Dear all,
I'm happy to share the results of 11 months of training KataGo on Igo Hatsuyoron 120 problem, probably the "hardest Go problem ever imagined".
This run capitalized on “g146 run” from David Wu (@lightvector, programmer of KataGo) when he ran a one-week experiment with 27 powerful GPUs on this same Igo Hatsuyoron 120 (IH120) problem in 2019, starting from KataGo's best 20b network at that time. More details on David’s run are available here:
https://lifein19x19.com/viewtopic.php?f=18&t=16995

Starting on 25th of August 2020 and until end of July 2021, with crucial help from David and Thomas Redecker, one of the best World experts of this problem, I started the same approach with a more recent and stronger version of KataGo: the starting point of this run was the best 40b network available at that time, 40b-s5G09, published in June 2020 and coming from g170's run.
The g170 run was much stronger than previous g146 run, with the best 40b net (40b-s5G09) almost a thousand Elo points stronger than the net previously used by David for IH120. And this g170 net was also able to play with many different rule sets: Tromp-Taylor, Japanese, Chinese and New-Zealand rules, as well as rules including a button (a 0.5-point bonus for the first player to pass the game). Many more details on the available rules can be found here: https://lightvector.github.io/KataGo/rules.html

**Short summary of this new Igo Hatsuyoron 120 run:**
Here are the main results of training KataGo 11-months on IH120 problem:

- **After about 4-months of training, KataGo had apparently found the best line of play (and numerous equivalent variations)! But that result was of course very uncertain at that time**
- **After 7 additional months of training, KataGo’s opinion was much clearer and opiniated: the best line of play remained unchanged, all alternative lines were either equivalent or significantly worse, so Igo Hatsuyoron 120 problem appeared to be (weakly) solved.**

**Of course, with such a complex problem, we cannot exclude that KataGo overlooked a sequence of moves fundamentally different and better than this preferred line of play!** To limit that risk, I did my best during the training to encourage KataGo to explore alternative moves, with precious advice from David: despite this broad exploration, the best line (and its variations) remained unchanged, so the result seems to be solid.

You will find the link to the final, and probably best, 40b IH120 network, and also a 20b network:

- **This 40b net is able to play the optimal line of play and its main variations perfectly at low visits** (e.g. 100 visits) but also most of the times with a single visit!
- It is also able to play and win easily with low visits if a single error is done by its opponent
- But this net is not an Oracle either: it is less opiniated on some alternative moves leading to unclear positions
- The 20b net seems very strong too, but was just quickly tested, so it may have significantly more weaknesses than the 40b net.

**Approach to train this new KataGo network, dedicated to Igo Hatsuyoron 120's problem:**

The approach used was comparable to the one used for g146: generating selfplay games starting from a large panel of random IH120 positions, mixed with "normal selfplay games" on square and rectangular boards with sizes from 7 to 19.

The aim of the normal selfplay games was to smooth the evolution of the initial network (only trained with normal games) and make it learn progressively IH120 problem and a myriad of positions from this problem, while still intensively playing "normal Go on normal positions"! During the run, the proportion of IH120 positions evolved from around 30% in the beginning of the run to about 70% in the end, and even 90% in the last few weeks.

**Learning process:**

The learning loop was the following:

- Initial net 40b-s5G09 played 70% of normal selfplay games and 30% of games from a panel of 20K IH120 positions.
- As expected, the first games on IH120 positions were of poor quality, as the net was not in a position to understand this extremely complex problem and its unusual positions.
- Once enough selfplay games were available, they were used to train a new net
- This new net was then used to generate additional selfplay on normal and IH120 games. These new games were of a better quality on IH120's positions, as the net improved its evaluation of these positions, thanks to the results of the first batch of games.
- Then, the cycle went on:
    - Latest net was used to generate new IH120 games (of better quality) and normal games
    - These games (mixed with some previous ones) were used to train a new net
    - And so on, for an 11-months loop!

**Hardware used:**

During these 11 months, I used two GPUs:

- An RTX2080 + RTX2070 for the first 2 months (August to end of October 2020)
- An RTX3080 + RTX2070 for the remaining 9 months (until the end of July 2021)

Around 3K selfplay games were generated between each training iteration and, with the 2 GPUs I was using, a new network was produced approximately daily.

