British Go Journal

Special Computer Go insert covering the AlphaGo v Fan Hui match

Fan Hui is on the right, Toby Manning centre
Photo courtesy of Google DeepMind

Number 174 Winter 2015 - 2016 Price £6.00

Introduction

Jon Diamond president@britgo.org

This special insert into the British Go Journal celebrates quite an achievement – the first win by a computer program (AlphaGo produced by the company Google DeepMind) in a match against a Go professional, the current European Champion Fan Hui.

Although the British Go Association wasn't officially involved in this event, our Treasurer, Toby Manning, was appointed as the independent Referee for this match. His report with the games and comments by various of the team involved and Fan Hui, in addition to Korean professional Hajin Lee subsequently, is the main feature of this issue. [These games are also published in SGF format on our website - see http://www.britgo.org/deepmind2016.]

To round out this issue we've put together some background to the history of Go-playing programs.

For those technically minded, there's a peer-reviewed article in the scientific publication Nature, written by Google DeepMind, about the software. The article states that AlphaGo uses Convoluted Neural Networks to suggest moves and Monte Carlo Tree Search playouts to decide on the actual move to make. DeepMind has used millions of games from KGS (and possibly elsewhere), adjusting the game weightings according to the grades of the players, to train the CNN, with little Go knowledge specifically embedded in the program.

The History of Go-Playing Programs

Jon Diamond president@britgo.org

This article has been synthesised from a number of online sources (referenced at the end), with some additions, mostly from my files. It is published online[1], where you can see the earliest program's game.

Go has long been considered a difficult challenge in the field of Artificial Intelligence (AI) and is considerably more difficult to solve than chess. Mathematician I. J. Good wrote in 1965:

Go on a computer? In order to programme a computer to play a reasonable game of Go, rather than merely a legal game, it is necessary to formalise the principles of good strategy, or to design a learning programme. The principles are more qualitative and mysterious than in chess, and depend more on judgment. So I think it will be even more difficult to programme a computer to play a reasonable game of Go than of chess[2].

The first Go program was probably written by Albert Zobrist in 1968 as part of his thesis on pattern recognition. It introduced an Influence function to estimate territory and

[1] http://www.britgo.org/computergo/history

[2] http://www.chilton-computing.org.uk/acl/literature/reports/p019.htm

Zobrist hashing to detect ko. It could just beat a beginner. [There are references to Go playing programs by H Remus (partially complete only in 1962) and D Lefkovitz in 1960 but no more information about them is known.]

Jon Ryder produced a program in 1971, which lost to a novice, so was probably no stronger.

The first computer-computer match was between the programs written by Jon Diamond (Institute of Computer Science, London University) and Jack Davies (University of Cambridge) in 1973 - the game was unfinished and no record of it has been found. Jon's program was probably the first to use the alpha-beta search algorithm and also beat a beginner. In strength it was about 20-25 kyu.

In 1978, Walter Reitman and Bruce Wilcox reported on their Interim.2 program, having started on it in 1972. It beat a 22 kyu player and used lookahead which was not full-board, rather it was a selective, goal-driven process.

Jonathan K Millen published an article in Byte in April 1981 discussing **Wally**, a Go program with a 15x15 board that would fit within the KIM-1 microcomputer's 1K RAM. Bruce F Webster published an article in the magazine in November 1984 discussing a Go program he had written for the Apple Macintosh and included the MacFORTH source.

The first computer tournament that we know of, the *Acornsoft Computer Go Tournament*[3], was held in London in 1984 with the British Go Association as organiser. It was sponsored by Acornsoft, and the programs all used their popular **BBC Micro** microcomputers on 13x13 boards. The name of the winning program is not recorded; its programmer was Bronyslaw Przybyla.

Later that year, the Unix user group Usenix sponsored the first of a series of Computer Go tournaments. You can read about these and many more on the *Computer Go - Past Events*[4] page. This 1984 event was won by Bruce Wilcox's **Nemesis**, which later evolved into the commercial product **Ego**.

The first time a computer competed in a human Go tournament was in the 1980s, **Nemesis** at the Massachusetts Go Club.

