# A Human-Computer Team Experiment for 9x9 Go

Darren Cook

Software Developer/Researcher, Tokyo, Japan

**Abstract.** Monte Carlo Tree Search has given computer go a significant boost in strength the past few years, but progress seems to have slowed, and once again we have to ask ourselves how can computers make effective use of the ever-increasing computer power. In 2002 we started a human-computer team experiment with very long thinking times and no restrictions on procedure, to see how strong such a team could be. We will introduce our experimental method and show the results so far.

## 1 Introduction

We take the view that a human-computer team, with long thinking times, can give us a very high standard of play, and can give us insights into how to tackle the weaknesses of today's computer programs.

Between 2006 and 2008 Monte Carlo Tree Search (MCTS) gave a boost to computer go of about five ranks on the 19x19 board. The algorithm improved the strength on the 9x9 board even more so that, on reasonable hardware, computer go programs are now a match for strong amateur players. But even on the 9x9 board, even using a supercomputer, their play still has weaknesses.

9x9 go is often considered a learning game, one that offers no serious challenge for players of intermediate strength and above. Study of professional 9x9 games quickly proves it is much deeper than generally realized: the higher the professional player's 19x19 rank the more likely he is to win a 9x9 game, and this is even more distinct in the amateur dan ranks [1]. However the consequence of this general belief is that there has been little development of opening theory, compared to 19x19.

Ingo Althöfer has performed a number of experiments in the computer chess domain [2][3][4] and has tried extending these experiments to computer go [5]. Unlike the current work these are quite restrictive, only allowing the human to choose between suggestions by the computers.

There is also Advanced Chess [6], introduced by Gary Kasparov, which allows a human and computer to work together however they wish, but the human is ultimately in control. Freestyle Chess allows teams, with any number of human and computer members on the team, and an "anything goes" philosophy. The current work is closest to Freestyle Chess.

This paper is organized as follows. Section 2 describes our initial experiment and section 3 describes what is has evolved into. Section 4 presents results and section 5 describes future work.

## 2 The Experiment

### 2.1 Motivation

The experiment began, in 2002, with a couple of questions: How much stronger would using computers, and lots of thinking time, make the author. And secondly, what is the strongest opening move on 9x9? We have answered the first question, and gained some good evidence to suggest an answer for the second.

### 2.2 Experiment Set-up

The experiment has been conducted solely on Little Golem, a turn based server, where players get 240 hours initially and then 36 hours per move. Multiple games are played simultaneously. There are various types of tournaments, including a championship where players are organized into divisions. Japanese rules, and 5.5pt komi, are used. There are currently 1,400 registered go players, with a little over 260 of those playing in the championship. An alias ("sm9") was used for our player, and no hint that computers were being used was given. For more information see [7].

At the time the computer part of the team was Many Faces 11, which was ranked at around 10 kyu, meaning it could comfortably beat beginners, but most players could beat it after a couple of months of serious study of the game. The human part of the team was at the 5-7 kyu level on online servers. In other words, able to consistently beat computer programs, but not really very strong. The human player was plagued by tactical oversights which conveniently was Many Faces's biggest strength. The other key part of the team was the opening book, described in section 2.4.

### 2.3 Procedure

Opening moves were chosen from the opening book, described in section 2.4. Once out of the opening the procedure was to explore different variations, usually deeply, then use the computer to score the game. This scoring function shows stone status and territory for each side, in addition to showing the score, making it easy to see when it has misunderstood something. When that happened a few more moves were played to clarify the situation and allow more reliable scoring. Once happy with a move the human would play his chosen sequence against the computer to check no tactical flaws.

In addition to helping avoid tactical blunders, the computer's ability to quickly score allowed the human side of the team to explore many more variations.

### 2.4 Opening Book

Our game collection contains four types of games: 1. Commented professional game records, with variations. These were input by the author, from watching the Minigo series on Japanese TV's Go and Shogi Channel, with a few from 9x9 games on NHK's go channel. 2. Other professional game records (primarily other Minigo games). 3. Downloads from online go servers, restricted to just games between reasonably strong players where possible. 4. Studied games.