The RTX2070 was used for training (training the 40b net was taking about 2 hours and about 40 min for the 20b net), while the RTX 3080 was only used to generate selfplay games.

**Initial panel of positions used:**

With the help of David and Thomas, I initiated the run using a panel of ~20K Igo Hatsuyoron positions they kindly shared with me. These positions were mainly coming from Thomas and a few other experts of the problem, as well as from David when he studied the problem with previous IH120's version of KataGo from g146 run.

More generally, Thomas shared his sgf files and the files accumulated from many experts of the problem, as can be seen there:

http://igohatsuyoron120.de/2015/1177.htm#%C3%BCAcknowledgementsII

For selfplay, I used the initial starting position of IH120 (see Annex A) in about 10% of the games, while the other selfplay games were played from a random position selected among the subsequent moves (i.e. from the panel of 20K positions)

**Updates of the panel of positions (used to start selfplay games):**

During the run, many more positions were added to train the network, with an updated panel of more than 50K positions at the end of the run (from 20K positions in the beginning of the run).

Positions added were mainly coming from:

- Thomas Redecker: I shared with him several intermediate versions of the net, which he used to analyse the problem and evaluate the performance of the nets. When he identified inefficiencies or misevaluations on some positions, he sent them to me: these positions (and variations of these positions) were added in the panel of starting positions for selfplay!

- Myself, manually: I am not an expert of the problem (and just an amateur Go player, around 1 kyu / 1 dan when I was playing on KGS several years ago) but when some variations intrigued me or when the net was hesitant to evaluate a position, I added a panel of variations around these positions.
- Positions generated from a python tool I coded: it systematically explores a given position, dedicating visits to every legal move (in order to identify potential blind spots) and creates a tree of interesting moves (with some parameters to adjust the broadness of the tree). These trees of moves were created for several interesting IH120 positions and were used for selfplay.

All in all, a bit more than 50K positions were eventually used to start selfplay games during the last few months, and helped the net better evaluate the myriad of complexities of this IH120 problem!

**Measure of progression:**

To measure the progression of the net on IH120's problem, I benchmarked regularly the latest net against the older g146 net, starting from a panel of random IH120 positions selected among the 20K initial positions.
The benchmark panel of 20K positions was not extended with additional positions along the run, as these new positions were not studied by g146: including them in match games would have introduced an obvious bias to the benchmark.
All the match games were played with Tromp-Taylor rules, as the older g146 net was only able to play this set of rules.
Having benchmarked many nets and Go engines over the previous years (LZ, Elf and KataGo), I noticed that results at low visits are quite representative of results at higher visits, and of course much faster. With only 2 GPUs, I decided to benchmark the new net with only 25 visits vs. g146 with only 100 visits. The few verifications I did with 10x more visits for both nets gave comparable results, so I sticked to these low-visits tests for my benchmarks.

**ELO results vs previous g146 run:**

When I really began to benchmark the new nets at the end of September, after 1 month of training, these new nets were already quite competitive in strength against the older g146 net! So, I decided to benchmark the new net with 4x less visits.

By December 2020, after about 4 months of training, the new nets were able to crush g146 on about any IH120 positions with the same number of visits. Even with extremely unbalanced visits (e.g. only 100 visits vs. 100K visits for the old g146 net!), the new net was either able to take the lead, thanks to a few notable improvements on its understanding of the problem, or to draw the position.

Here is a graph of the evolution of g170 vs. g146 (with generally 1000 games per dot, so the standard variation is around 12 Elo points). The yellow and orange dots are moving-averages.

