March 2016

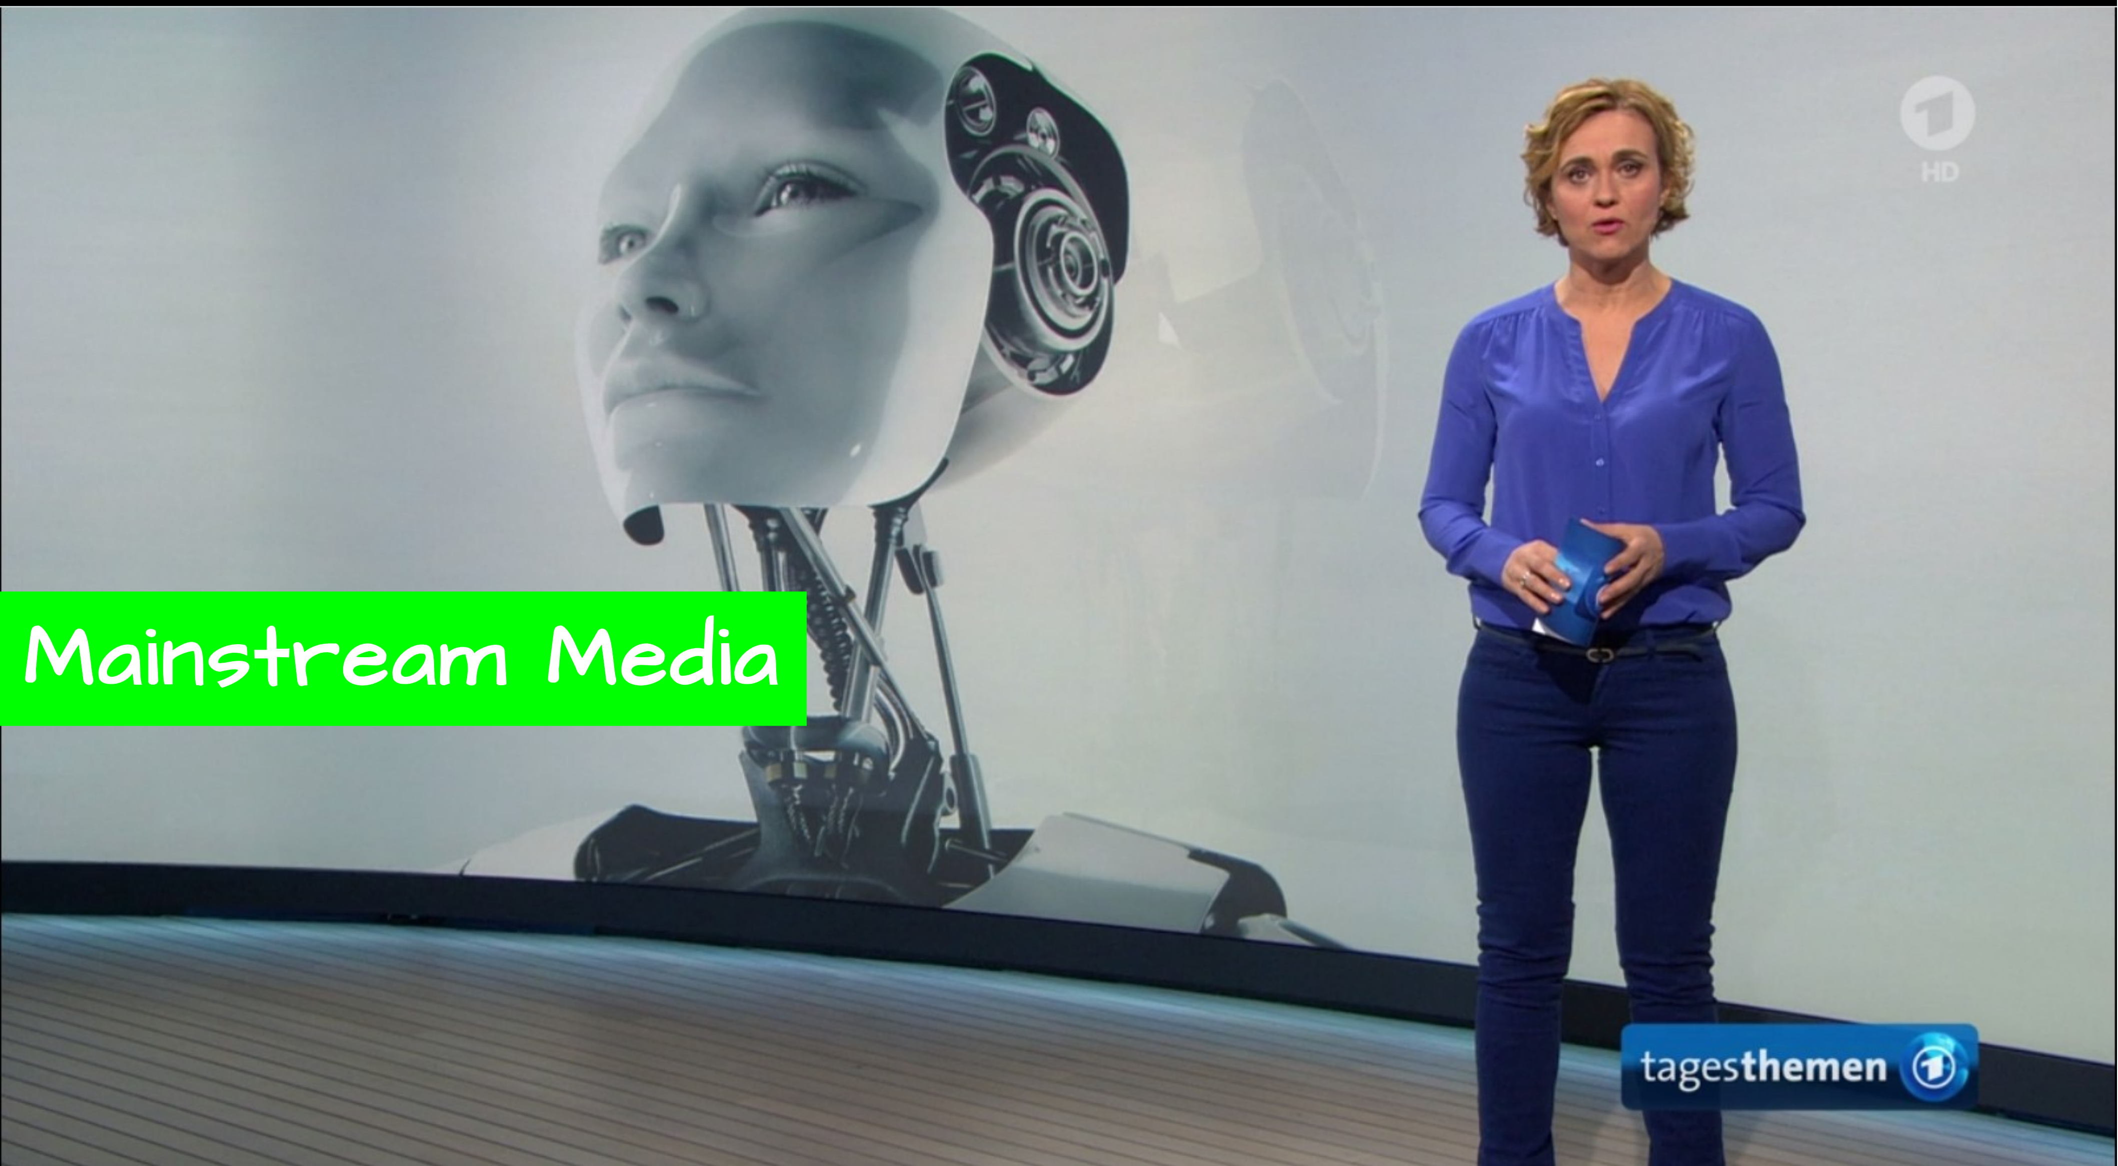
Mainstream Media
HD
tagesthemen

AlphaGo


1997
IBM
RS/6000 SP


ING
`2000 "應氏杯" 世界電腦圍棋錦標賽
`2000 ING CUP WORLD COMPUTER GOE CHAMPIONSHIP
主辦: 應昌期圍棋教育基金會
承辦: 貴陽市棋類協會
貴陽應氏圍棋活動中心
(1985-2000)
1
3
THREE

