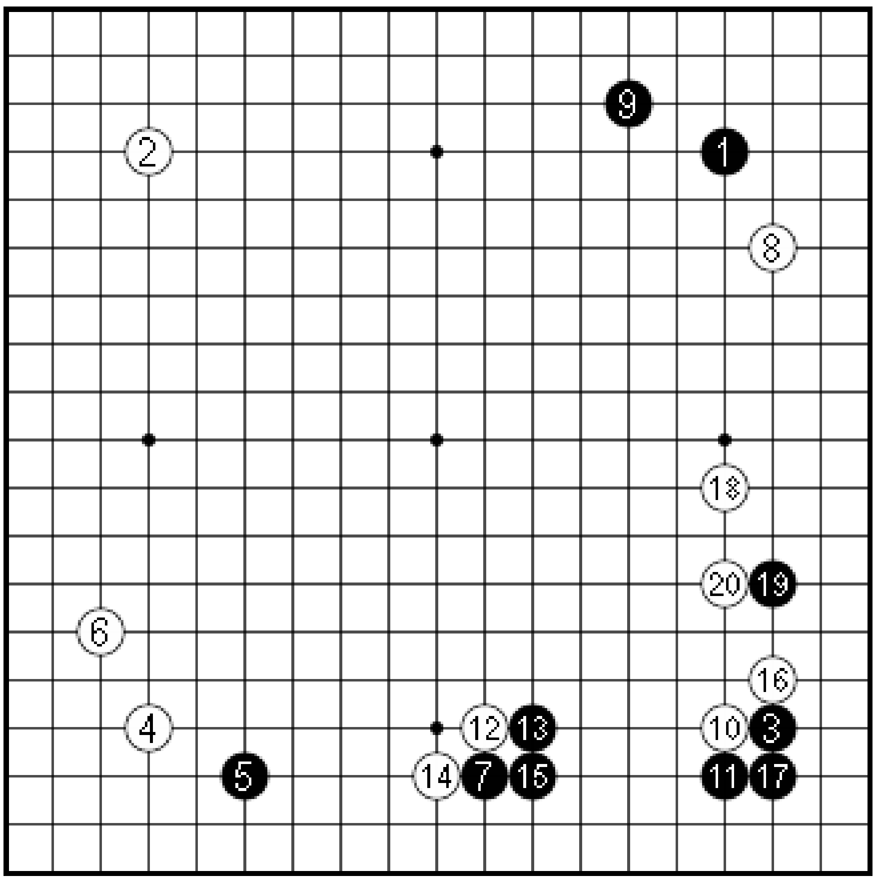

Figure 3: Alphago self-play game 1

The moves 10-16 may be considered a sequence of sente probes by white. These lead to move 18 taking an ideal and perfect point developing the stone in the upper right corner and helping the weak stones in the lower right. Furthermore, the efficiency is greater than a normal play just above the star point if the invasion isn't a worry. The probes also work well in the game when black approaches at move 19. Moves 10 and 16 may help white more efficiently fight with move 20. This makes move 19 look like invading in white's sphere of influence.

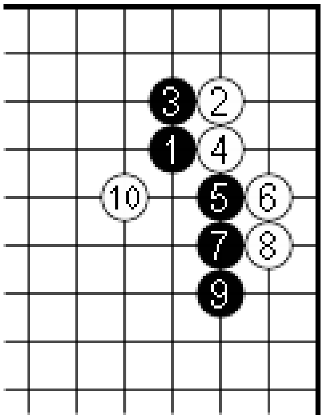

Figure 4: 4-4 point 3-3 invasion joseki

This joseki (1-9) used to be seen as too good for black as black gets a large thick wall early on in the opening and white's territory is too small. However alphazero turns this judgement on its head, judging that white living in the corner is sente is good for white. White's corner is thick as it is alive. But black's outside is weakened as it is too close to white's thickness, and the wall is a lot less thick than it appears. Since we have a relatively new proverb that "any group which doesn't have two eyes is not thick". Hence while all white's moves are useful and necessary, black's moves lead to an inefficient result for having spent an extra move. Also, alphazero taught us a new move at move 10 to follow up. Peeping the cutting point in black's wall puts black into bad shape and may serve as the preparation for an attack on the whole black group. Alphazero is good at measuring small subtle weaknesses such as this and endeavours to give the opponent cutting points that can be exploited later.

profit in gote, or gain guidance as to direction of play in sente. One example is pressing low stones even lower and offering the choice to the opponent to either risk a fight or take a smaller territory.