In 1987 the Ing Foundation of Taiwan sponsored the first of a series of annual Computer Go tournaments. They provided generous sponsorship, with the winner of each annual tournament competing, using handicap stones, against inseis (trainee professional players, with strengths around amateur 6-dan) for further prizes. The fewer the handicap stones needed by the program, the bigger the prize it could win, on a progressive scale with a maximum of 40,000,000 Taiwanese dollars (worth over US $1,000,000) for a program able to win against the inseis with no handicap. This "million-dollar prize" was never won, the sponsorship from the Ing Foundation ended after the 2001 tournament, and only the prizes for handicaps of 11 stones and more were ever claimed.

GNU Go was published in 1989 as the first open source program.

Very strong players were still able to beat programs in 1998, while giving

[3] `http://www.computer-go.info/events/acorn/1984/index.html`
[4] `http://www.computer-go.info/events/index.html`

handicaps of 25-30 stones. There was also a case in the 1994 World Computer Go Championship where the winning program, **Go Intellect**, lost all 3 games against the youth players while receiving a 15-stone handicap. In general, players who understood and exploited a program's weaknesses could win even when giving much larger handicaps than typical players.

The Computer Go Olympiad, organised by the International Computer Games Association, was started in 1989 for 9x9 and in 2000 for 19x19, with the initial tournaments both being held in London and won by **Dragon Go** (9x9) and **Goemate** (19x19).

In 2003, **Go++** beat a 5-kyu amateur in a 9-stone-handicap 19x19 game.

WinHonte in 2005 appears to be the first program using neural networks.

In 2006, advances in the strength of Go programs were still being made, though the rate of advance had slowed. Processor speeds were continuing to double every two years in accordance with Moore's Law, but this did not help, as the algorithms used by the best programs did not scale well, if at all. However, in this year Kocsis and Szepesvari published their seminal paper Bandit based 'Monte-Carlo Planning'. This describes a Monte-Carlo based algorithm that was effective for computer Go (in fact a French team was working on a closely-related algorithm at the same time). This not only led to a rapid advance in the strength of Go programs over the next few years, it allowed them to use a method that did scale well, so now Moore's Law was working with the programmers again.

MoGo, developed by the French team mentioned above, beat an 8-dan professional in a 9-stone-handicap 19x19 game in 2008. It was running on an 800-node supercomputer. He estimated the playing strength of Mogo as being in the range of 2-3 amateur dan. In the same year the program **Crazy Stone** running on an 8-core personal computer won against a 4-dan professional, receiving a handicap of eight stones.

In 2009 **Zen** playing on the KGS Go server achieved a rating of 3-dan, playing 19x19 games against human opponents. [The KGS rating scale is slightly weaker than the European rating scale, close to the American scale, and rather stronger than the Japanese scale.] **MoGo** and **Many Faces of Go** beat professionals taking a 7 stone handicap.

Through 2010 and 2011 programs showed steady improvement with **Zen** beating a professional with 6 stones. In July 2010 **MoGoTW** won an even 9x9 game as white against a top professional. However, at the end of 2010 John Tromp, approximately 1 dan, beat **Zen** in a $1000 challenge in a best of 5 match; he lost a rematch in early 2012 comprehensively.

In March 2012 **Zen** beat top professional Takemiya Masaki 9p at 5 stones by eleven points, followed by a stunning twenty point win at a 4 stone handicap. Takemiya remarked "I had no idea that computer go had come this far." It also reached the rank of 6 dan on the KGS Go Server playing games of 15 seconds per move. However, it's not clear how seriously professionals have been taking these exhibition matches.

At the 27th Annual Conference of the Japanese Society for Artificial Intelligence in June 2013, **Zen** defeated another top professional with a 3 stone handicap with a time setting of 60 minutes plus 30 seconds byoyomi. In March **Crazy Stone** beat Yoshio Ishida with four handicap stones.

In 2014, for the *codecentric go challenge*, a best of five match was played between **Crazy Stone** and eleven times German Go champion Franz-Jozef Dickhut, 6 dan amateur, without a handicap. Dickhut won as was expected by most observers and the contender himself before the match. However **Crazy Stone** won the first game by 1.5 points, which was a resounding mark that the top programs have reached top amateur level.

This was reprised in October 2015, this time with **Zen** playing and Dickhut won again 3-1 with **Zen** winning the first game, again by 1.5 points.

Zen has been champion of the Computer Olympiad from 2011 to 2015 in all board sizes, but it should be noted that **Crazy Stone** did not take part.