5d win 1998

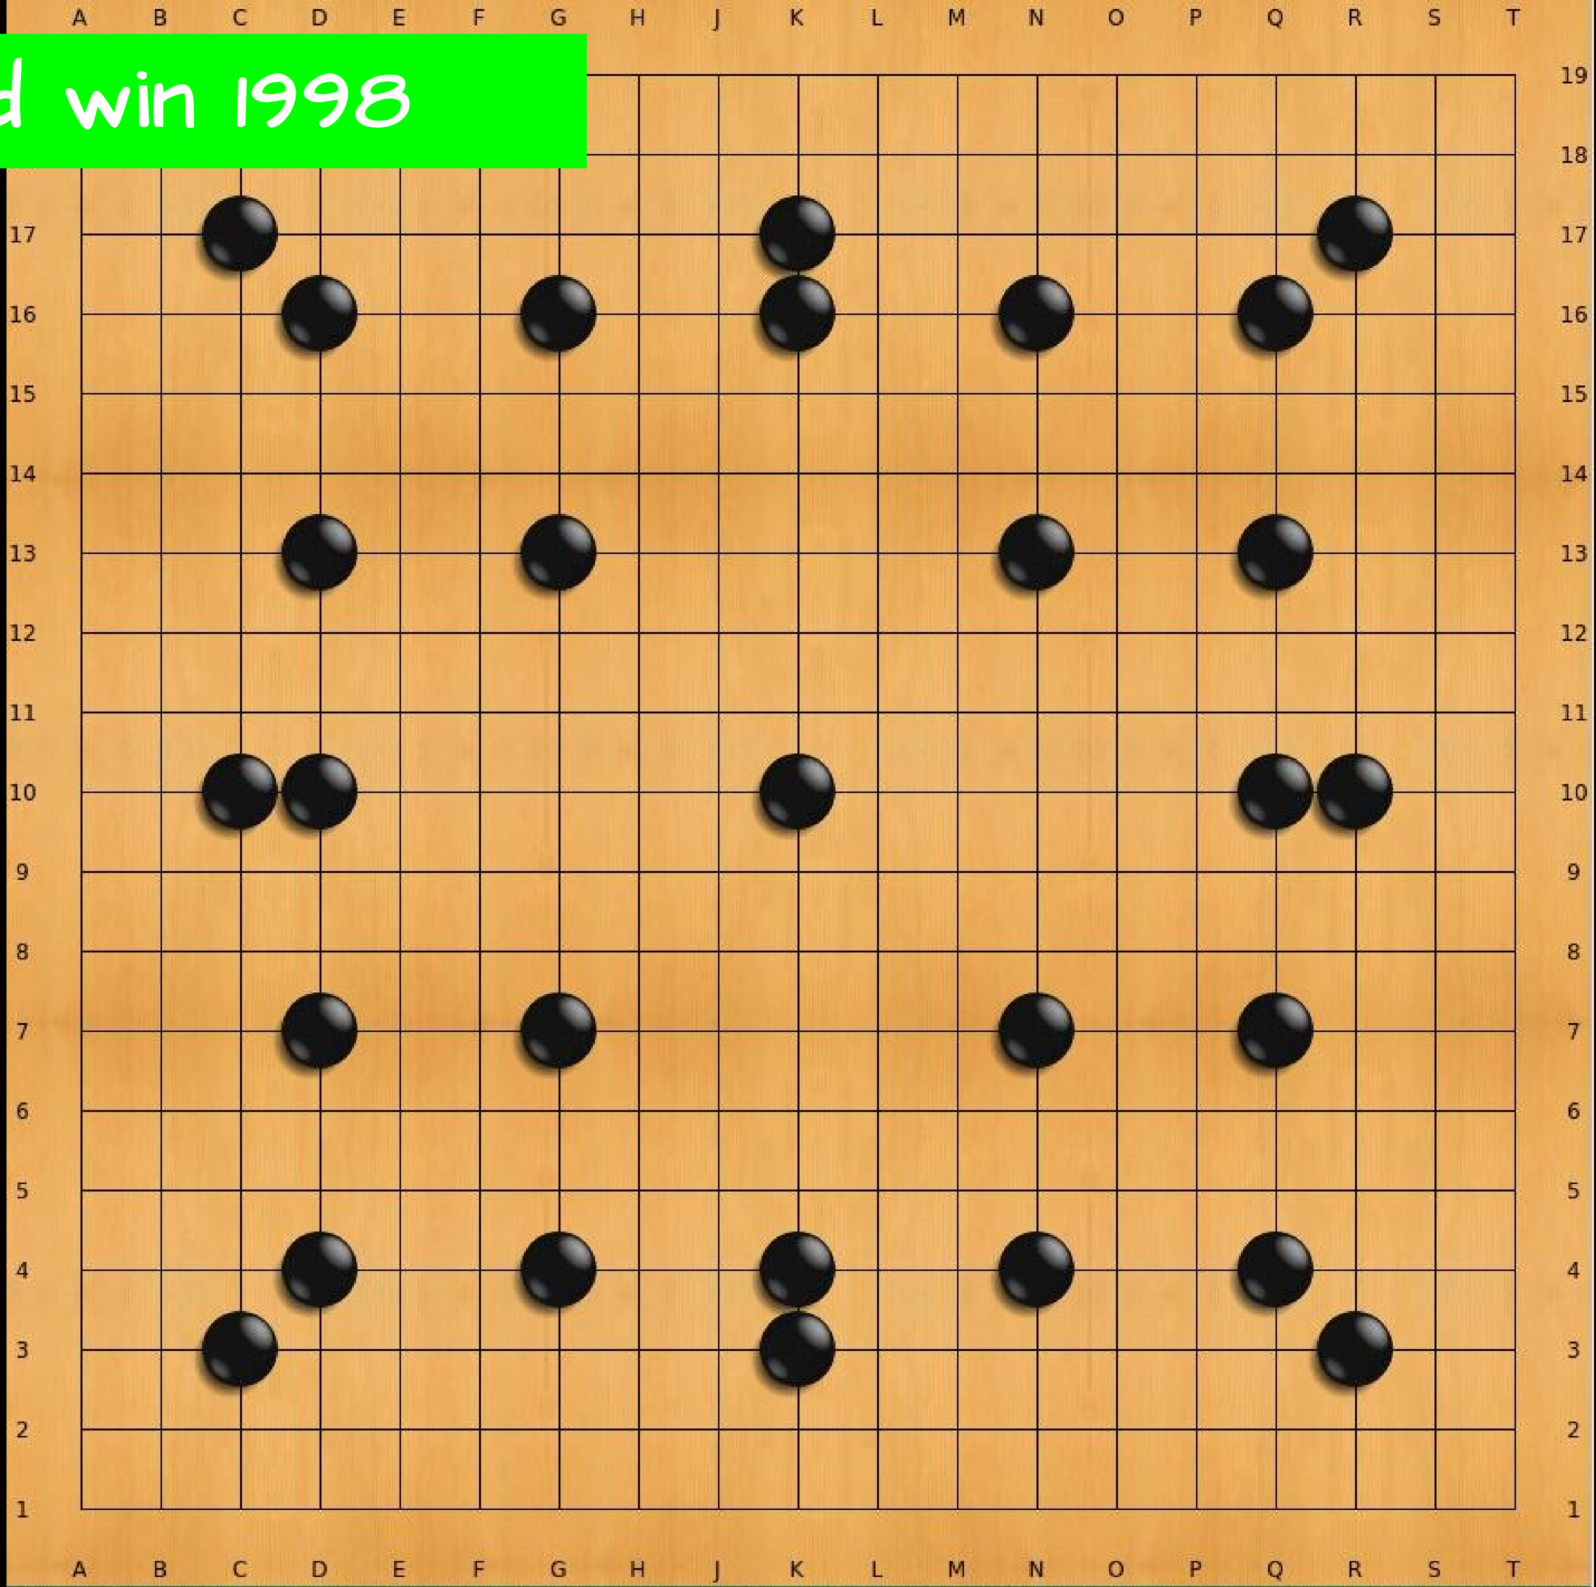

October 2015

January 2016

This is the first time that a computer program has defeated a human professional player in the full-sized game of Go, a feat previously thought to be at least a decade away.

Silver, D. et al., 2016. Mastering the game of Go with deep neural networks and tree search. Nature, 529(7587), p.484-489.

November 2015

Move Evaluation in Go Using Deep Convolutional Neural Networks

Chris J. Maddison
University of Toronto
cmaddis@cs.toronto.edu

Aja Huang[1], **Ilya Sutskever**[2], **David Silver**[1]
Google DeepMind[1], Google Brain[2]
{ajahuang,ilyasu,davidsilver}@google.com

Abstract

The game of Go is more challenging than other board games, due to the difficulty of constructing a position or move evaluation function. In this paper we investigate whether deep convolutional networks can be used to directly represent and learn this knowledge. We train a large 12-layer convolutional neural network by supervised learning from a database of human professional games. The network correctly predicts the expert move in 55% of positions, equalling the accuracy of a 6 dan human player. When the trained convolutional network was used di-

RUBYCONF
SAN ANTONIO | 2015

What did AlphaGo do to beat the strongest human Go player?