See figures 6, 7, 8.

Sacrifices, exchanges and kos The mark of a strong player is how well they judge the value of exchanges (by definition), and hence the ability to play very complicated kos that require the combination of knowledge in all parts of the game tends to be a marker of how strong a player is. They require reading out long complicated sequences combined with the skill of accurately counting and judging differences in value. Often exchanges can be very large and difficult for humans to count, but I think AI have an unfair advantage in this area as they are built to be good with large numbers. Whereas a human brain finds it very hard to deal with and understand.

It used to be suspected that AI were bad at such complicated ko judgements, but it seems that alphazero finds it comparatively easy compared to humans and has many new judgements to offer. They particularly revolutionise the opening theory and overturn many classic josekis. See alphago self-play games, especially game 10 of 50.

2 AI ideas

META So where do we go from here? Is it the fate of humans to be worse than computers at every game and skill? In my opinion, given how brilliantly go was cracked despite how deep and beautiful it was, AI will be able to dominate in almost any skill it chooses, but I think all hope is not lost. I think there are always meta-concepts beyond the programming of a computer. While it may learn to overcome any single one, or utilise them when a more advanced programmer changes the programming, there will always be more blind spots and weaknesses.

For example, one possibility that alphazero is not directly trained for is time management. It could of course be easily trained by tournaments as usual to train a function that decides how much time to use. But I think this is an arena where there is no "right" answer, and may confuse neural network bots. I think such training may not necessarily lead to good results that **generalise** to new situations. But also that as the programmers, we are fundamentally flawed in our foresight, and there will always be a "meta" concept that may be critical to a bot's performance, such as my example of time management, but

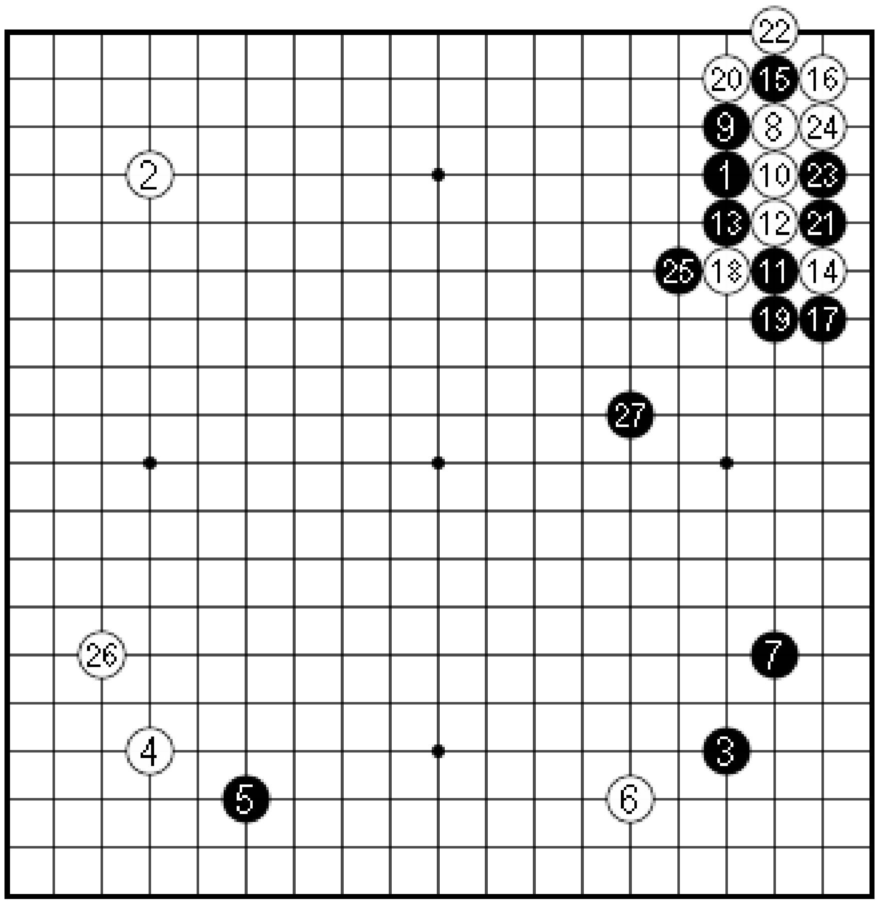

Figure 5: Don't play near thickness

The upper right is a new alphazero joseki for the 3-3 invasion in to the 4-4 star point. Of note is move 27. Before, we were told by professionals to capture a laddered stone such as 18 as a thick ponnuki (taking it off the board) to prevent bad aji and profitable forcing moves later for the opponent. But alphazero fights with this move and says that just capturing is too passive and that black needs to ask for more. This far away move is an attempt to maximise efficiency, and perhaps also follows the proverb "don't play near thickness". Also of note is how alphazero seems to treat the upper right group as extremely thick (presumably because it is alive) and is happy that move 7 is so far away, making a 6 space extension. It seems to think the spacing is ideal and tries to make everything in between as territory. Again this follows the proverb.