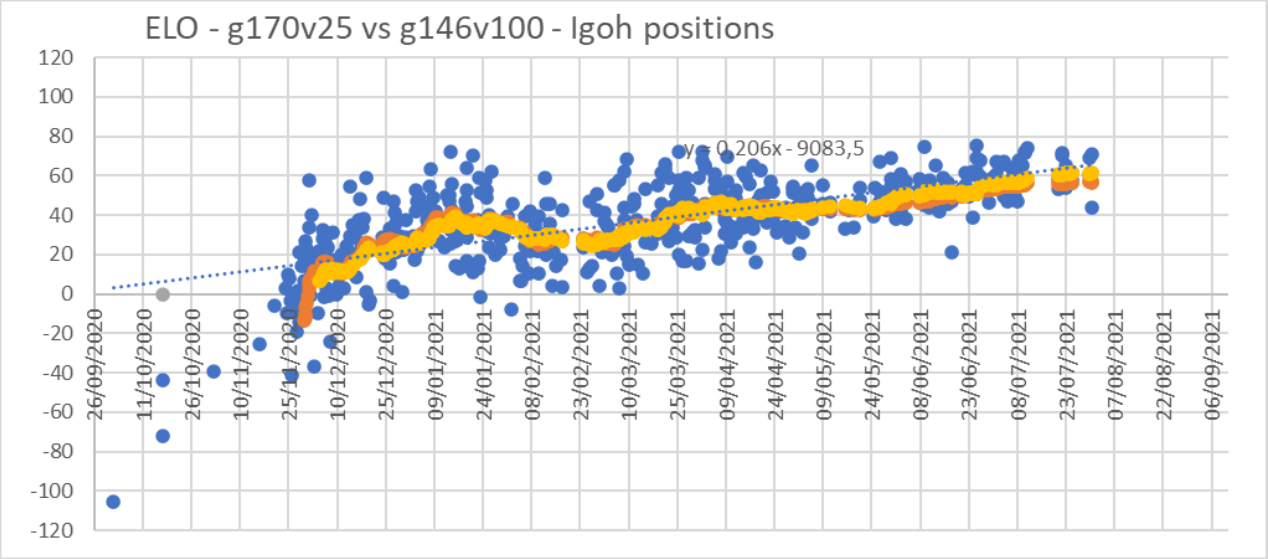

As can be seen on this graph, the new g170 net with 25 visits was already close from g146 with 100 visits at the end of September 2020, just around 100 Elo behind the previous net despite 4x less visits.
By December, it began to be better than g146 (still using 4x less visits).
Then, it continued to improve slowly, to eventually reach about +60 Elo vs. g146 (still 4x less visits).

A few comments on this benchmark: it is only a small part of the story, as the match process increasingly disadvantages the new best net:

- In the match process, when a match game is launched on a random IH120 position, both nets give their evaluation of the fair komi and the game begins with the average of these 2 komis and random colours. When the weakest misevaluates the komi and get the favourable averaged komi, it may win the position quite easily despite being weaker, limiting the ELO difference
- Well known IH120 positions lead to draws: once the nets improve, the proportion of well known positions increase, mechanically limiting the measured Elo difference between g170 and g146. And among the panel of 20K positions, a significant percentage of them are quite end-game positions played perfectly (and thus leading to draws)

I am not able to assess the impact of these issues, but it plays for sure a role to explain the limited Elo improvement of the new KataGo over the last 6 months *on this benchmark*, while the net clearly improved a lot in its understanding of the problem (as will be illustrated thereafter).

**Status of Igo Hatsuyoron 120 problem at the end of 2020 (after 4-months of training):**
Around the end of 2020, the impossible was possibly within reach: **the new nets were apparently converging to a reliable solution of the Igo Hatsuyoron 120 problem!**
Quite surprisingly, the result was the same for Chinese, Japanese, Tromp-Taylor and New-Zealand rules: without komi and assuming White started the game and no prisoners, White wins by a single point!
**So, at the end of 2020, the proper komi was apparently -1, leading to a draw with perfect play.**

At that point, it was not reasonable to expect the problem was "solved", but that was becoming a credible possibility. However, many positions were still misevaluated, as Thomas managed to find, and the confidence of the network was still quite low on several positions.