In November 2015 there were published articles indicating that Facebook as well as Google were developing Go-playing programs, with Facebook's available for play on KGS.

See the Computer Go pages on Wikipedia [5]*, Sensei's Library* [6]*, Jay Burmeister and Janet Wiles Technical report* [7] *(good for historic stuff up to about 1996) and computer-go.info* [8] *for more details, the references and discussion of the problems and techniques involved in programming Go.*

AlphaGo

Toby Manning toby.manning@dsl.pipex.com

It was while I was travelling to the Isle of Man Go Tournament that I received a strange telephone call from Jon Diamond "Are you free for the week of October 5-9?" I responded Yes, and my request for more details was met with "I can't tell you".

I was then contacted by Google DeepMind, who asked me to sign a Non-disclosure Agreement; it was only after signing it that they would tell me what it was all about.

DeepMind, a British Artificial Intelligence Company acquired by Google in 2014, had been developing an AI computer program to play Go. They reckoned that their program **AlphaGo** could beat any other software that was publicly available, and they wanted to test it against a professional Go player.

They were arranging a match against Fan Hui from Bordeaux, who is one of the strongest players living in Europe; he won the European Championship for the third successive time at Liberec this year. They wanted someone from the British Go Association to see "fair play" and Jon Diamond had "volunteered" me!

[5] http://en.wikipedia.org/wiki/Computer_Go
[6] http://senseis.xmp.net/?ComputerGo
[7] http://staff.itee.uq.edu.au/janetw/Computer%20Go/CS-TR-339.html
[8] http://www.computer-go.info

They wanted me to remain independent, so instead of payment they agreed to sponsor the London Open in 2016. This resulted in an arrangement that satisfied all parties.

I had a chance to play AlphaGo as part of the preliminary discussions – I lost by about 17 points. It was clearly of dan-strength, but I was not convinced it was of professional strength (it seemed to make a couple of elementary errors, but I was not strong enough to take advantage of them). However, it may be that it knew it was ahead and was simply playing conservatively.

So in early October I went down to London to act as referee. In order to make the playing conditions as natural as possible – we all know that playing on a computer is not the same as playing on a board – the game was played on a normal Goban. Aja Huang (5 dan) who works for DeepMind placed the computer's moves on the board, and then communicated Fan Hui's moves to the computer. A representative of DeepMind pressed the clock.

There were two separate matches, each of five games. During the mornings the time limits were 1 hour, with 3 periods of 30 seconds byoyomi; the afternoon games were played completely in byoyomi (also 3 periods, 30 seconds).

The commentary below deals with the five games played in the morning, which are the ones included in the scientific publication **Nature**: Silver D. et al. Mastering the game of Go with deep neural networks and tree search. Volume 529, issue 7587, pp 484-489: http://www.nature.com/nature/journal/v529/n7587/full/nature16961.html.

Comments are by me, following a discussion with Fan Hui and Aja Huang, supplemented by some subsequent ones from Hajin Lee (a Korean professional). Unfortunately, since I didn't record all the details at the time there might be some errors though...

SIDEWAYS LOOKING PERSONS

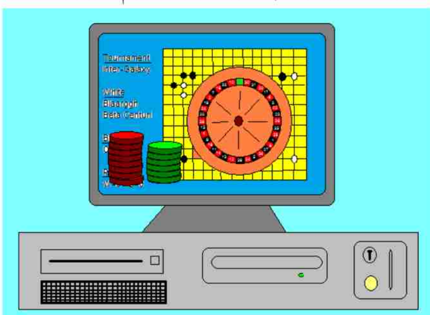

I added the Monte Carlo system to myGo playing program ..... and all went well until it discovered the on-line casinos.

Game 1: 5 October 2015
Black: Fan Hui
White: AlphaGo
Result: White wins by 1.5 points

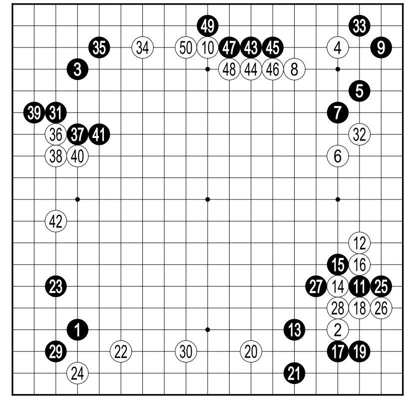

Diagram 1 (1-50)

The first game is shown here. It was a very quiet game, with very little fighting; I think Fan was trying to get the measure of AlphaGo. The crucial part came when Fan invaded at 43 and AlphaGo let him connect out; this seemed to result in White getting a wall that was not doing much. Indeed, this seemed to represent AlphaGo's style: it is not very aggressive as long as it's not behind.