Studied games refers to game records made while analyzing positions within this experiment. They contain a large number of variations, most played out deeply, as described in sections 2.3 and 3.2. Early on in the experiment the professional games were most useful, but soon the studied games became invaluable. Professional game records turn out to be not as useful as they would be in 19x19 go because the games are generally played under blitz time controls.

When a particular opening lost, instead of abandoning it, it was played again in the next championship only changing what we considered to be our last mistake: the move closest to the end of the game where our analysis shows we could have won if we had played differently. E.g. move 10 might be changed. Then, if that failed again, yet another alternative move 10 might be tried, or if none seem good then a different move 8 might be tried. As we will show in section 4 the ever-improving database of studied games has been a major contributor to sm9's strength.

One way to view the growth of the database over time is as a count of unique game variations starting from the two main openings: 5-5 and 3-4.

| | All 5-5 | All 3-4 | Pro 5-5 | Pro 3-4 |
|---|---|---|---|---|
| Sep 2004 | N/A | N/A | 70 | 69 |
| Feb 2005 | 9,062 | 2,795 | 81 | 106 |
| Mar 2006 | 11,854 | 4,182 | 107 | 162 |
| Nov 2007 | 16,970 | 5,993 | 261 | 285 |
| Dec 2008 | 19,259 | 7,198 | 261 | 285 |
| Jan 2010 | 25,032 | 10,519 | 261 | 286 |

## 3 The Evolved Experiment

After playing in various tournaments in 2002 and winning the 2.1.1 championship, we concentrated on just the championship: fewer games studied intensively against the strongest available opponents satisfies the experiment's goals best.

The MCTS algorithm gave computer go programs a dramatic strength increase. Combined with ever-faster hardware, the computer part of the team improved at 9x9 go from about 10 kyu in 2006 to about 1 dan or even higher in 2008. Naturally the sm9 player took advantage of this. In 2006 an early version of Crazy Stone was used to give an alternative opinion to Many Faces. Overall it was similar strength to Many Faces at 9x9, but it's strength profile was different: the closer the game got to the end the stronger its moves became.

From November 2007 Mogo was released, and that was then used for the main analysis, with Many Faces 11 being used to play out the sequences and check the score. In November 2008 Many Faces 12, with MCTS and a big jump in strength, was released and used in the team. In November 2009 Fuego 0.4 was also included, giving three strong computer programs in the team. All are using MCTS but are just different enough to be worth using together.

### 3.1 The Current Process

This section will describe the current move selection process. The process is deliberately kept informal and flexible (we are interested in discovering the best move in each position, not artificial restrictions even if they would allow the experiment to be more repeatable). However there are still three parts we can identify:

1. Opening book
2. Move selection and validation
3. Unbalanced quiescence search

The opening book, described in section 2.4, is still used in preference to any of the computer programs. Once we leave that database, an sgf file for the game is created and opened in Many Faces. We may spend some time exploring some variations to understand the position, or we may immediately choose a good-looking move. Once a top-level move has been chosen we enter the validation stage.

Many Faces is then asked to play as the opponent. While it is thinking the position is set up in Mogo and Fuego. When Many Faces chooses a move, analysis is started in Mogo and Fuego, and while they are thinking the prime variation suggested by Many Faces is input and considered. Then the prime variations suggested by Mogo and Fuego are also input and considered.

At this point we have three scores, where 50% is an even game, less than 50% means good for sm9 (as the programs are playing from the opponent's point of view), and greater than 50% means bad for sm9. As an additional, informal, input we have the human opinion of each of the three prime variations with regard to potential nakade misunderstandings or complex semeai or ko that might cause the scores to be less reliable. Special attention must be also paid in situations where the scoring is different in Japanese and Chinese rules. I.e. a point of territory in a seki, or a ko where pass is the only available ko threat. Many Faces is the only of the three programs that fully plays with Japanese rules. Fuego supports Japanese rules, but is using Chinese rules for simulations.

If all three scores are below 45% (*"good"*) it means all three programs agree the human move seems good; if nothing seems suspicious the move will be played with no further analysis.

If all three scores are above 55% ("*bad*"), and it wasn't previously known to be a losing position, it most likely means the move chosen by the human is a mistake. The chosen move will be undone and an alternative tried, repeating the whole validation stage.