**Additional training in H1 2021:**
From the end of 2020 to July 2021, the Elo gains vs. g146 were limited (about 60 Elo) but new networks improved significantly their understanding of the problem, thanks to selfplay and regular interactions with Thomas Redecker.
With Thomas' crucial help, I was able to provide multiple new interesting positions to KataGo's panel of positions during H1 2021, improving its understanding of these difficult lines of play and forcing it to reassess some badly understood positions. You will find here a detailed description of some of the improvements done during that period of time, with the very active help from Thomas:
http://igohatsuyoron120.de/2015/0513b.htm#%C3%BCVariationsOurSolutionIIIB

With some technical advices from David, I also extensively used the "Hintpos" functionality of KataGo: this tool analyses games and suit of moves (from the panel of IH120 positions and, occasionally, from deep games played by KataGo) and identifies the moves were its evaluation changes dramatically. These changes of evaluations identify "blind spots" for the net: they are then selfplayed and forked several times. Thanks to this process, the net can update its inadequate evaluations and accelerate its learning process.

**Final status of Igo Hatsuyoron 120 problem:**

**The main line of play and proper -1 komi were unchanged from December 2020 to the end of the run in July 2021. Unless a dramatic move has been overlooked, which cannot be excluded:**

**Igo Hatsuyoron 120 is probably *weakly* solved!**
("*weakly* solved" means the optimal line and main variations appears to be solid,
while alternative lines are not perfectly played)

**More specifically, it means that *if the best line is really optimal*, which seems credible:**

- **With a few dozens of visits, KataGo can play the best line and its main variations perfectly. And a single visit is almost always enough for that too!**
- **When exploring some less common variations, a few hundred visits are enough to play virtually perfectly**
- **If one or two significant mistakes are done, leading to less explored lines, a few thousands visits are generally enough to get an extremely solid evaluation of the position, but the score evaluation may be slightly wrong (e.g. estimating a +2 score vs +1 in the end).**
- **In some cases, mainly after 2 or more mistakes, KataGo's evaluation may become more biased:** its score evaluation may be offset by one or a few points, even with deep visits (100K or more), possibly until late in the end game. In several cases, the net can estimate a position is for example won with a +1 score, while it only realises lately that the game is drawn.
- Eventually, if explored moves lead to fundamentally different variations of the problem, KataGo is faced with a more "normal" game of Go. In that case, KataGo is still a formidable and super-human opponent, but about anything can happen!

**How did we arrive to this tentative "conclusion" that the problem appears weakly solved?**
Several elements reinforce this feeling:
1) By construction, KataGo's intensely explores alternative lines of play in selfplay: initial komi has some random uncertainty so, when the komi is estimated unfair by a side, this side has a strong incentive to generate complex and uncertain positions to get some chances to win! Moreover, the preferred move is often not played in order to encourage exploration, some games are forked, some totally random moves are occasionally played to force exploration, etc. All this contribute to a large exploration of alternative moves: for a given position, if an alternative move is better than the best known move at a given time, it should eventually be discovered!
2) In the end, despite this intense exploration of alternative lines of play, none of them led to a change of the "preferred solution", the variations were either unfavourable or equally favourable, so **the best line** (and its variations and permutations) **remained unchanged**.
3) In order to gain more confidence on this result, we tried to explore as intensively as possible alternative moves:

- We increased random exploration (with significant help from David, of course).
- I added a systematic exploration of thousands of opening moves.
- I added thousands systematic variations of several complex positions.
- I generated several deeply asymmetric games (eg 100K visits for Black vs. 100 visits for White, with a komi favouring White, and the opposite) to give more opportunities to explore new hard-to-discover lines, and used these positions for selfplay.
- I used the hintpos tool regularly to speed up KataGo's learning process, especially when it badly estimates some positions

All these attempts failed: not a single move modified the "best line / preferred line". Hence, this failure to find alternative lines of play increases the confidence in the best line of play!
Of course, there are for sure many more unexpected ideas in Igo Hatsuyoron's problem and KataGo is not an Oracle, but these months of training by this superhuman entity are quite convincing: the problem appears weakly "solved"!