Tobias Pfeiffer

@PragTob

pragtob.info


STRANGE
GROUP

Go

Computational Challenge


Monte Carlo Method

Neural Networks

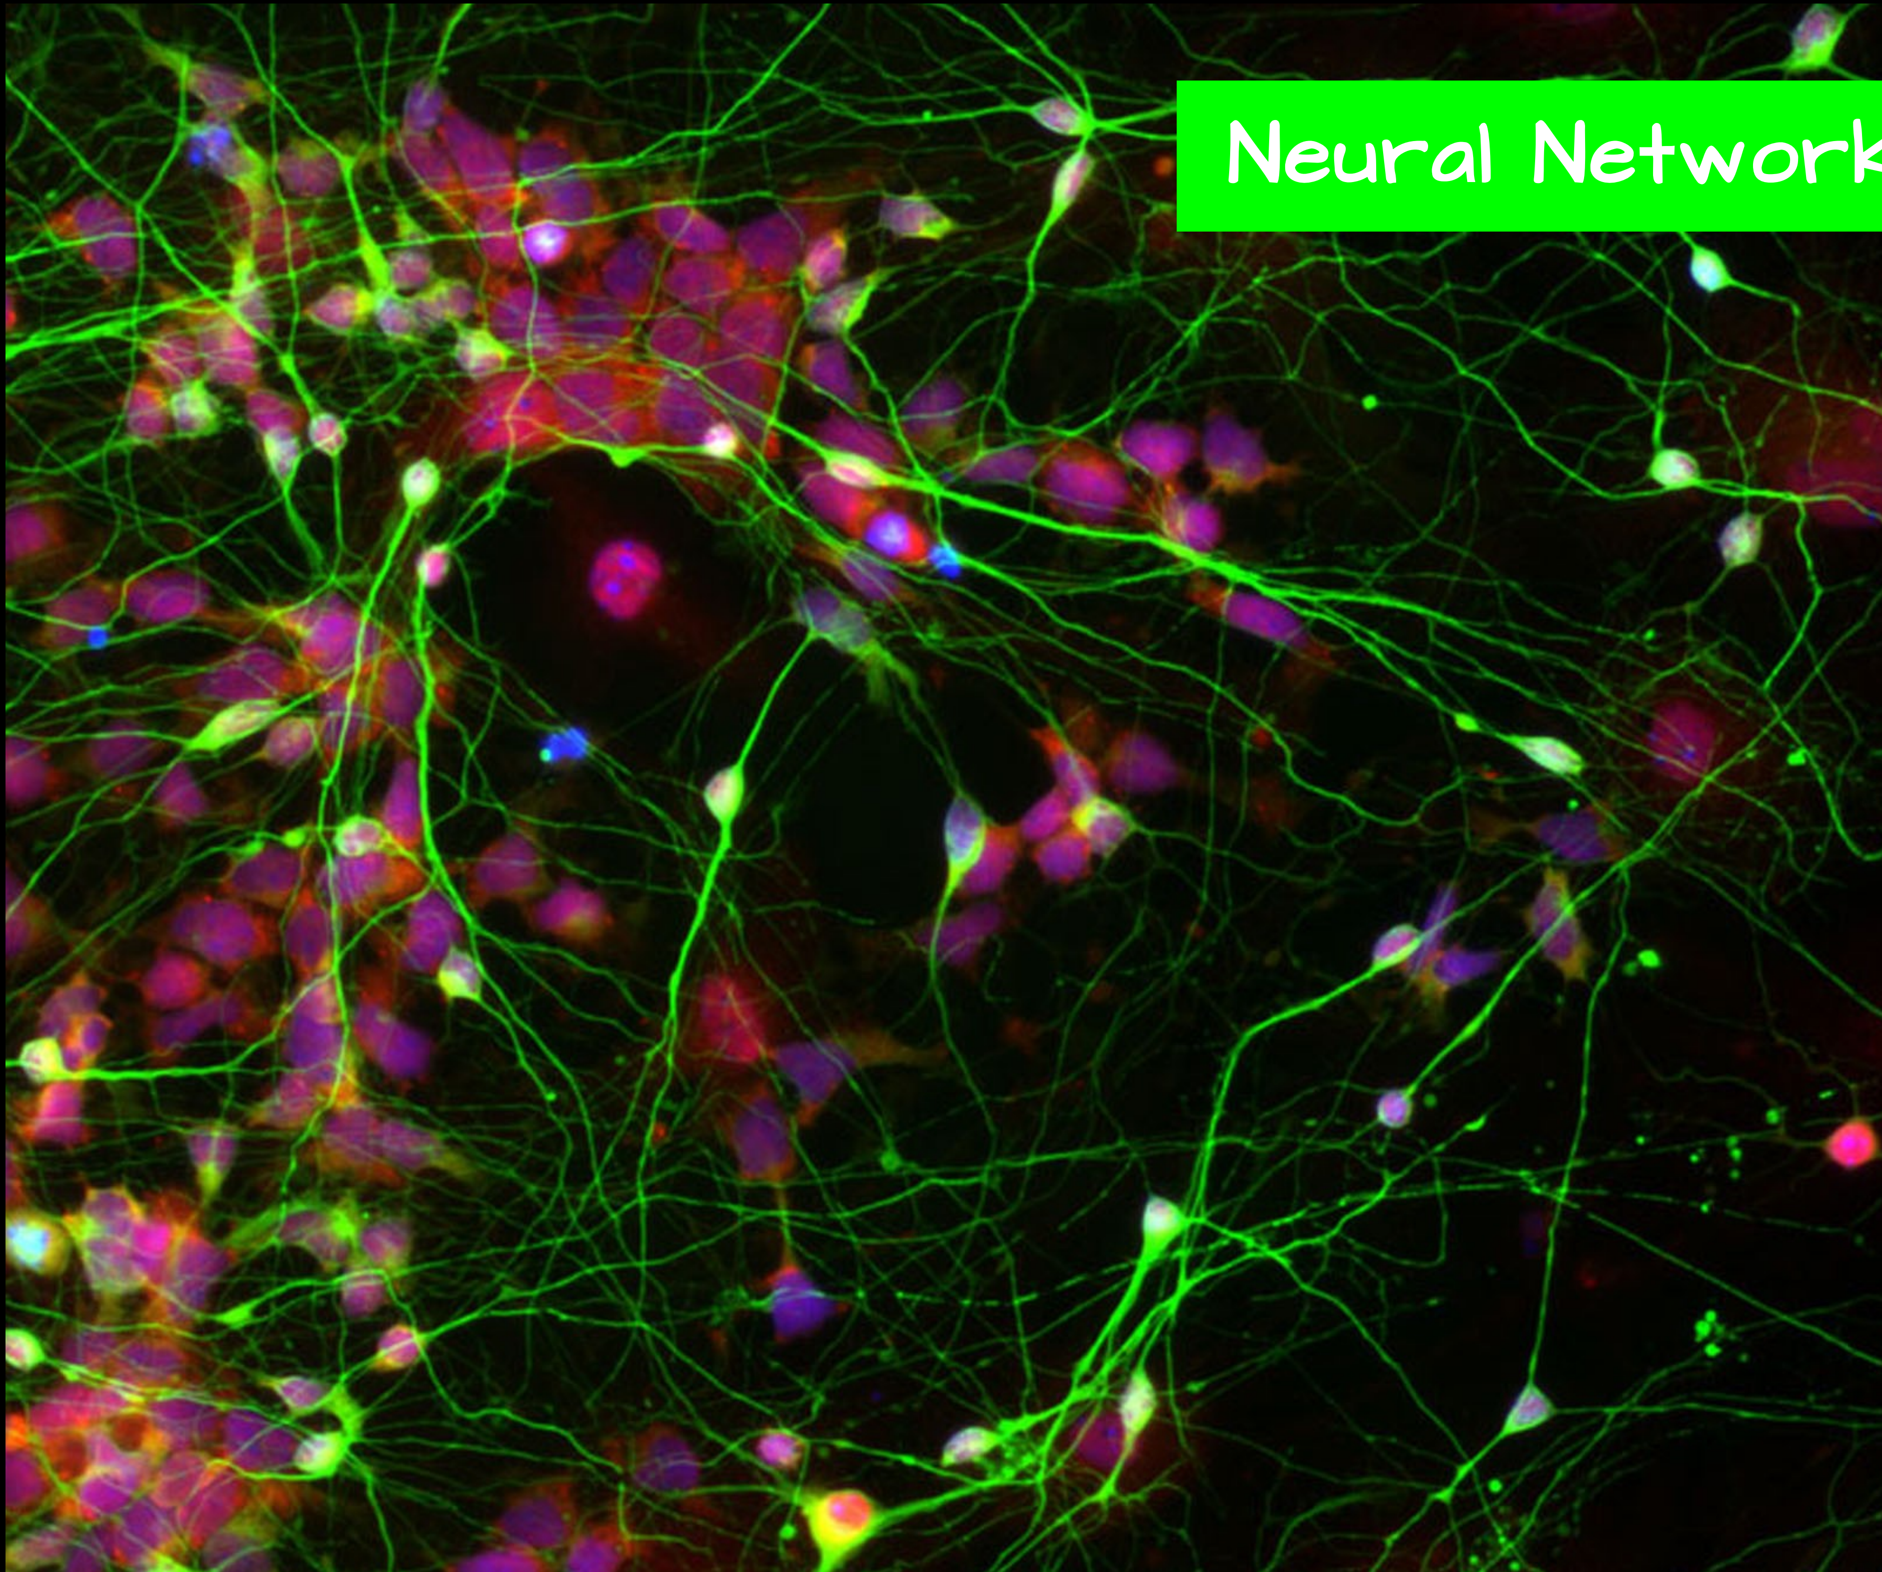

Revolution with Neural Networks

AlphaGo

What did we learn?

Go

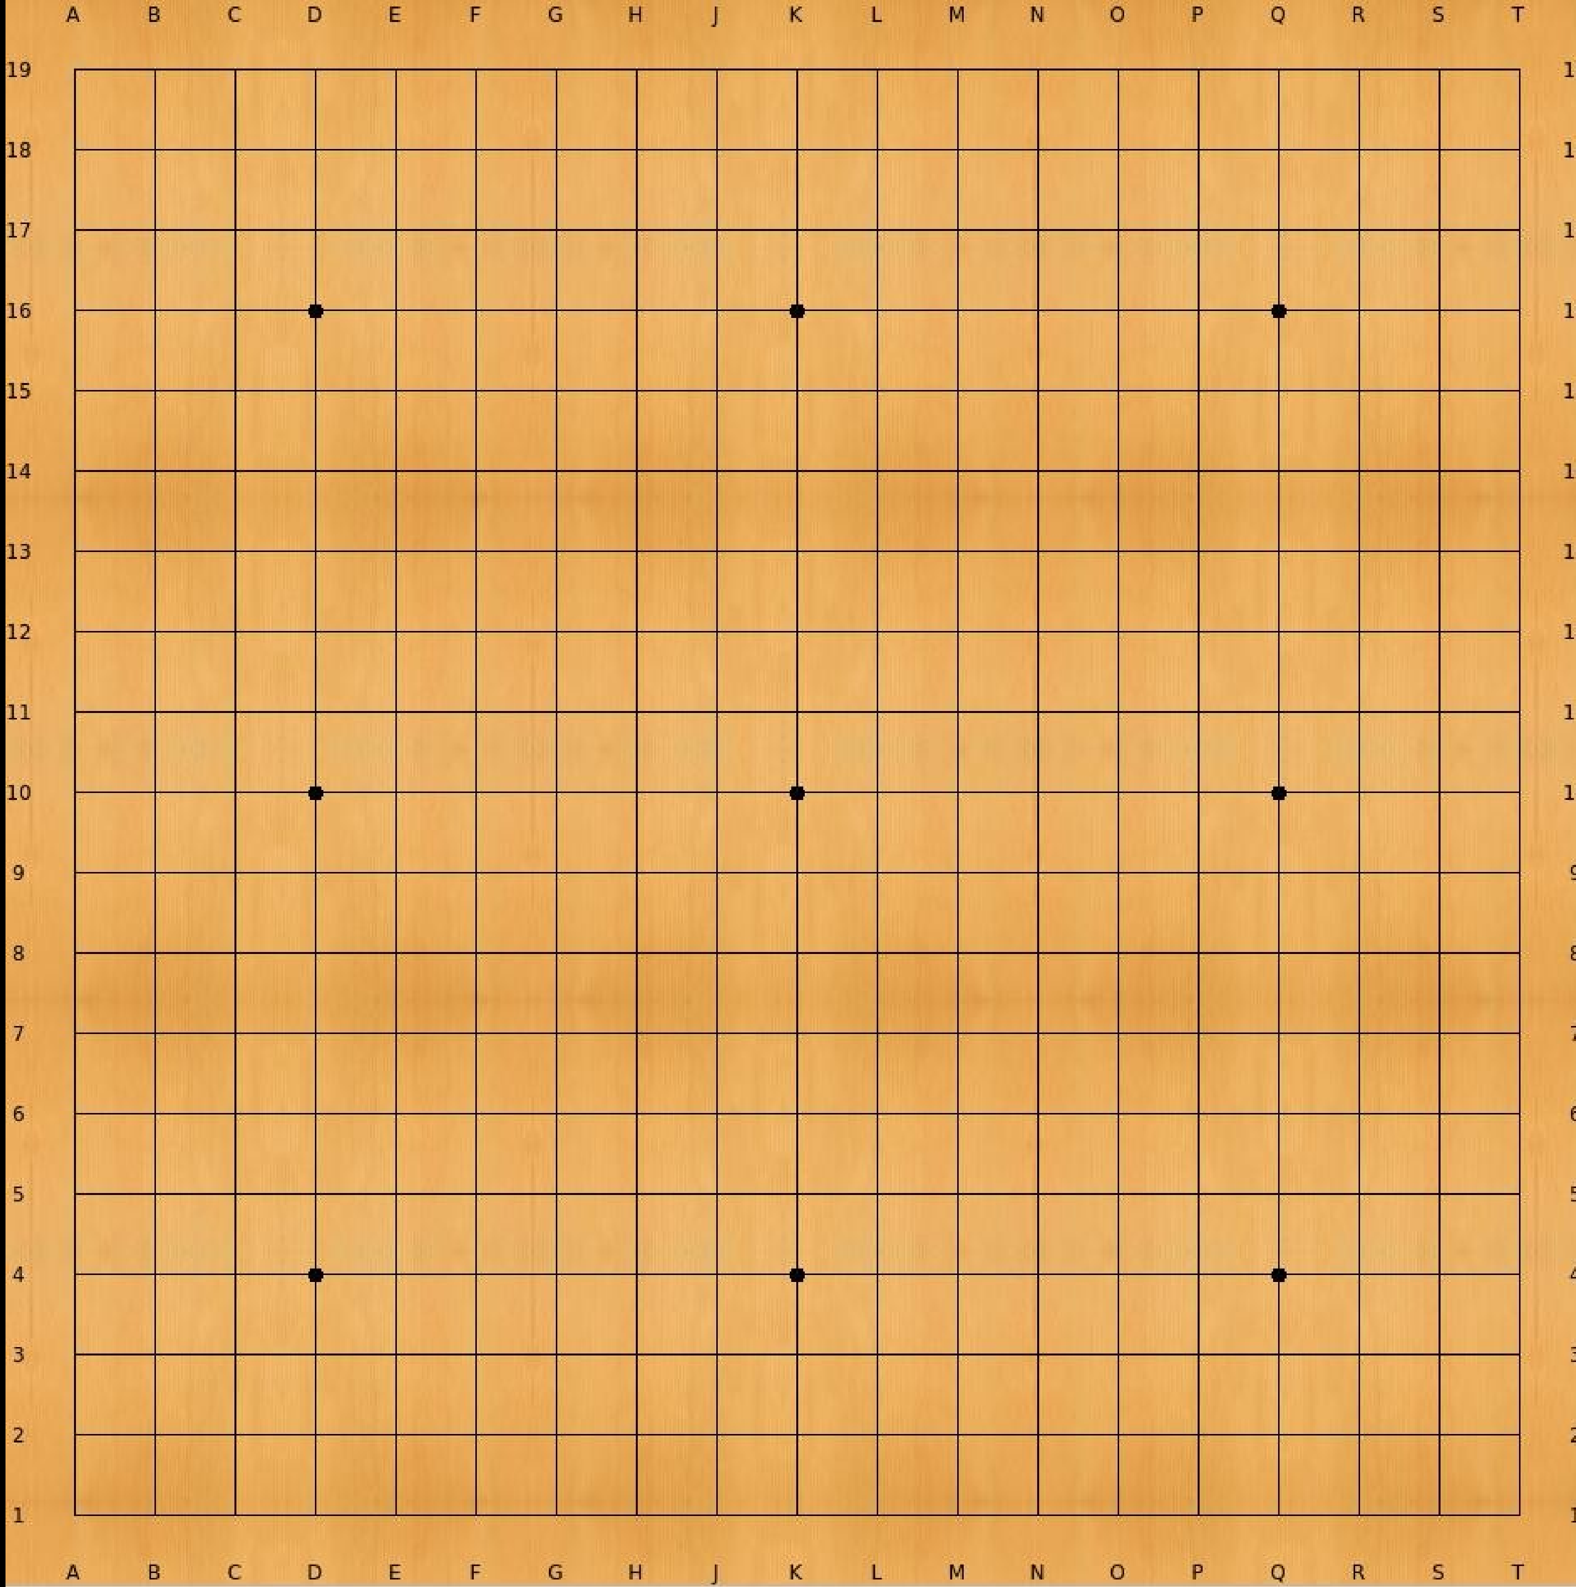
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