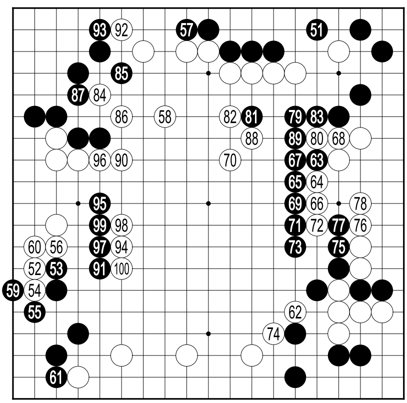

Diagram 2 (51-100)

However, Fan then relaxed a bit, and when White played the sequence 128-134 to get the large yose in that corner in sente; the game was close.

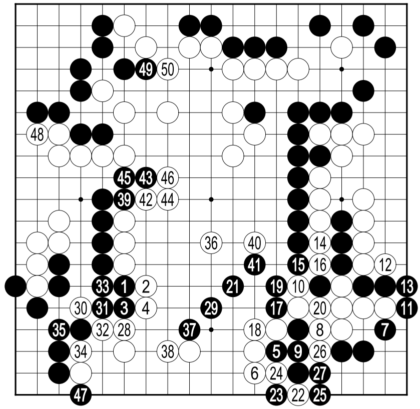

Diagram 3 (101-150)

A tight yose resulted in AlphaGo winning by 1.5 points.

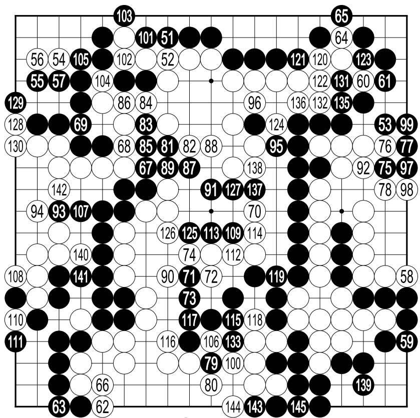

134 at 79.

Diagram 4 (151-245), White fills ko

Fan Hui used all his time and went into the third period of byoyomi; Alpha Go used about 45 minutes.

Game 2: 6 October 2015

Black: AlphaGo

White: Fan Hui

Result: Black wins by Resignation

On the second day AlphaGo took Black, and played the onadare ④-⑫. It used to be considered joseki, but according to Michael Redmond is now thought to favour White (i.e. Fan Hui). I was told that AlphaGo did NOT have a joseki dictionary: it was working it out from first principles (although it has used a lot of professional games for training purposes).

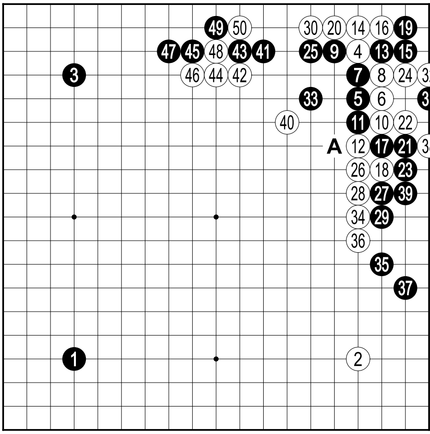

Figure 1, Moves 1-50

⑯ was, according to Ishida, invented by Go Seigen and was a "revolutionary move". ㉖ was considered an overplay (the old joseki is at **A** in Figure 1, according to Ishida, but the analysis is complex and beyond this article), however it's now the new joseki played by many professionals.

45 was a mistake, but Fan failed to take advantage of this: ㊻, instead of being a push along the top, should have been the push and cut shown in Diagram 1, where White captures the 6 black stones and has enormous thickness. If Black captures ③ by playing at **A**, White plays at ❹ then at **B** and gets a good result, the ladder being good for him.