**KataGo's degree of confidence and proper komi:**

KataGo is very confident about IH120 initial position: the win/draw/loss initial status is extremely dependant on the komi value! This table illustrates the high degree of confidence of KataGo:

| Black's winrate (after 1M visits) | Japanese rules | Chinese rules | Tromp-Taylor rules | New-Zealand rules |
|---|---|---|---|---|
| **Komi -1.5** (0.5 advantage for Black) | 95.4% | 94.3% | 93.7% | 94.7% |
| **Komi -1** (fair komi) | 49.0% | 47.8% | 47.8% | 48.4% |
| **Komi -0.5** (0.5 advantage for White) | 2.3% | 1.0% | 1.1% | 0.8% |

Some comments on these results:
1) As can be seen, there is absolutely no ambiguity on the proper komi: -1 is fair, while a 0.5 difference leads to extreme differences: even with the smallest possible difference of 0.5 komi, KataGo estimates the game could/should be resigned by the losing side from the beginning!!!
2) All main rule sets have the same fair komi of -1 (assuming White started the game and no black prisoner). KataGo is also trained with button Go (a 0.5 points advantage for the first side to pass) and when button is activated, the proper komi is -1.5 points.
3) The initial IH120 position appears to be a bit easier to play for White: its initial win rate is around 52%. This is also confirmed by the results with different komi values:

- When the 0.5 komi difference favours White, its win rate is ~98.5%.
- When the 0.5 komi difference favours Black, its win rate is only ~94.5%, which is high but significantly less confident than the 98.5%.

Another illustration of KataGo's confidence comes from very unbalanced games: KataGo playing with a single visit (and White or Black colour) is able to get a draw when facing KataGo with 100K visits! Two examples of such unbalanced games are presented here:


Igoh-Bv1-Wv100K - 2 Draws.sgf

**Remaining identified difficulties for KataGo:**

This KataGo network still has several weak points, which you'll be in a position to explore by yourself!
The main identified difficulties remaining for KataGo on this IH120 problem are related to badly explored positions: by construction, it regularly explores suboptimal lines but this exploration is more limited each time a mistake is added: it explores regularly positions with 1 mistake, sometimes positions with two mistakes, but rarely positions with three or more mistakes. Hence, these lines are badly explored and KataGo's evaluation are less accurate in this case.
Thomas found positions where the net thinks the score is N for a long time, before seeing it is in fact N+1 or N-1. When N=1 or -1, this inaccuracy can change the outcome of the game.

**Where to get more information on Igo Hatsuyoron 120 problem?**

Thomas Redecker has built an extremely detailed web site with thousands of variations of the problem, explaining its main complexities and KataGo's best solution: http://igohatsuyoron120.de
This is very probably the best available source of information on this marvellous problem!

As we cooperated a lot in 2021, Thomas had a regular access to these improving KataGo networks: this allowed him to make some deeper analysis of the problem and, reciprocally, it allowed me to improve KataGo thanks to the additional positions he sent me (which I added in the panel of positions for selfplay).

Many positions deeply studied with the help of this new 40b network have been added on his site since the beginning of 2021!

**Where to find details on the (probable) solution of Igo Hatsuyoron 120 problem?**

The identified and probable solution is available here: http://igohatsuyoron120.de/2015/0009.htm, see "**Igo Hatsuyorōn's Problem 120 (2021), Main Semeai**".
And the direct link is there: http://igohatsuyoron120.de/2015/0058.htm#%C3%BCVOSIII2021i1

**Where to find these new g170 KataGo's nets?**

The latest and best 40b net is available here: https://magentacloud.de/s/8i9BmTSKjf2ipCr
And a 20b network was also trained (with the same inputs as the 40b net) and may be useful for people with slower computers and/or without a GPU: https://magentacloud.de/s/c55fdK2ZJoiwNFZ

**What next for KataGo and Igo Hatsuyoron 120 problem?**

This KataGo network probably plays perfectly the main lines of play, but its degree of certitude on alternative lines can be improved for sure!
A new run with a larger network (e.g. a 60b KataGo network) for a longer period of time and with more GPUs would for sure increase the confidence on this result. David also added several notable improvements in KataGo since then, such as an improved MCTS algorithm, a Graph-search version (helping KataGo in situations when there are many permutations, which is the case for IH120) and is also experimenting new types of (hopefully) stronger networks. All these improvements should help in a future run and, who knows, possibly improve the best line of play of IH120!
If any anyone is willing to spend time and energy to set up a new run, I'll be happy to provide some technical help to set up the infrastructure.