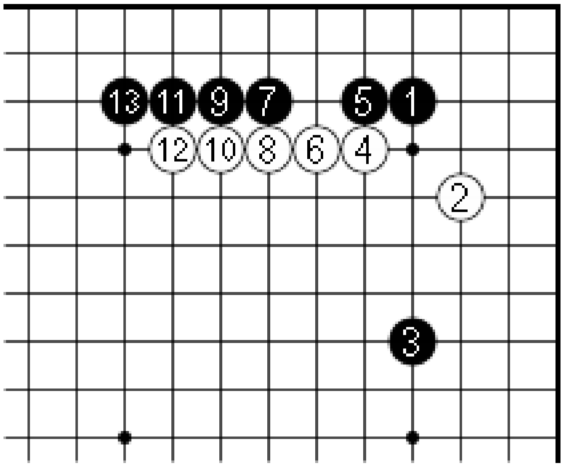

Figure 6: 3-4 low approach press joseki

This is a classic joseki following on from 4; 8-13 are optional. Alphazero likes move 4 extremely much wherever the position the position of move 3's pincer including tenuki (playing away). It also much prefers the low approach of 2 over the high approach, just as human players (such as Shusaku) preferred pre-1960s, thinking that the high approach is slack on territory. The low approach tends to lead to fighting as black has little left in the corner after the approach. It seems that alphazero tends to think of move 4 when playing move 2. Perhaps it is a case of pressing stone 1 down low.

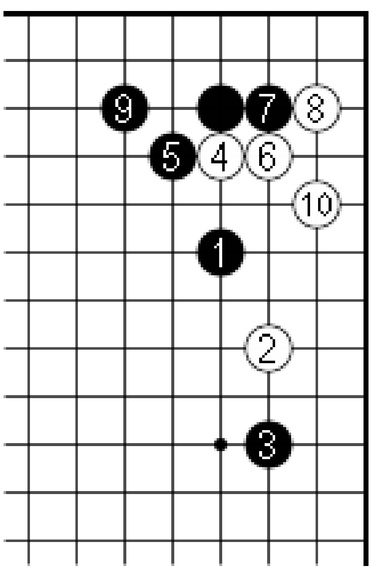

Figure 7: 3-4 large shimari joseki

Alphazero's favourite way to develop from the 3-4 stone is the two space high extension at move 1. Move 2 is a good move to undercut the high stone and threaten a new move and tesuji at move 4 that exploits the large gap between the black stones. The following sequence is a notable new joseki that it invented. See the games of alphago zero vs alphago master.

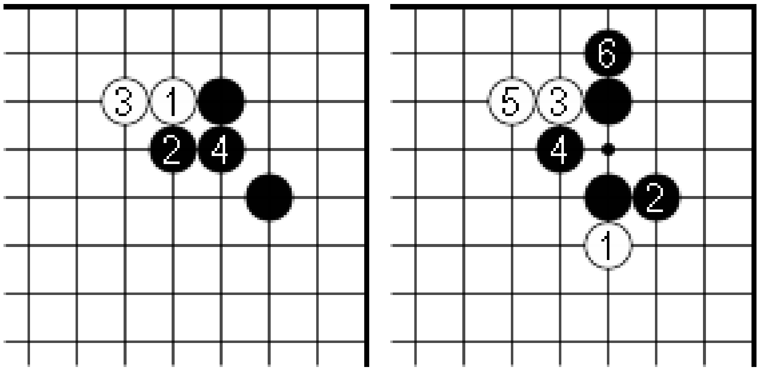

Figure 8: Classic shimaris (enclosures) for the 3-4 point

The classic shimaris for the 3-4 point are the knight's move and the one space jump. However, alphazero likes to make early attachments to probe and make these small enclosures even more over-concentrated in only taking the corner territory. This seems to be why alphazero prefers larger shimaris.
These exchanges used to be considered ajikeshi as they make black's territory more solid for uncertain profit, so it is interesting that alphazero judges the timing of these exchanges to be the earlier the better. Presumably if black gets to add another move, these exchanges won't be good exchanges, so white must immediately help the light stones. Once white is reinforced, black's local efficiency becomes fairly low.

that haven't been trained for. However, truly pondering this problem gets us into deep realms of metaphysics and maybe Zen, so I will stop here with this.