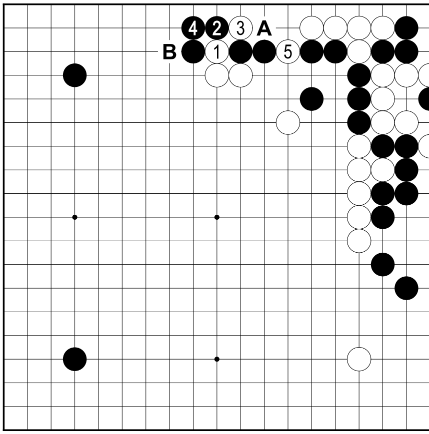

Diagram 1

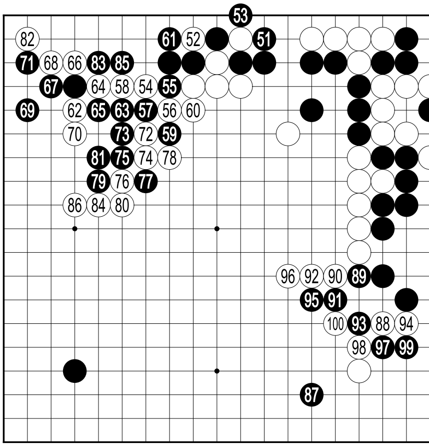

Figure 2, Moves 51-100

Fan then compounded this mistake by playing tsuke with 62. If he had played at 65 the game may have been easy. But with 70 (over which Fan took a long time) he had a choice: to live in the corner and give Black a lot of thickness which would nullify his own strength in the centre, or to sacrifice the corner and increase his dominance in the centre.

⑧② should have been at ⑧④ immediately to seal off the side in sente. In fact analysis afterwards showed that he could live in the corner with ko, but the analysis is complicated.

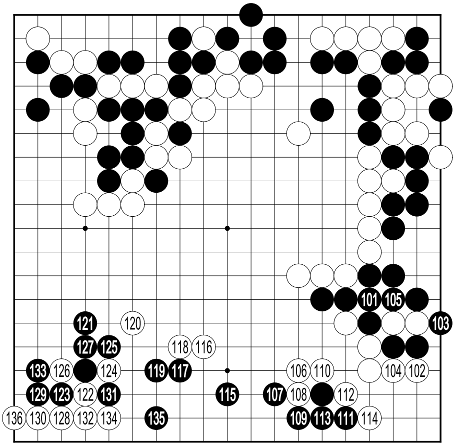

Figure 3, Moves 101-136

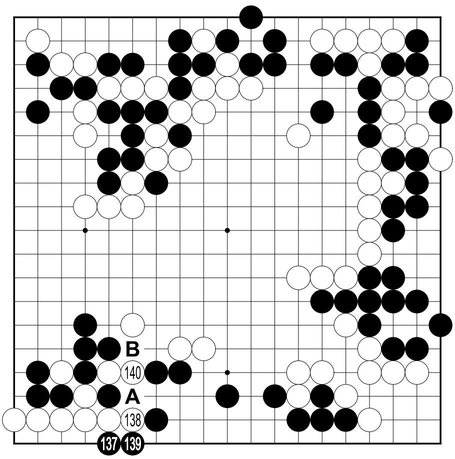

Diagram 2

After ⑬⑥ the White group is alive (see Diagram 2). If Black tries to kill with ❶❸❼, the sequence to ⑭⓪ results, after which A and B are miai.

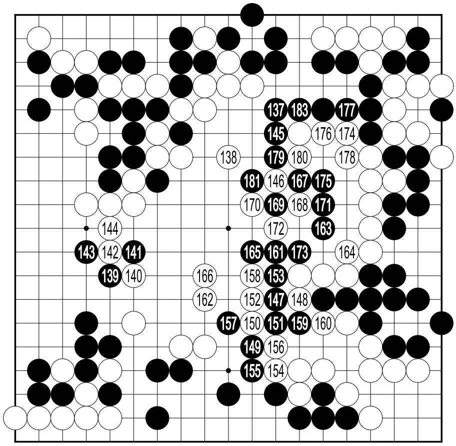

182 at 169.

Figure 4, Moves 137-182

The game proceeded and with 147 Black attempted to prevent White making a large territory in the centre. Black did this successfully and when it created the seki in the middle White had no hope and resigned.