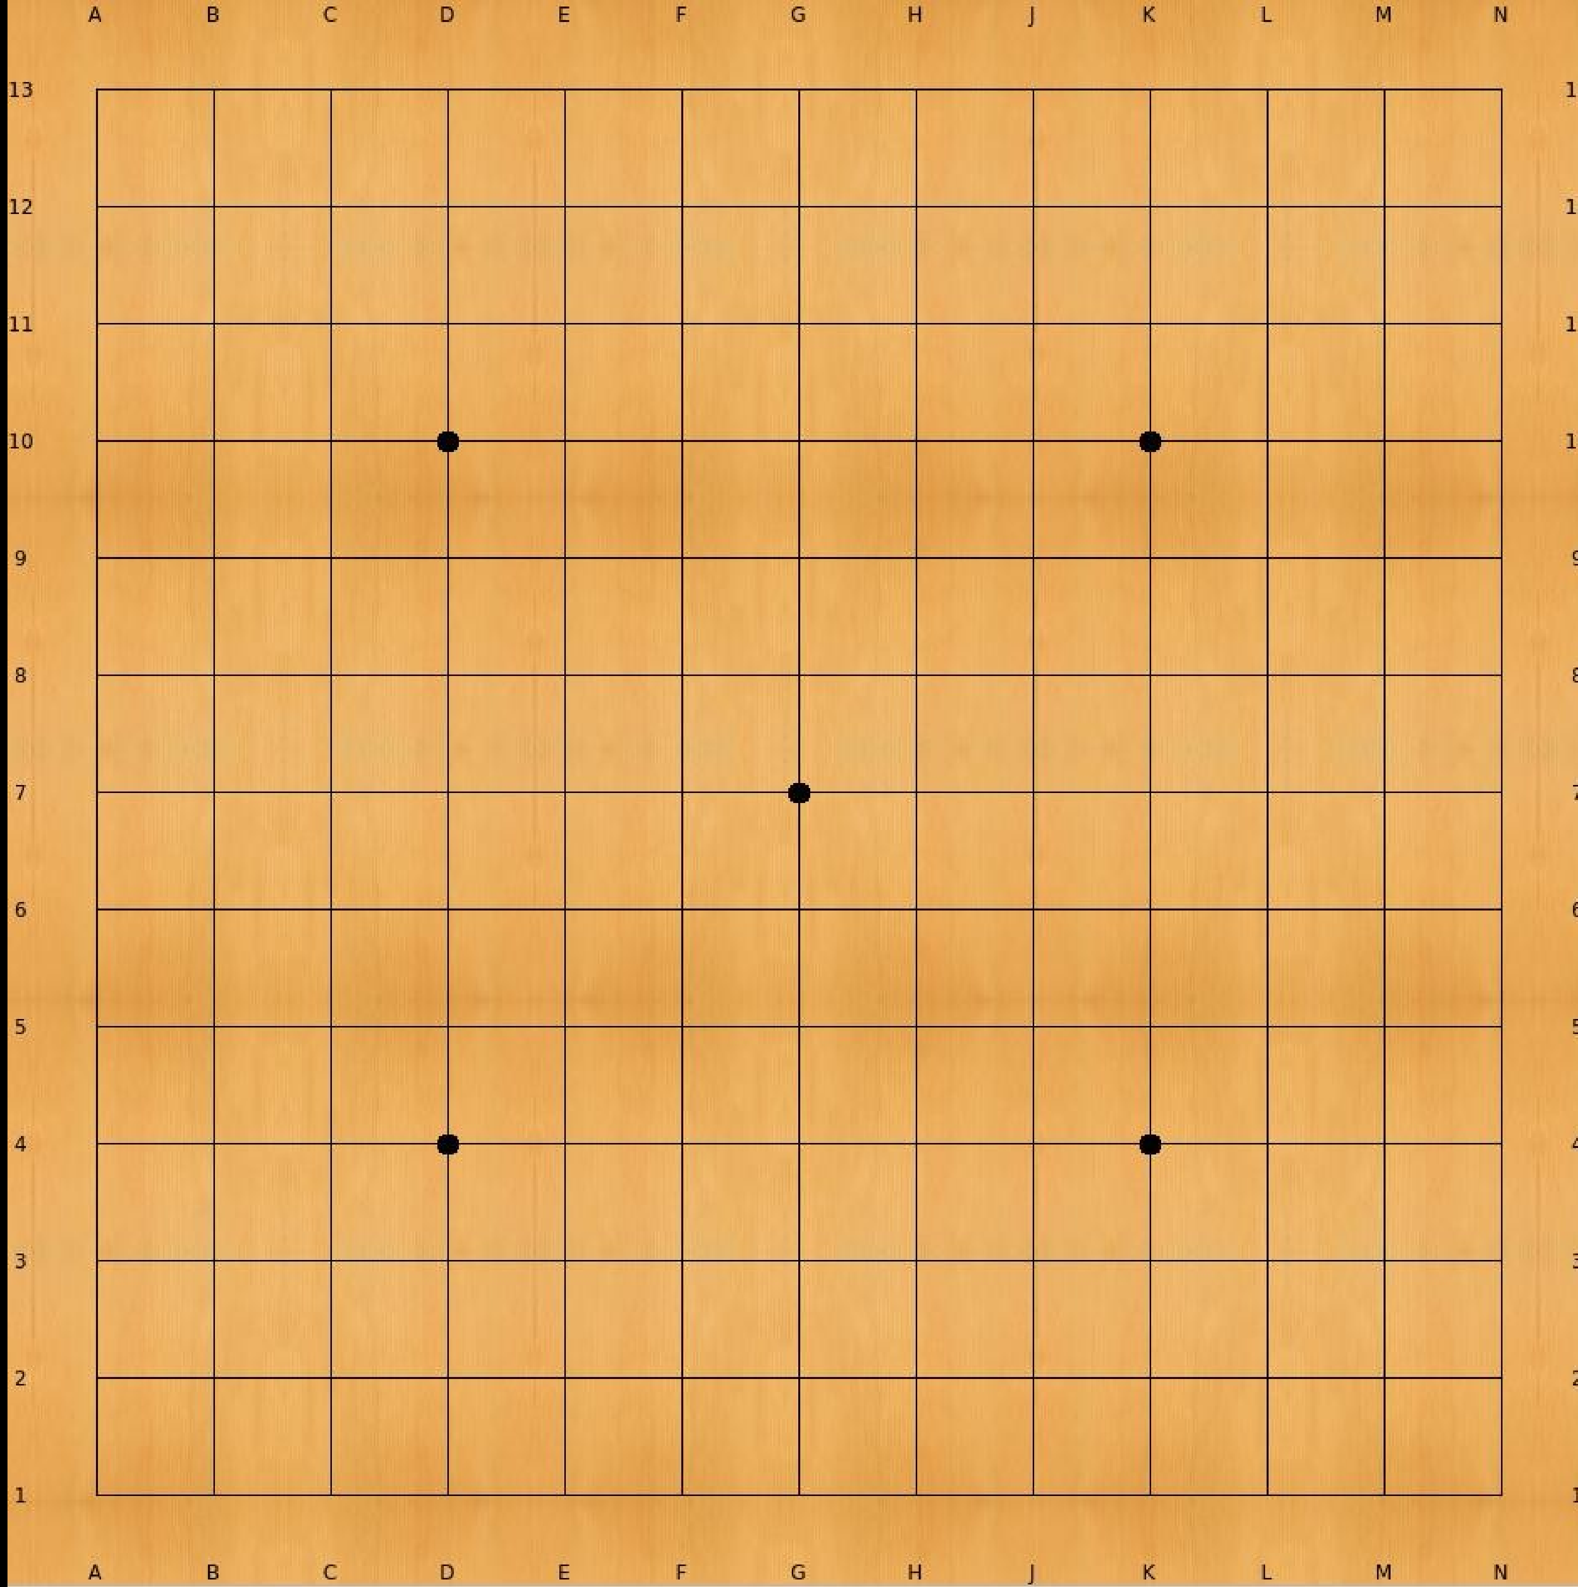
A B C D E F G H J K L M N
13 12 11 10 9 8 7 6 5 4 3 2 1