However, training for time management would also be training a bot to value what positions needs attention. For example, a bot may play better if it pays more attention to dangerous and complicated variations, but such a bot may be beaten by a bot that pays more efficient attention to good variations, as they do now, so it is still a matter of intelligent programming.

Strengths/weaknesses For me it remains to be seen what the strengths and weaknesses of AI are, especially as they change over time. I offer some Hegelian dialectics that hopefully converge on truth.

- One obvious weakness is that they are expensive. They take a lot of time and energy and resources to train and build. The general Alphazero took 5000 special TPUs running for 13 days to train, looking at 700000 steps of batches of 4096 positions or around 9 million games.[5]. In human history, there are only of the order of 100,000 recorded games. And once built, it still takes space to store the configuration and they remain computationally expensive to run. And yet, these requirements appear to be going down over time rapidly too as only a few GPUs and CPUs are required to play the latest versions at high level. The general alphazero examines 60,000 positions per second, far more than a human, but far less than rule based chess engines.
- Another obvious weakness is that they are not systematic. Their training by experience is systematic, but they don't learn from the design of previously well-performing programs, and don't build a systematic "theory" of go. And yet, perhaps they do in a subtle and ambiguous way, hidden beneath all their rows of weights. Perhaps they do have a systematic evolution. But even if that is the case there is no solid way to interpret it and measure its progress (yet) except by the battle-hardened way of playing in tournaments. But ELO ratings exist and can compare strengths so that is not too bad a measure, but still they don't (yet) set down theory for later bots to follow. They don't seem to talk in language. And yet maybe they do since the games they play are used as training for later bots and may count as a "the language of culture".
- They may be fragile as they are training towards perfection. Just one transcription error between weight files may cause the whole thing to be broken. It feels like training an equilateral triangle by iterating each vertex to move as an isosceles triangle. You may "converge" to an equilateral triangle by some well-designed measure. But along the way, many things beyond the measure may be broken or randomly thrown out, such as area or position. And yet there can be further measures which measure the instability of such things.

Onwards I am interested in how we can develop the technology of alphazero in order to better understand go and its place in the universe. Right now, AI takes as input the last 8 board positions with indicators for black and white stones of each colour. It learns to better output two features: policy (probability of each next move being played) and value (win rate) networks.

What I wonder about is how to connect these mathematical concepts to human traditional concepts. One idea would be to try to "look into" the neural network weights and try to figure out what each neuron layer is doing, for example by deconvolutional neural networks, which may study how "meaning" percolates throughout the structure. It is also fairly easy to study what an individual neuron is doing by looking at when it activates and when it doesn't, like in neuroscience. Unfortunately there is no guarantee that an individual neuron represents a "concept" as each concept may encompass the combined output of many neurons, or not be used at all. The other direction would be to start from our concept, and compare it to the neural network's judgement by analysing real games that it plays or inside the network. Unfortunately, our conceptual notions are very fuzzy and seem to be a matter of judgement. They are very hard to define especially in terms of mathematics except for liberties and life and death. A middle road is perhaps to work to define our concepts mathematically (at least sufficiently to get a training set) and then ask the neural networks to learn to recognise our concepts and put a value on each group or each position for each concept, so we can measure them.

Concepts that I think deserve a deeper look into are connection, "good shape", thickness, aji, weak group, attack and defense and territory. This is because they are very important and relatively hard concepts but potentially defineable. I think progress can be made by developing the "language" of how we communicate with neural networks.

I dream that this would lead in the direction of making a bot to solve math problems.

A few disconnected thoughts:

- Maybe just a bot that can be taught is sufficient to make progress e.g. concept, addition, theorem, proof, problem solving
- When we learn maths, we learn by connecting new ideas to old ones that we already have either intuition or algebra for. For example, this could be tracing back the ideal and abstract concept of number to our real world or calculus to physical motion. This could be the equivalent of adding a new concept to a neural network, or adding a new neuron. We should be able to find a measure to determine how easy it is to delete a neuron, and hence see how easy it is to squeeze the channel there, and hence see how "congested" the network is in holding different concepts
- How do we discover new laws of physics? QM is such a mathematical description of reality, is it possible that it is really close to the truth?
- Are neural nets still too human to work out the mysteries of the universe?