When all three are scoring close to 50%, but suggesting different moves, we might try each move in a different engine. For example, if Many Faces suggested F2 and Fuego suggested D3, we try the F2 in Fuego and ask it to evaluate that. And try D3 in Many Faces.

When the programs disagree on who is winning, for instance one says *bad*, the other two say *good*, then we set it up as a contest. The side that thinks sm9 is losing will play for the opponent, and the human and the other two computers will work as a team to prove it wrong. The whole validation algorithm is repeated at each ply. When two programs say *bad*, and one says *good* we will usually go back a move and try to find something less controversial. If nothing obviously better we will come back and play each of the programs that are saying *bad* to try and prove each of them wrong.

Sometimes the computers are asked for their move suggestion at the root, but usually it is the human who suggests the moves. One reason for this is it gets us one move deeper in the tree. Another is that it helps prevent the sloppy moves leading to a half-point win that MCTS tends towards. However, like MCTS, the first move that seems to win is chosen: there is no attempt to maximize the score.

The final element, a kind of quiescence search, is used when the dust has settled and we think we have a win. This will be at around move 20 in a quiet game, at around move 30 in a more unsettled game. It may be used at the root, or after using the above-described technique for a few ply. The human player plays out the position with simple, passive moves; the standard moves to tidy up the boundaries between territories. It is called *unbalanced* quiescence search because the player who appears to be winning has to play at a disadvantage: if there is a ko we assume the opponent will get to fill it and if a move would leave a cut behind then a more solid move is chosen; the opponent moves do not have to have such good shape and possible weaknesses they leave behind are overlooked. There is no branching, and the moves are played quickly. If at the end of this we have a win of 1.5pts or more then we can feel confident in the result. If the win is 0.5pts we are less comfortable, and may go back and try again with different opponent moves, or in a different order. If the position is a loss we know that to win we have to play a more complex endgame and will revert to the main move validation stage for a few more ply.

When we are losing, the quiescence search will be performed from the opponents point of view, and if we lose even after getting all the best endgame moves we should resign.

### 3.2 Alternative Approaches

Our approach is just one possible one; many decisions have been made arbitrarily, but the evolved experiment represents eight years worth of tuning those decisions.

A voting system is deliberately not used. It would be tempting to go with the majority view when two programs think one thing, and one thinks the opposite. But this is hiding one's head in the sand: a disagreement means there is something not clear. It is much more effective to play out the prime variations until the program with the misunderstanding is discovered. Especially with disagreements in the early middle game it is possible for two, or even all three, programs to be wrong.

### 3.3 Environment

The teams consists of: Many Faces 12 with about 75 seconds per move, and 2-4 cores; Mogo, November 2007 release, with 75s and 3 cores (i.e. 225 seconds of CPU time); in the early game this might be doubled to 150s (450 seconds of CPU time); Fuego 0.4 is given 90 seconds and 3 cores (i.e. 270 seconds of CPU time). There is little science behind these times; they were chosen as a compromise between getting the best possible suggestion and human patience.

### 3.4 Time Investment

Wall clock time for the move selection process varies depending on the stage of the game and who is winning. When using moves from the opening book it is 0 to 5 minutes. In the early middle game (e.g. moves 8 to 12) it is 5 to 15 minutes. In the main middle game (e.g. moves 12 to 24) it will be 10 to 15 minutes if winning, 15 to 60 minutes if losing. In the endgame (e.g. moves 24 to 40) it is 1 to 10 minutes for a winning position, and 5 to 20 minutes for a close losing position (if the position is hopeless we will resign).

Average time per game is approximately 5 hours, roughly 15 minutes per sm9 move. For tournament 22.1.1 the total time spent on 8 games (371 moves) was 46 hours, spread over two months. The shortest time spent on a game was 1.5 hours. The longest was 10 hours, spent on the game described in section 5.2. 8.5 hours of that was after move 30, in other words, after the game-losing mistake had already been made. For tournament 21.1.1 25 hours was spent on 6 games (two players dropped out) (189 moves). For 20.1.1 37 hours was spent on 7 games (274 moves).