So after two days the score was 2-0 in AlphaGo's favour.

In discussion, Fan thought he would do better if the time limits were longer. In particular, AlphaGo was playing relatively quickly which further reduced the time available to Fan for thinking.

Game 3: 7 October 2015

Black: Fan Hui

White: AlphaGo

Result: White wins by Resignation

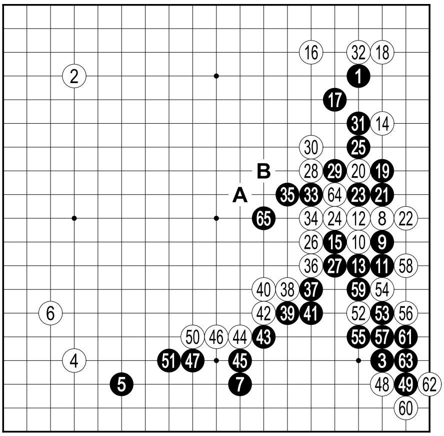

Figure 1, Moves 1-65

Fan took Black in this game, and a complex position rapidly developed on the right hand side (professionals often play immediately at 20 rather than at 18 to solidify the White group).

The conclusion was that this position was bad for AlphaGo, which gave away a very large corner without gaining sufficient compensation. It could only win the game if it could use its thickness to attack, and in particular capture the two stones 33 and 35 in the centre.

60 was particularly crass (and one of the few times where AlphaGo seemed to make a particularly bad move); White cannot live in the corner and the response at 61 is a significant gain for Black.

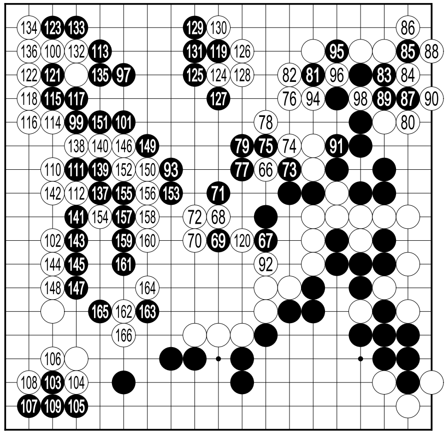

Figure 2, Moves 66-166

But then Fan made a catastrophic overplay when he played the kosumi (diagonal move) at 65, which AlphaGo duly punished. A one-point jump at **A** or kosumi at **B** should have sewn up the game for him.

Fan then compounded his mistake with 67 and by failing to make his group on the top right live unconditionally. Instead he allowed White to play atari at 94 and, although he salvaged a ko, the game was effectively over.

Fan was extremely upset with himself over these blunders, and had to go out for a walk to compose himself.

Game 4: 8 October 2015

Black: AlphaGo

White: Fan Hui

Result: Black wins by Resignation

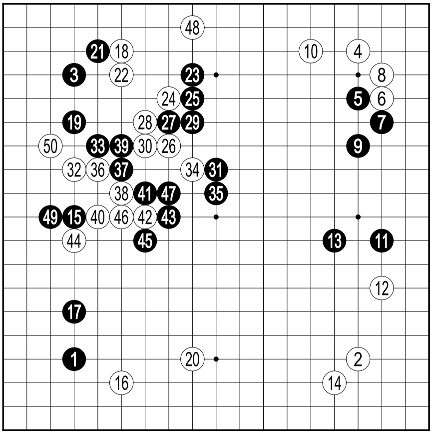

Figure 1, Moves 1-50

This morning Fan took White. The fuseki 1-13 had been played in one of the afternoon games, but Fan played 14 in the lower right corner to see what AlphaGo would do. It immediately made a san-ren-sei, but then tried to turn the moyo into territory; Fan afterwards suggested that 17 should be an attachment above 16 to further expand the moyo (and keep White to a low position on the bottom).

The attack with 21 and 23 may be good locally, but is meaningless in a global sense, because the resultant Black wall 23-31 is nullified by White's strength in the upper right hand corner. If Fan had simply run away by jumping to 47 then Black would have achieved nothing from his attack.

The invasion at 32 was an overplay, giving White two groups to look after. He made both of them live, but Black 45 was painful, and both White groups are in poor shape. However, when AlphaGo played at 49 and threatened play 51 or one point to the left of 32, White should sacrifice his group (at least temporarily), playing at 64 instead.