Theory of go Continuing on from the above, I wonder if AI could be developed to offer (and teach) us a more concrete and rigorous theory of how to play go well, perhaps by an algorithm of the following sort:

- When the play is not close-combat such as in the opening, play alive moves before less alive moves, big moves before small moves. Note that each point in the centre costs around three moves to surround, while the corners require one. So including the value of the life of the stone played, the centre is about twice as expensive as the corners in the opening. And play urgent moves before less urgent moves, where urgency is relative to how weak groups are, and if there is a mutual vital point between weak groups for space, flexibility or shape. Note that it is much easier to make life around the corners and the edge of the board so follow the proverb "The corners are gold, the sides silver and the centre grass." Perhaps this can be phrased as urgency makes a move bigger than it normally is. And interpreted in the context that big moves may be those that make life more easily or influence life and death of groups more, such as near the corners and edges. This is until a wall is made.
- .When in a fight, or if stuck on how to respond, ask where the opponent wants to play. And the opponent's good point is generally good for you. Of course, this strategy may be iterative to find what moves are good for the opponent and may require reading ahead.
  - If the opponent's good point is too small or passive for you to play, then fight back by disconnecting (cutting) or by threatening an enemy group in order to defend or respond actively.
  - If the opponent's good point was a defensive move, ask what it was defending against, and play there.
  - If you can profit in sente or probe in sente first, then do that before playing your good move.

More lessons TO ADD:

DIAGRAMS

Territory, tenuki, direct fighting, subtlety of leaving cuts, connecting, attacking, reduce vs invasion at last possible moment. Early attachment leaving a follow up and forcing a response. Probes. Tesuji, aji. Efficiency is what we all fight for, so how does it do it?

3-4 makes follow up of approach/shimari more important. 4-4 offers 3 good responses so follow up is less important

Take a position or joseki and play historical vs alphago move

Aiming for a definitive guide to lessons from the master/alphago.

have an accompanying sgf with all the key variations?

could analyse winrate changes with leelazero to look for more weaknesses and more concepts.

Concepts relating to accurate positional judgement are particularly difficult

——

Into evolutionary biology

the weights act like genes that mutate. This is like a model of evolution where the group is powerful. After all, there are the necessary components of initial randomness, inheritability of attributes, mutation (in a coordinated way here, controlled by group, via backpropagation), and competition (between groups, in tournaments if a network beats an the leader 55% of the time, then it becomes the leader).

We can even examine the difference between the theories and models of kin selection and group selection. I understand kin selection is a more popular theory these days to explain nature. But this neural network model uses competition and tournaments to select each network (a group), and the network produces its own training data (like the environment) and whatever the random distribution of weights (individual units of selection), they are trained, selected, and modified by the group for the benefit of the group. There may be a kin selection occurring for self-sacrifice, but also there is an arms race against time for producing training data that can better defeat the group, since the group tends to mostly be competing against the same “species” of high-performing networks with very similar weights.

Overall,maybe this highlights a problem with kin selection in trying to explain altruism: kin make up much of the environment, and hence there may be more competition (in terms of fitness) between kin as they are near than non-kin. So perhaps kin selection leads to much randomness in the gene pool, but perhaps a total order on fitness (like in go) depending on the distribution of the gene pool. Kin selection may actually select for defectors from kin more than cooperators, from a game theory perspective. Perhaps a lot of what we see as altruism is actually the opposite when it comes to kin.

In any case, I find the whole debate fascinating but confusing.

It would also be interesting to study if there are parallel concepts such as junk DNA (rarely activated neurons) and so on.

References

[1] Alphago online teaching tool https://alphagoteach.deepmind.com/

[2] Leelazero, a free open-source version of alphazero from Gian-Carlo Pasciutto https://zero.sjeng.org/home

[3] Classic joseki dictionary http://eidogo.com/

[4] Pattern search program for pro games http://ps.waltheri.net/about/

[5] Alphazero: Papers and games https://deepmind.com/research/alphago/alphazero-resources/, Alphago papers and games https://deepmind.com/research/alphago/

[6] Weiqitv alphago teaching tool https://weiqitv.com/v/1839644