## 4 Results

Figure 1 shows the ratings chart of sm9, overlaid with two players of known rank. The black line is a Japanese 7-dan, and the light gray line is a European 4 dan. sm9 quickly became one of the strongest players on the site, using just an early version of the opening database and the help of Many Faces for scoring. But then from 2003 a number of strong players (including a number of high-dan amateurs) joined, and sm9 struggled against them.

From the chart we can suggest that in 2004 the author + Many Faces 11 + small opening book is 4 ranks below a Japanese 7-dan/European 4-dan. This equates to 10 ranks higher than Many Faces by itself, and 6 ranks higher than the author by himself.

As of February 2010 sm9's rank of 6.2d is about 2 ranks higher than those same players. The author is now around 1-dan level at 9x9. Many Faces and Fuego are about 1 dan at 19x19, but we estimate MCTS go programs are one to three ranks stronger at 9x9. If we call the individual programs 4-dan at 9x9, we can say the team adds about two ranks over the strongest member of the team. As an aside, the human player is now comfortably the weakest member of the

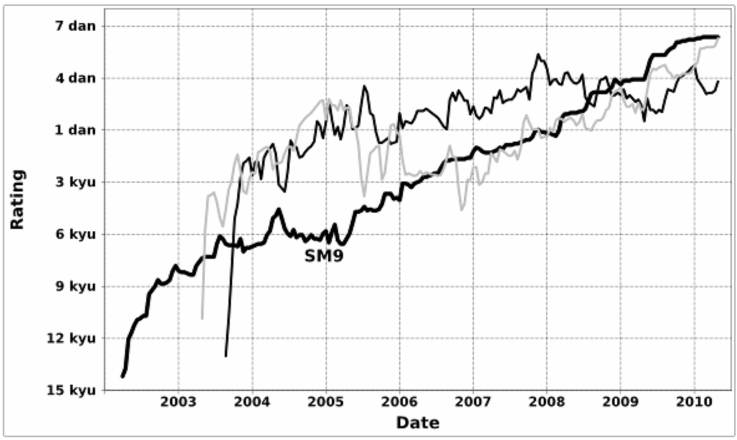


**Fig. 1.** rating chart

team, losing every game played against each of the computer programs at the settings used in the experiment. Only by reducing the CPU time by an order of magnitude is the author able to win some of the games.

While MCTS has certainly helped sm9's level a lot, Mogo did not start to be used until the end of 2007, so the improvement that began in 2005 we instead consider due to the advantage accrued from the ever-improving opening database, and the way losing openings were studied systematically to flesh out the database in a useful way. Recently sm9 already has a clear advantage in many games by move 8, and in worst case has a close game.

Since late 2008 sm9 has been the strongest 9x9 go player on the Little Golem site. sm9 won the 17th, 19th, 20th and 21st championships, with 2nd place in the 18th and 22nd.

It is important to point out that the human opponents effectively have all the thinking time they want, and should be playing go at a higher level than they normally do. They are less likely to make tactical blunders, and have time to accurately count the score. We also assume that the strong players most successful on Little Golem (i.e. sm9's top division opponents) are also those for whom this way of playing is most beneficial to their strength. At least some of the top players appear to be examining variations using a computer, which would also help raise the level of their play. Also, from March 2010 a Fuego bot has been openly playing on Little Golem, apparently using 10 minutes/move, so a little more than the Fuego settings used in our experiment. It will be interesting to see where its rank ends up, but it appears to be heading for 2-3 ranks lower than the 4-dan we had predicted. If that turns out to be the case it suggests that the benefit of our team is even more than the two ranks we estimated above.