**As a conclusion, I hope all the players fascinated by Igo Hatsuyoron 120 will find pleasure exploring this marvellous problem, by themselves and/or with KataGo's help.**
**Enjoy!**

Best regards,

**Karl Desfontaines, February 2022**

To contact me, I'm "Friday9i" on:

- Discord: https://discord.com/channels/417022162348802048/583775968804732928
- Life in 19x19: https://lifein19x19.com/viewtopic.php?f=18&t=16995

# ANNEXES

**Annex A: Initial "Igo Hatsuyoron 120" position:**

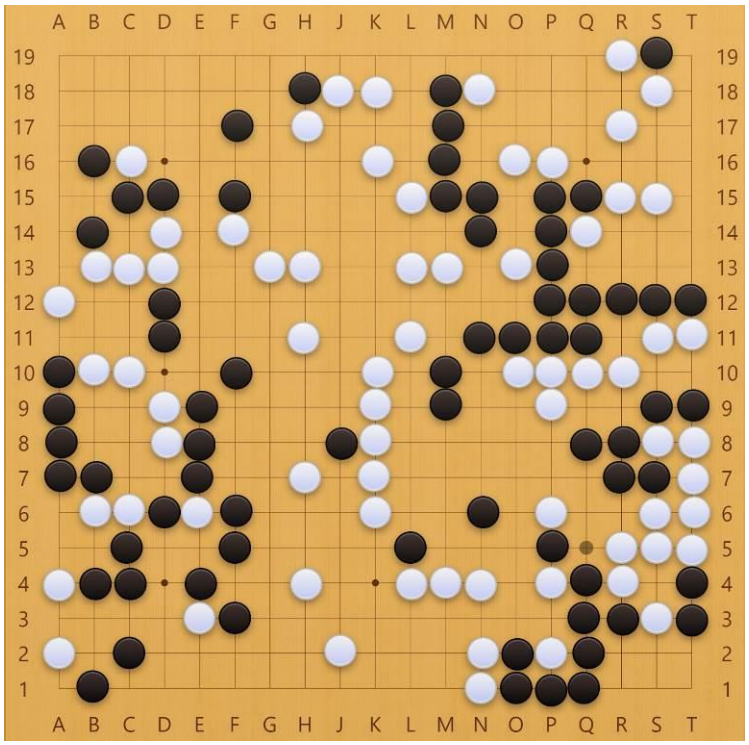


Sgf file:

**Black to play and win!**

**Precisions:**

- Komi is -1.5 (with komi = -1, Black can only get a draw).
- Any rules *without button* can be used (eg Japanese, Chinese, Tromp-Taylor or New-Zealand rules).
- Assumption: White started the game and there is no prisoner.

**In this position:**

- The only winning move for Black is Q5!
- The solution has around 205 moves and there are way too many lines of play to count them!!!
- In the end, with best play from both players, Black wins by 0.5 points!
- And KataGo is able to play all of these variations with low visits, and most of the times, a single visit is enough.

Amazingly, every single stone on the board is crucial for the problem (even removing the isolated and lost Black stone at S19, not directly involved in the best line, has an unexpectedly complex effect on the solution of the problem, thanks to an analysis done by Thomas with KataGo!)