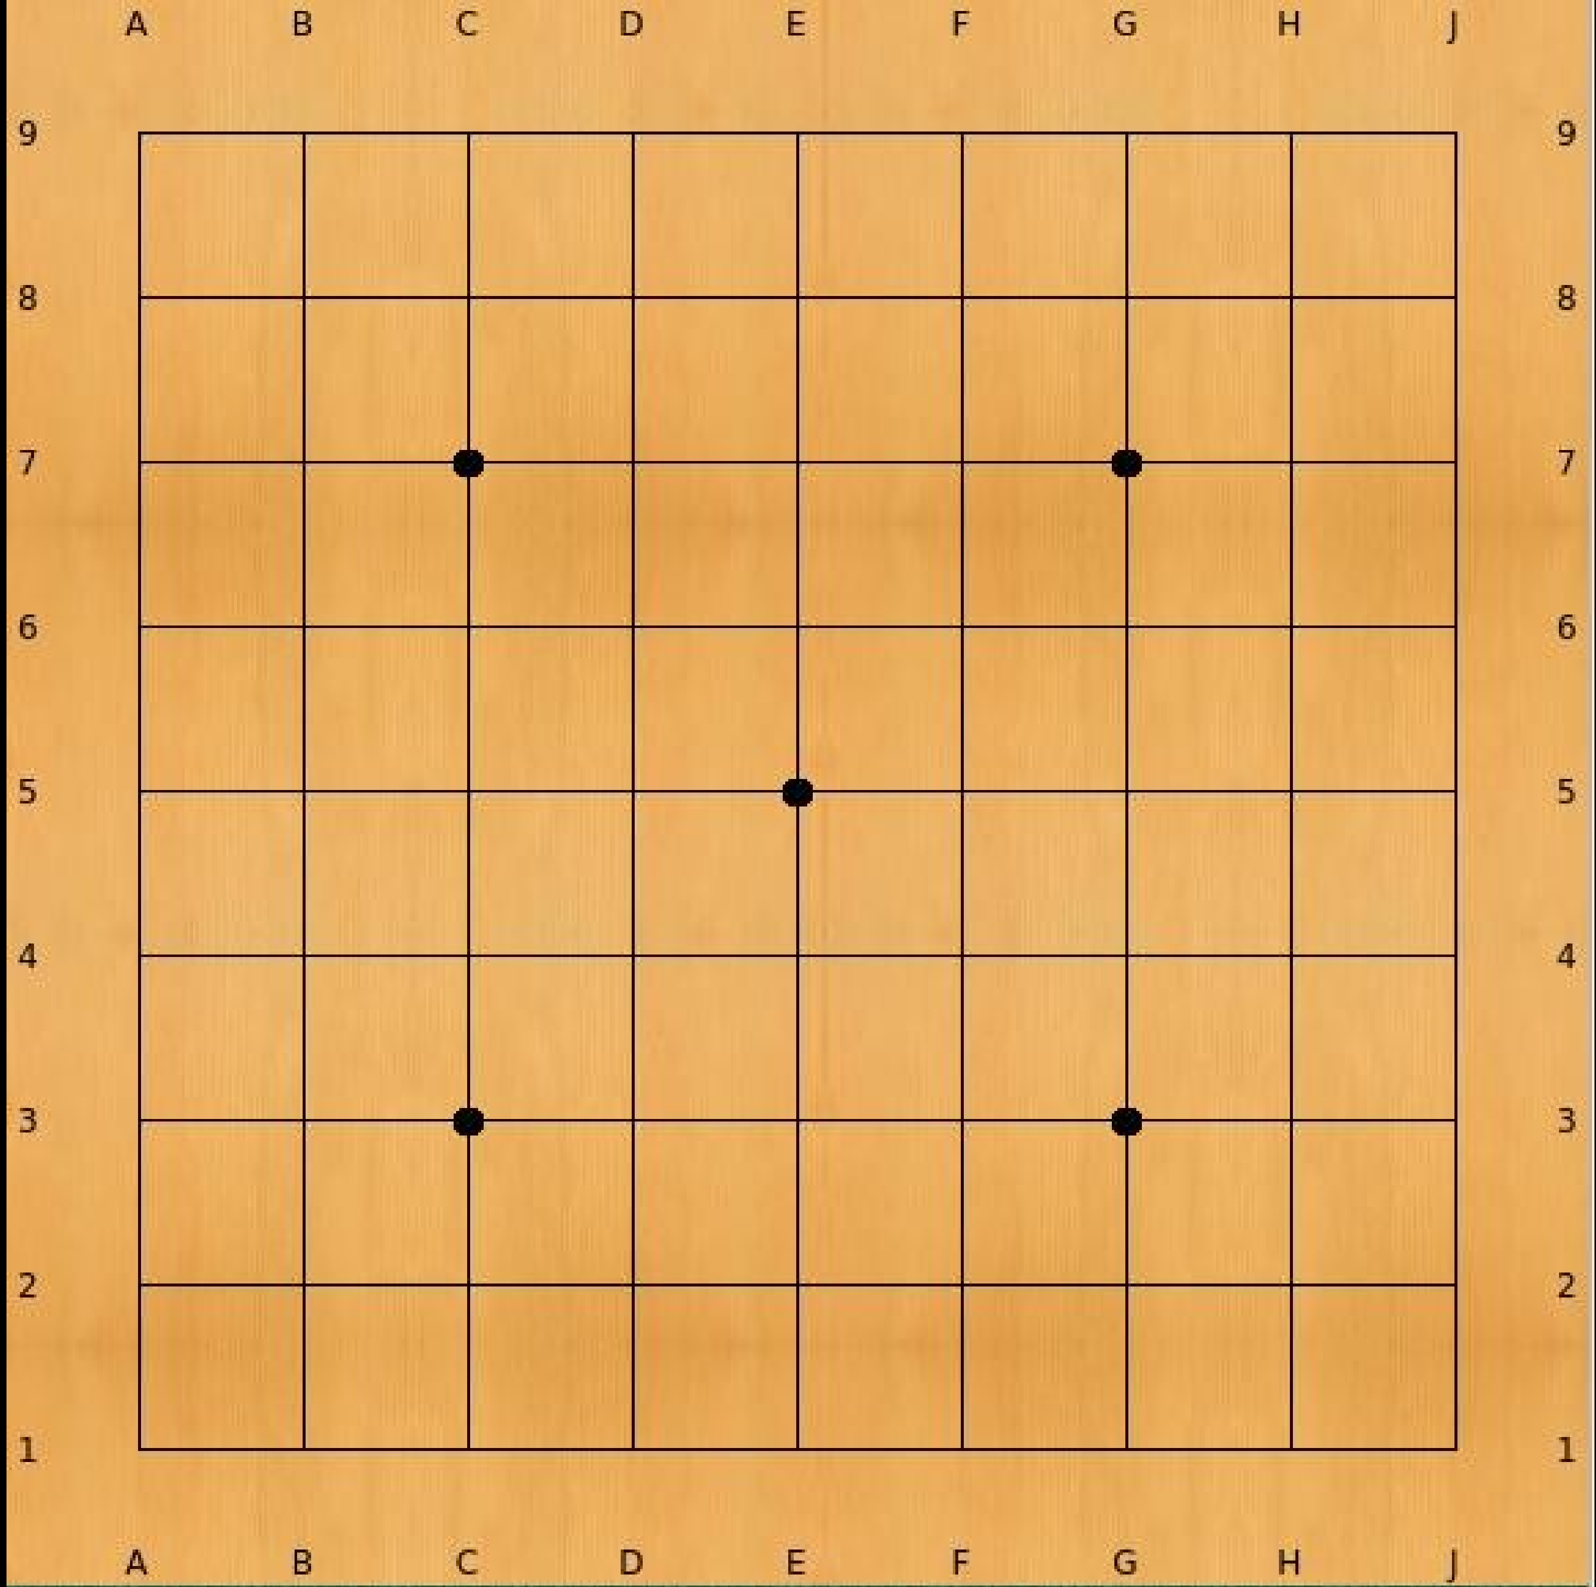
A B C D E F G H J
9 8 7 6 5 4 3 2 1