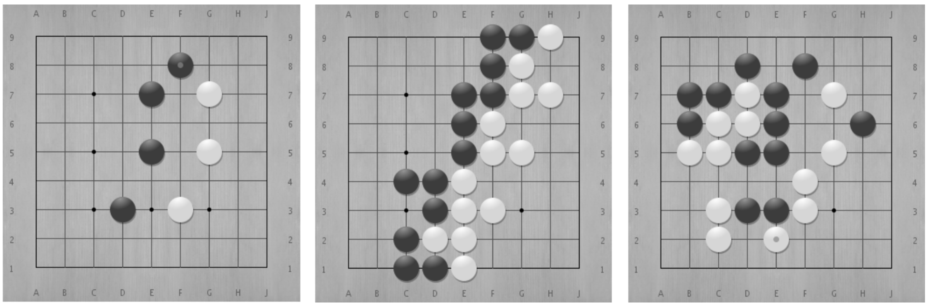

**Fig. 2.** example game positions

### 4.1 Examples

The positions in figure 2, taken from one game, show examples of when each team member makes mistakes that get caught by the current process. In the left position the game has been quiet so far and if white plays passively here the territories will solidify.

The human player briefly considered E2, but then chose D2. Many Faces thought 56.8% to black, and Mogo thought 54.3%. Playing it out gives us the middle position in figure 2, showing a 1.5pt win to black. The computer players have saved the human from a mistake.

The human then considered C3, understanding it will either connect around (and thus be more than 2 points bigger than white D2 was), or be very hard to kill. Now Many Faces thinks 51.2% to black, and Mogo thinks 50.1% to black; however C3 was actually played, and it led to a win. In this position Many Faces would have chosen white E3, though (correctly) thinking white will lose the game with it. Mogo would have chosen B7, though with only 50% confidence. In both the author's intuition and subsequent analysis white B7 fails to live and therefore loses the game. The human has saved the computers from a mistake. (Incidentally, Fuego would have chosen C3 and correctly thought white was winning with it; so in this case Fuego made the human redundant.)

The rightmost position shows how the game continued. Now the computers disagree the other way. Mogo and Many Faces think black D2 is only 45% to black. Fuego thinks D2 is good for black (though only just: 51%). This position is (probably) a 1.5pt win to white, so Fuego is wrong here. Taken together, these examples show how the team members nicely complement each other.

### 4.2 Study Of A Loss

In February 2010 sm9 lost its first game (by 0.5pt, as black) in 18 months. The opponent ("xaver", the light gray line in figure 1) also earned a perfect score in the top championship. Apart from demonstrating a very high level of play, this particular game nicely highlights the weaknesses of the current team, so we will look at it here.

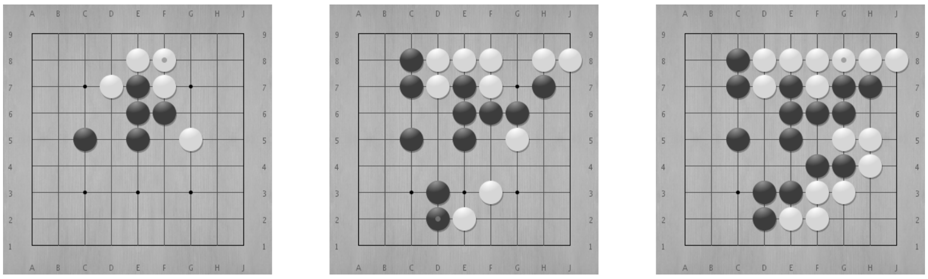

**Fig. 3.** example lost game

sm9, as black, started with a tengen opening, and the first nine moves were all from the book. White 10 at F8 (see left position of figure 3) was unusual, and possibly a mistake. sm9 briefly considered G4 and H5, before playing the obvious-looking G6 move. After G6 the score estimates by the three computer programs were: Many Faces: 47.5%, playing F3 Fuego: 44%, playing G3 Mogo: 45.6%, playing G3

By move 20 sm9 was feeling very comfortable. The human player suggested D2 (middle position), as the simplest way to win, and all three programs agreed it was good (Many Faces: 36.3%; Fuego: 36%; Mogo: 38.8%).

sm9 then played some simple moves to seal up the boundaries, and we reached move 30, as shown in the rightmost position. It was at this point that we performed the unbalanced quiescence search and discovered it led to a half-point loss for sm9. After intense analysis we discovered even the complex lines lead to white wins. The problem was the bottom-right corner involved nakade, seki and ko, and in some variations white was even able to force and win a whole-board semeai. The weaknesses of MCTS had led to under-confident scores for white at move 20, because precision play was required by white in the corner in order to win. Additionally a few variations ended with black having a point in seki, which would be worth one more point to black in Chinese rules.