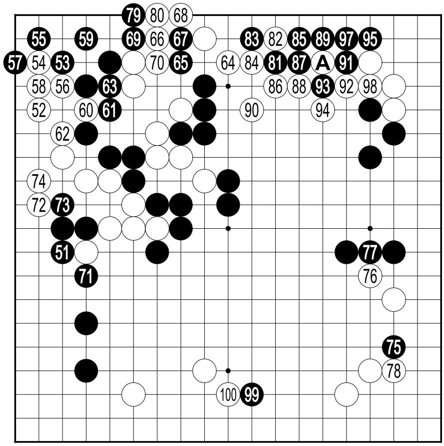

96 at A.

Figure 2, Moves 51-100

The capture of these stones is worth about 25 points, but a continuation at 90 would be worth nearly as much: more importantly, saving the left hand group leaves a weakness behind at 65 which Black later exploited mercilessly.

When AlphaGo sought to exploit this weakness, Fan made a mistake: 70 should be atari at 79; for Black to start a ko is very dangerous, as White can win the ko in sente, threatening to cut one point below 30 and kill the entire corner.

Subsequently then White could possibly play 90 at 97. This position seems to be yose ko (a ko that AlphaGo has to win twice). Fan could then get adequate compensation even if he lost the ko.

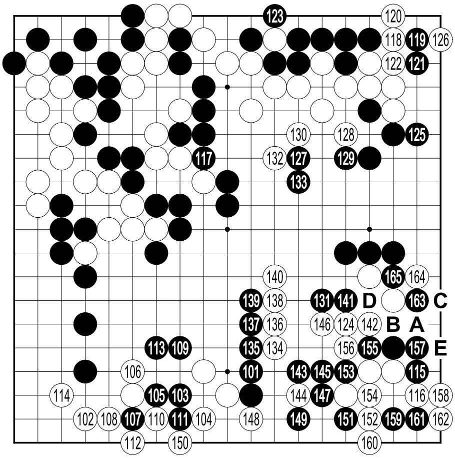

Figure 3, Moves 101-165

The game continued, but after 165 a ko in the lower right hand corner develops (White at **A**, Black at **B** etc.) Fan recognised that he could not win the ko: not only does he have fewer ko threats, but if Black wins the ko it is in sente as AlphaGo then threatens to kill the corner by playing on the 1-2 point.

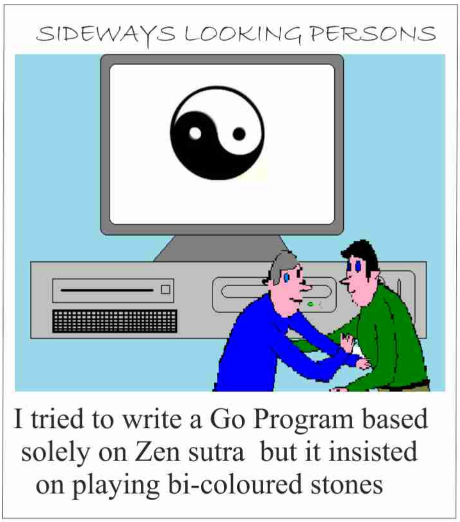

Game 5: 9 October 2015
Black: Fan Hui
White: AlphaGo
Result: White wins by Resignation

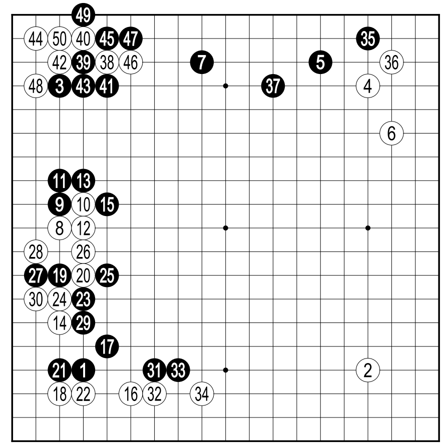

Figure 1, Moves 1-50

The fuseki in this game had previously been played in the afternoon games on Monday and Wednesday. 41 was mistake. It was better to block at 50, since White pulling out with 46 created many problems.