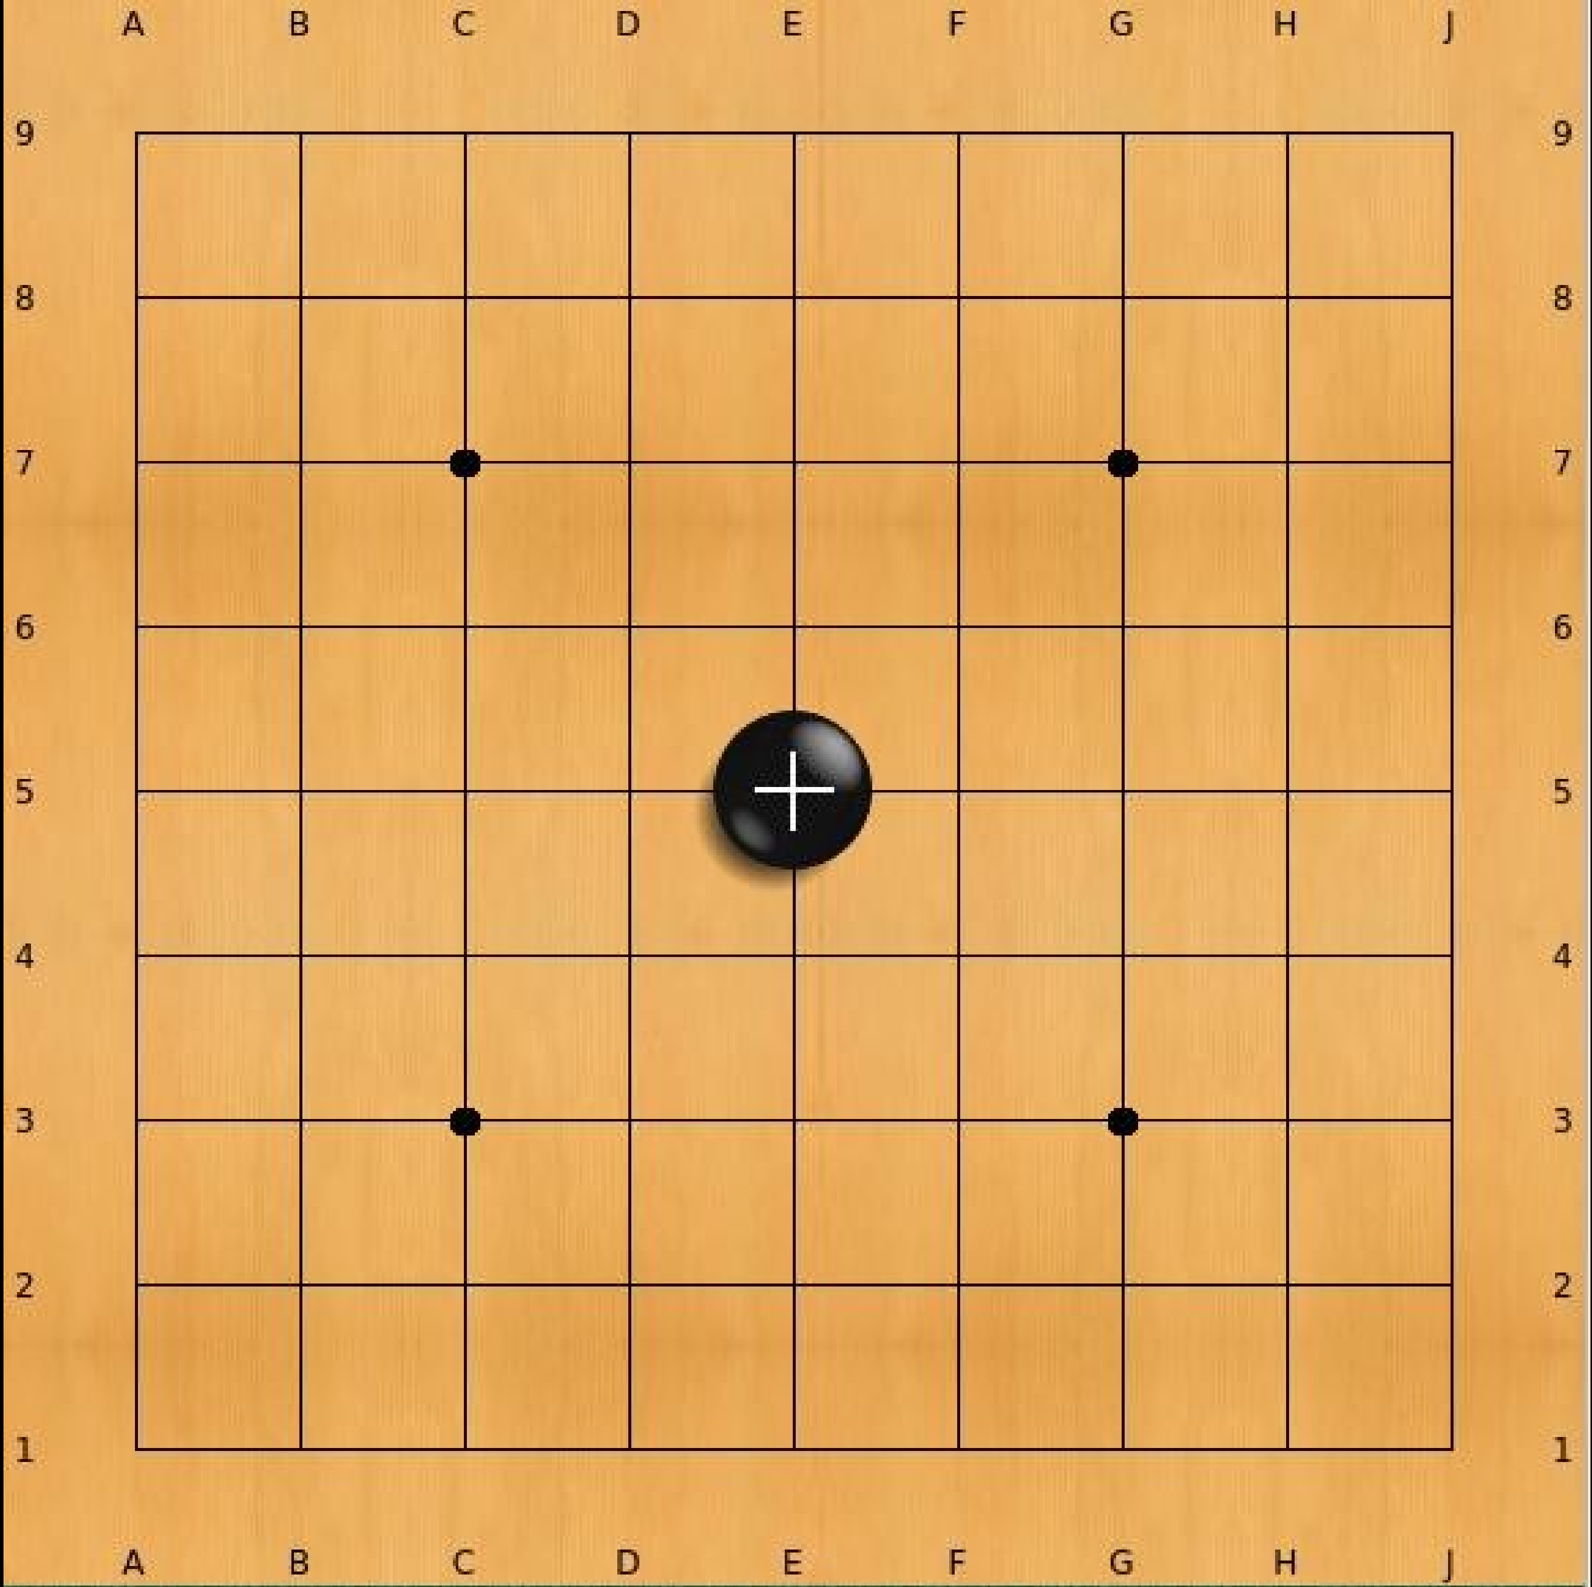
A B C D E F G H J
9 8 7 6 5 4 3 2 1

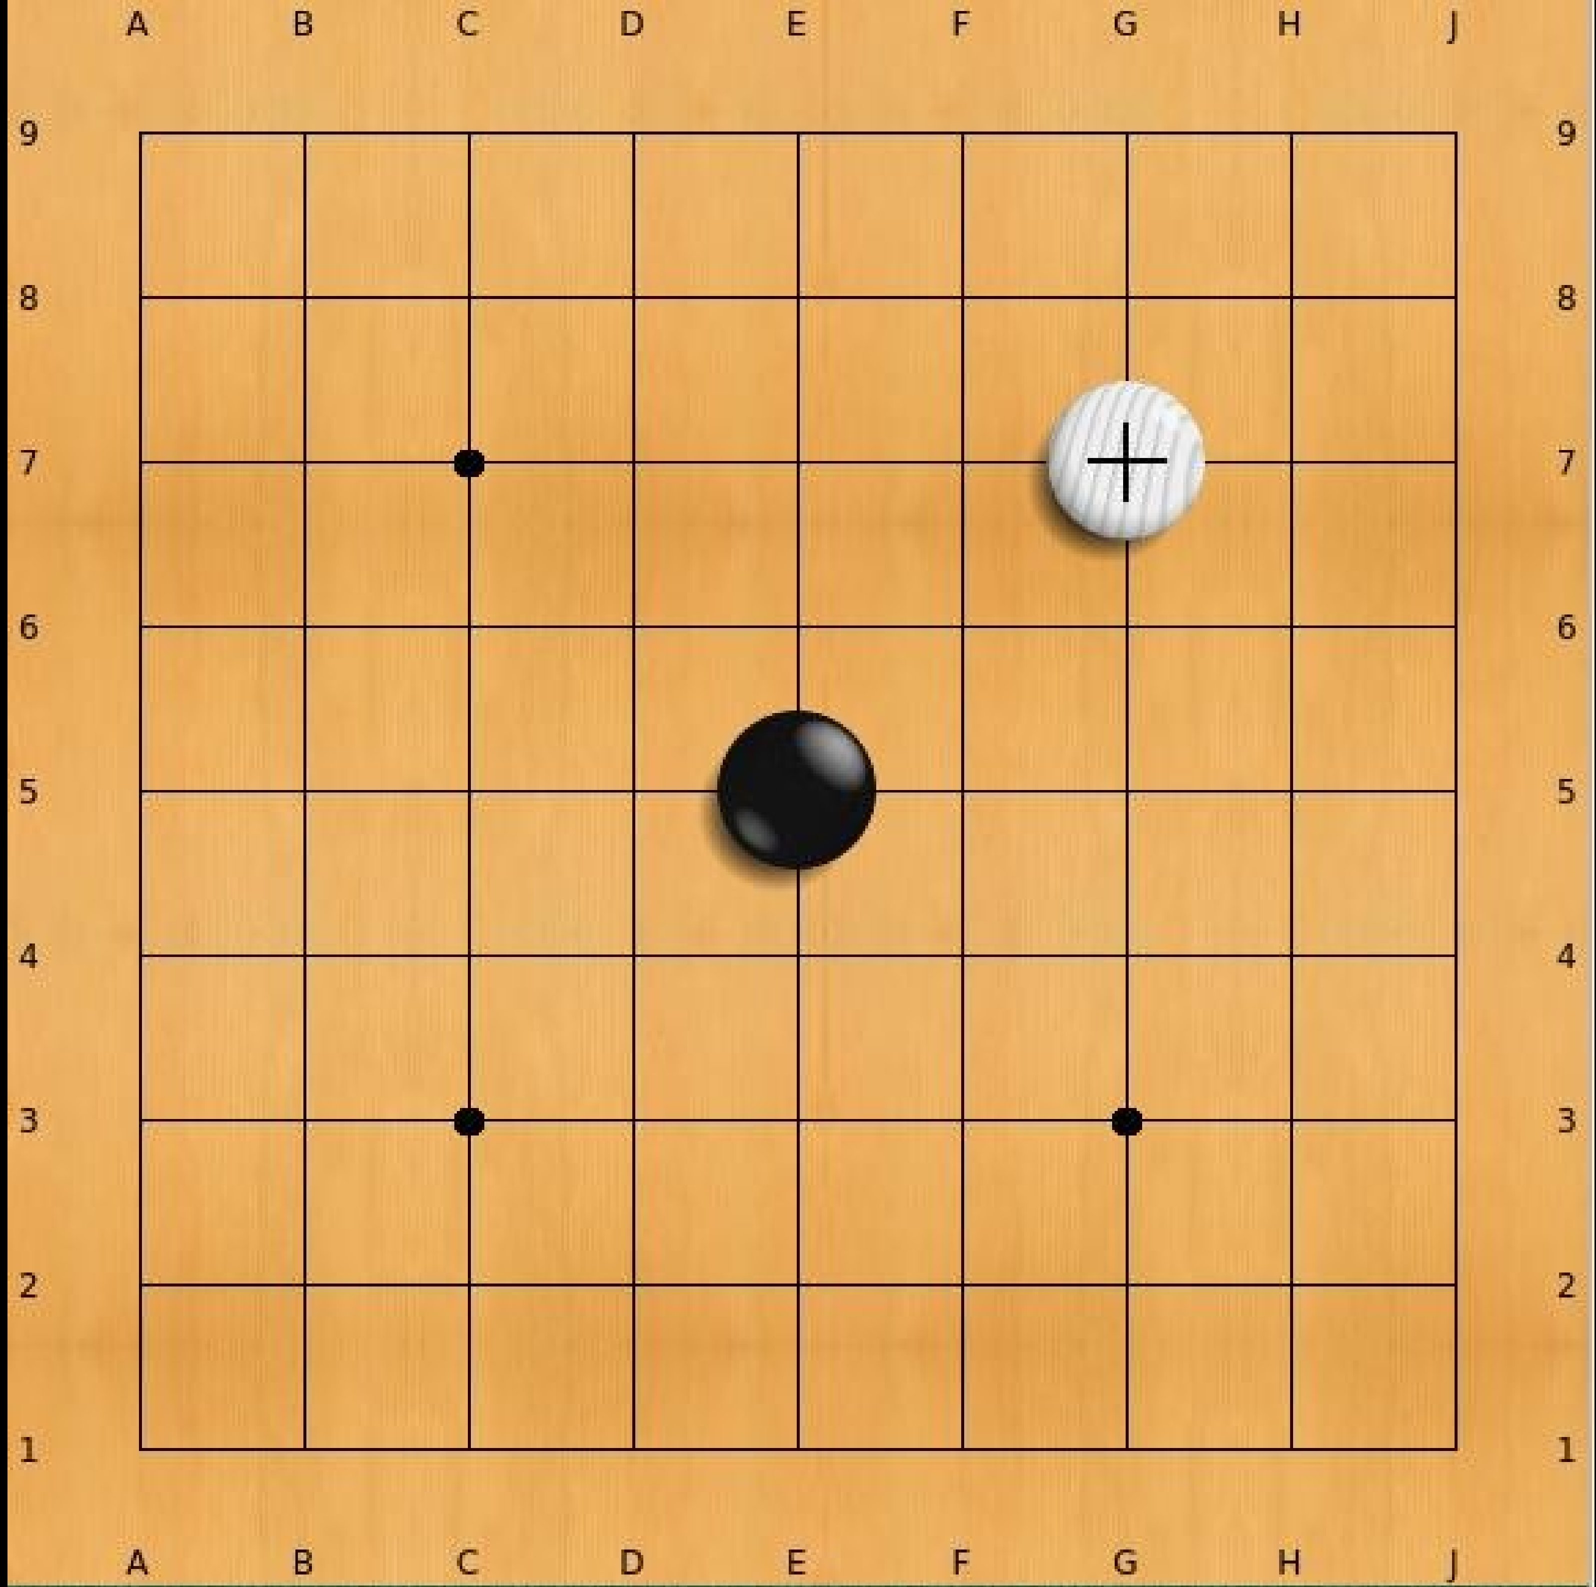