Their confusion persisted right to the end of the game. E.g. at move 45 Fuego and Many Faces thought 38% to white, whereas Mogo thought 70% to white. However Mogo's confidence was also a misunderstanding, as it could be beaten when it took white from that point (with 225s 2.2Ghz CPU time per move) against the author.

Black 21 at D2 was chosen by the human player, but it turns out all three computer programs also choose it. The obvious way to prevent this type of loss is introduce another computer engine which understands seki and nakade in its simulations. Another way would be to do the unbalanced quiescence search earlier on. The relatively weak GnuGo program takes 2 seconds to self-play to the end of the game and says it is 1.5pts to white, after black 21 at D2, so it could have raised a red flag.

### 4.3 Komi and Opening Move

Little Golem uses 5.5pt komi. Japanese pros switched to playing at 6.5pt well over a decade ago, and most computer go 9x9 tournaments use 7.5pts. However our game database is gradually giving more weight to the hypothesis that 5 points may be the fairest komi for the 9x9 board.

For the opening move sm9, as black, always plays the 5-5 (tengen). The author was expecting to be able to prove this is inferior to alternatives (because pros stopped playing it when they increased the komi to 6.5), but so far has not been able to. The circumstantial evidence is that 5-5 is either the strongest opening, or the joint strongest opening.

Another 9x9 opening study [8] also shows evidence that 7.5pts komi gives white an advantage, and that the 5-5 move is best. However this study only uses Mogo, with no special-handling of Mogo's blind spots.

## 5 Future Research

We have demonstrated how a team of human and computer players can play a level of go that is not just very high but also appears to be distinctly above the level of any individual team member.

Two interesting questions to ask are can we remove the human from the team, and does the process apply to larger board sizes? We feel an automated algorithm can capture most of the benefit of the current move selection process, and that this algorithm can apply at all board sizes, and are currently working on developing such a system.

In addition we are doing offline analysis of our game database (a long-term project involving years of CPU time) to improve the opening book still further, and we are introducing another program into the team that does heavier playouts and has more hard-coded knowledge about seki and nakade.

## Acknowledgments

The author would like to thank Prof. Reijer Grimbergen and Prof. Ingo Althöfer for their helpful feedback on an early draft of this paper.

## References

[1]: http://dcook.org/compgo/strength_correspondance_9and19.html

[2]: Althofer, I. (1991). Selective trees and majority systems: two experiments with commercial chess computers. In "Advances in Computer Chess 6" (Editor D.F. Beal), Ellis Horwood (1991), Chichester, 37-59.

[3]: I. Althofer. 13 Jahre 3-Hirn - Meine Schach-Experimente mit Mensch-Maschinen-Kombinationen. 3-Hirn-Verlag (1998), http://www.3-hirn-verlag.de/books.html .

[4]: Ingo Althofer and Raymond Georg Snatzke. Playing Games with Multiple Choice Systems. Lecture Notes in Computer Science Verlag Springer Berlin / Heidelberg Volume 2883 (2003), Computers and Games 2002, 142-153.

[5]: http://www.minet.uni-jena.de/fakultaet/iam/personen/StrausTalkd.html

[6]: http://en.wikipedia.org/wiki/Advanced_Chess

[7]: http://dcook.org/compgo/sm9_experiment.html

[8]: Audouard, P., Chaslot, G., Hoock, J., Perez, J., Rimmel, A., and Teytaud, O. 2009. Grid Coevolution for Adaptive Simulations: Application to the Building of Opening Books in the Game of Go. In Proceedings of the Evoworkshops 2009 on Applications of Evolutionary Computing: Evocomnet, Evoenvironment, Evofin, Evogames, Evohot, Evoiasp, Evointeraction, Evomusart, Evonum, Evostoc, EvoTRANSLOG (Tübingen, Germany, April 15 - 17, 2009). M. Giacobini, A. Brabazon, S. Cagnoni, G. A. Caro, A. Ekárt, A. I. Esparcia-Alcázar, M. Farooq, A. Fink, and P. Machado, Eds. Lecture Notes In Computer Science, vol. 5484. Springer-Verlag, Berlin, Heidelberg