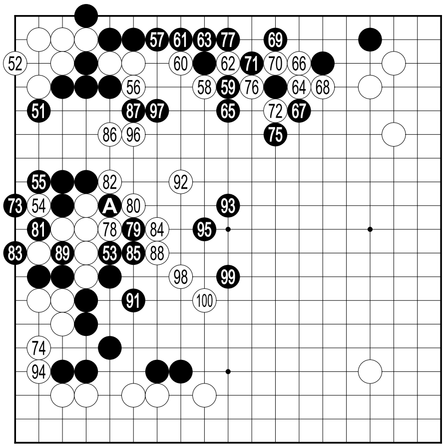

90 at A.

Figure 2, Moves 51-100

The game proceeded until 78 when AlphaGo threatened to break out through the Black wall. Fan afterwards thought he should have played simply, answering the ataris; however he played at 81 instead, but this was a total waste of a move - a catastrophic mistake: it should have been directly at 83.

93 started an attack on the White central group, but White had time to *take the money* with 94 before defending. Black should have played there himself, before attacking, as 94 weakened the Black group on the lower left and made it easier for White to escape with his central stones.

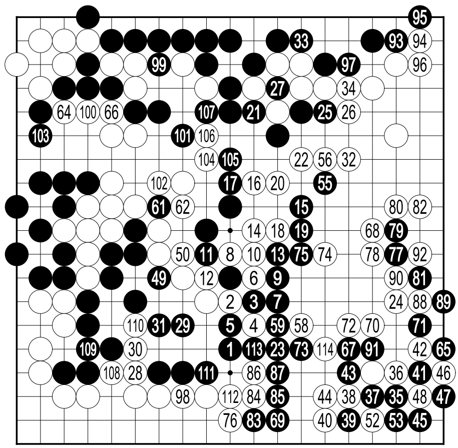

51 at 41, 54 at 48, 57 at 41, 60 at 48, 63 at 41.

Figure 3, Moves 101-207

It was always going to be difficult for Black to attack this central White group successfully, and when the attack petered out Fan knew he was well behind. He struggled on for a while, hoping to salvage something from the wreckage, but eventually resigned: he was over 20 points behind.

Conclusions

Jon Diamond president@britgo.org

Hajin Lee, who commented on the first 4 games, said *AlphaGo's strength is truly impressive! I was surprised enough when I heard Fan Hui lost, but it feels more real to see the game records.*

My overall impression was that AlphaGo seemed stronger than Fan, but I couldn't tell by how much. I still doubt that it's strong enough to play the world's top pros, but maybe it becomes stronger when it faces a stronger opponent.

I agree, it's an impressive achievement and it looks like a human player – when I first played through the games I didn't know which side was AlphaGo and couldn't tell. In retrospect this isn't too surprising as AlphaGo has been training using human games.

I've got two conclusions based on these games – it seems to play unnecessary sentes on occasion and it definitely plays conservatively when it's very confident it's ahead and aggressively if it's definitely behind. Apart from that I can't see any obvious weaknesses... but maybe these issues will be fixed by now!

One significant aspect of this match was that AlphaGo analysed orders of magnitude fewer positions than IBM's Deep Blue did in the Chess match in 1996 against Gary Kasparov. Deep Blue also had a handcrafted evaluation function, which AlphaGo does not. These indicate the general improvements in AI techniques that Google DeepMind have achieved. I think the techniques used, which include Convoluted Neural Networks and MCTS, are definitely applicable in other artificial intelligence application areas, such as Facial Recognition and perhaps autonomous cars, but that's another story...

The technical article is **Nature**: Silver D. et al. Mastering the game of Go with deep neural networks and tree search. Volume 529, issue 7587, pp 484-489: www.nature.com/nature/journal/v529/n7587/full/nature16961.html

If you're interested you can read the abstract free of charge, but you'll have to subscribe to read the rest... There's not much Go stuff and it's quite heavy going, but that is only to be expected of a scientific article.

Finally, how does this affect humans playing Go? Well, I think not very much. The loss by Kasparov against Big Blue in 1996 didn't really affect Chess, although there are programs that help Chess players with databases of games and analysis, and I think the same will apply to Go.

It doesn't feel like we need to worry too much about how this technology will affect face-to-face games though, since, apart from anything else, the gain from the improvement of an odd move isn't as high as in Chess.

So keep on playing!