A B C D E F G H J
9 8 7 6 5 4 3 2 1

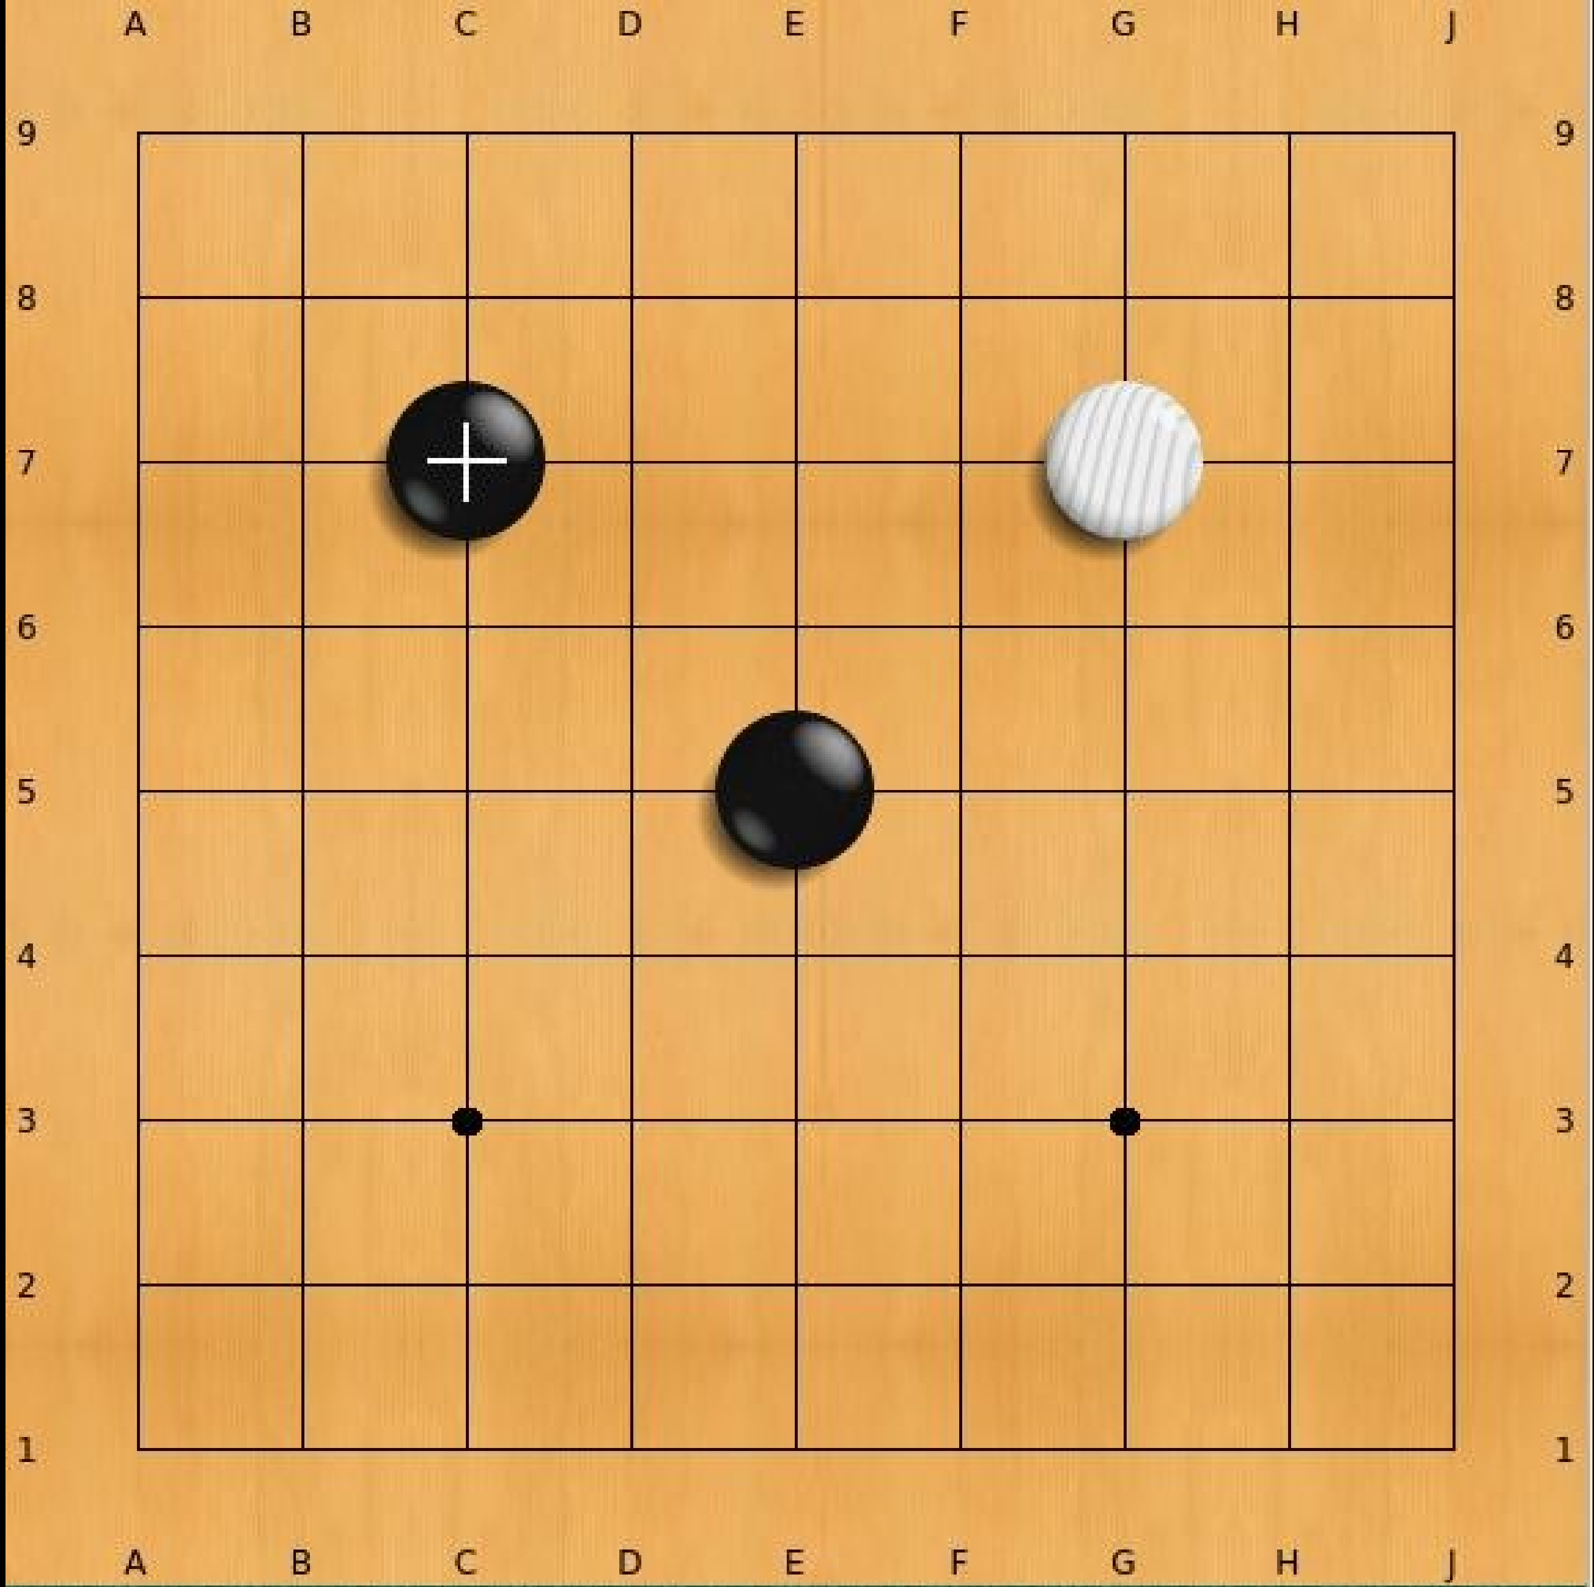
A B C D E F G H J