**Annex B: two examples of fast 2500 visits games with proper -1 komi, leading to draws:**

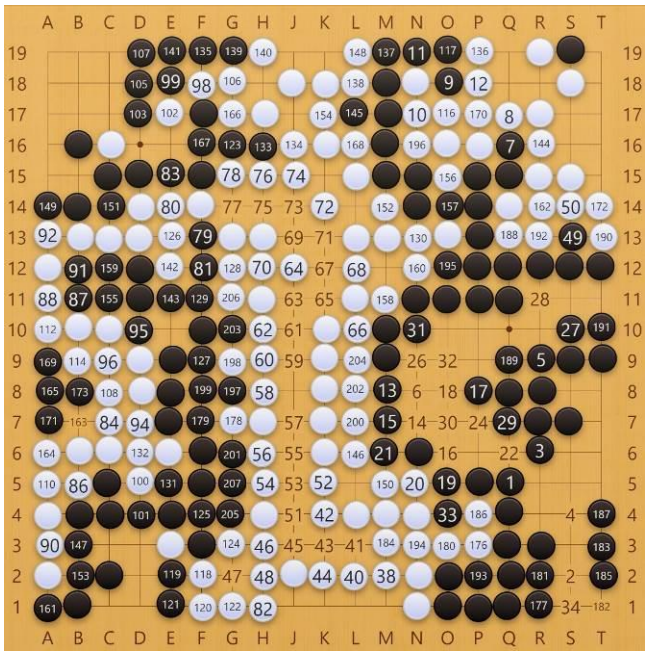

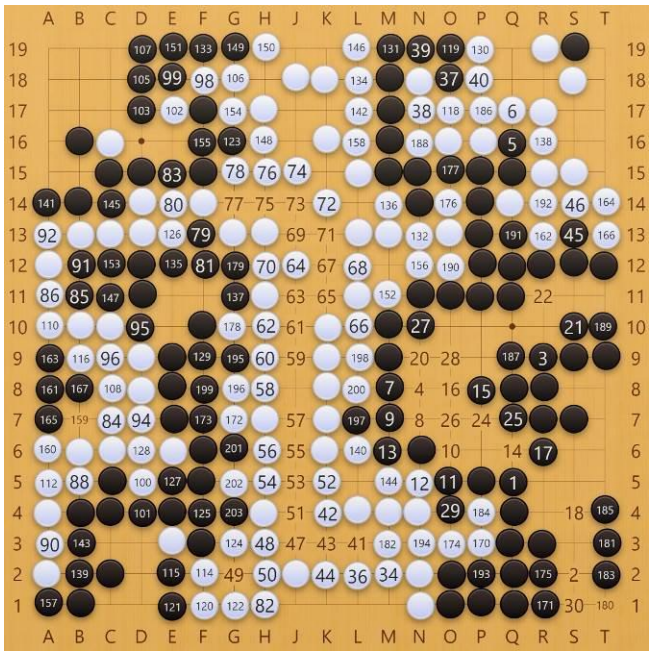

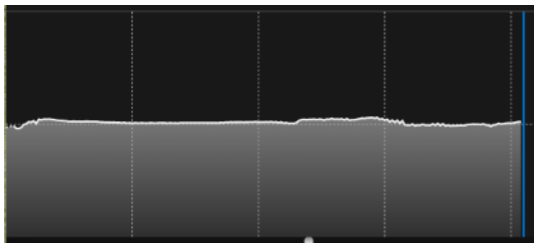

Notes: the win rates during these games are almost totally flat. KataGo has virtually no doubt that the games will be drawn, despite huge fights, big captures and a total of ~80 captured stones! There are some significant differences between the 2 games in the order of moves played, eg stone O18 played at move 9 vs move 37, and some differences in final territories (in the second game, Black has 2 more points around E10 and White 2 more points around J17). But the results are of course identical: draws!

And here are nine games played with Japanese rules and one with New-Zealand rules and some wideRootNoise (to encourage diversity) up to 0.2. All the games are quite comparable and end in a draw:

**Annex C: the Very early Guzumi** (B S13 on move 39) **followed by White G2**:

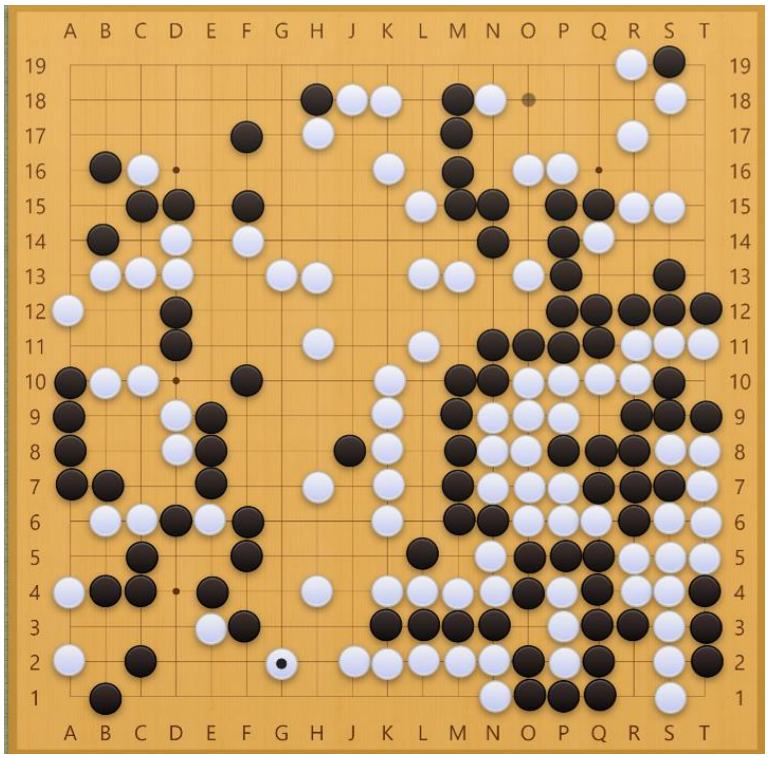

This Very Early Guzumi S13 and the White answer at G2 lead to a difficult game for White, while the simple White answer at S14 is a "classical" variation of the main solution, leading to an "easy" draw (easy for KataGo, at least).

Here is an example of the final result of one game after this very early Guzumi and the G2 answer. As can be seen, the final result is very different from the classic IH120 solution:

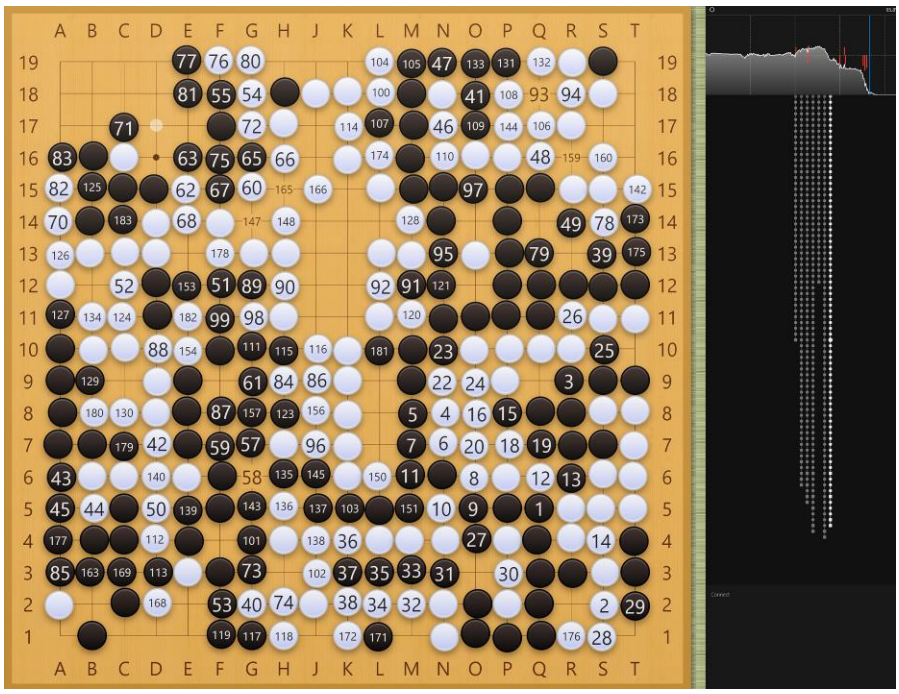

It's a more "normal" game everywhere on the board, excepted in the bottom-right quarter, where a large hanezeki remains (the 4 black stones including K3) and the resulting status of the bottom right groups up to R11 is complex. This game was lost by White but, after the early guzumi and G2 answer, the position is not solved by KataGo (as it did little training on this sequence): some games are won by White while others are won by Black. It should be noted that the uncertainty on this position does not

change IH120 solution at all, as both players can avoid it: Black can avoid the early Guzumi and, if it plays it, White can simply answer with S14 if it prefers, leading to the standard solution.

Here are some games played with 1000 visits from this position:

IgoH-VeryEarlyGuzumiVariation-Wins&Lc