9
8
7
6
5
4
3
2
1

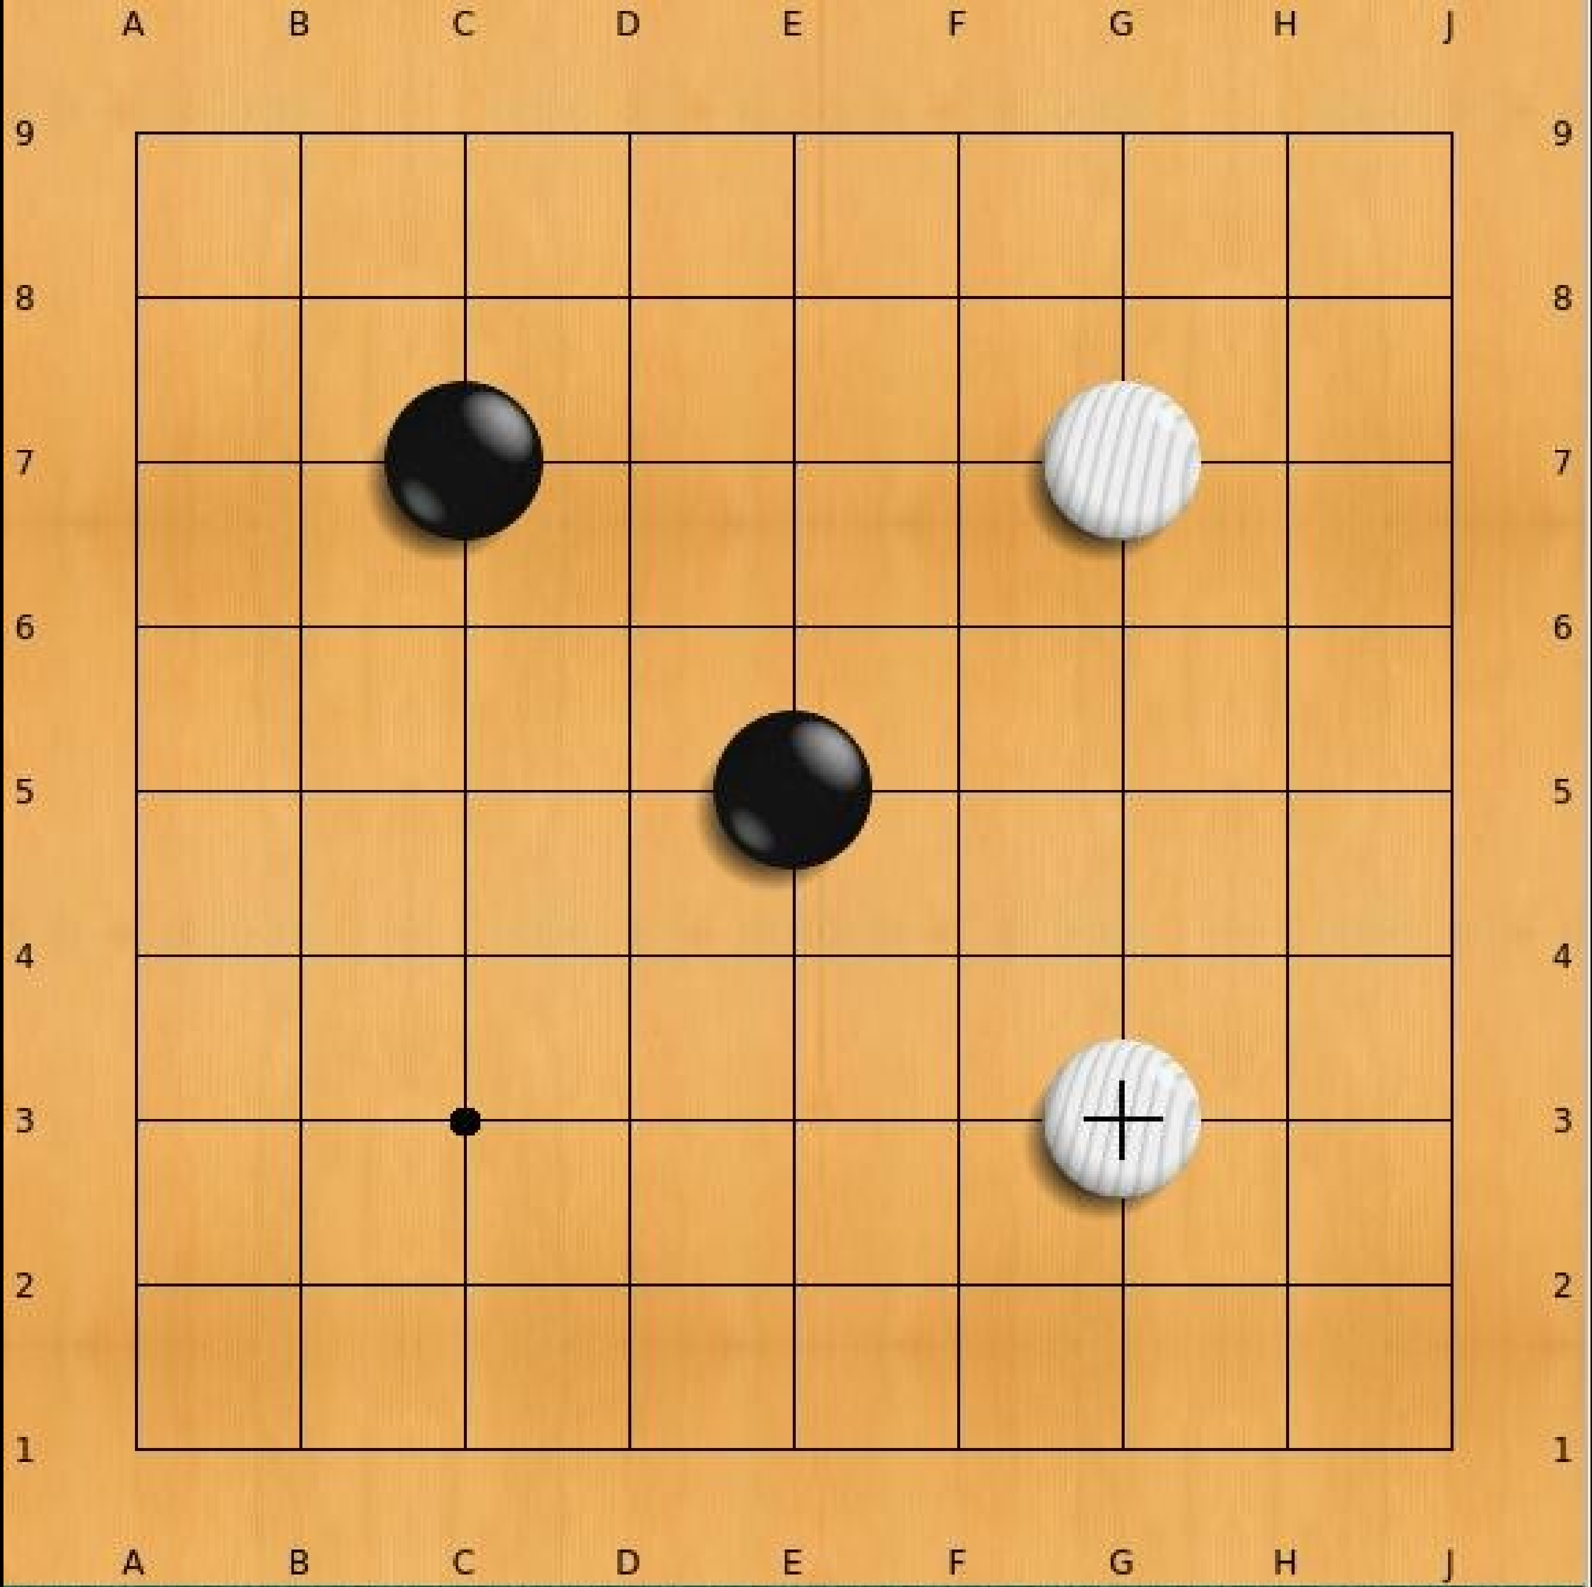
A B C D E F G H J
9 8 7 6 5 4 3 2 1

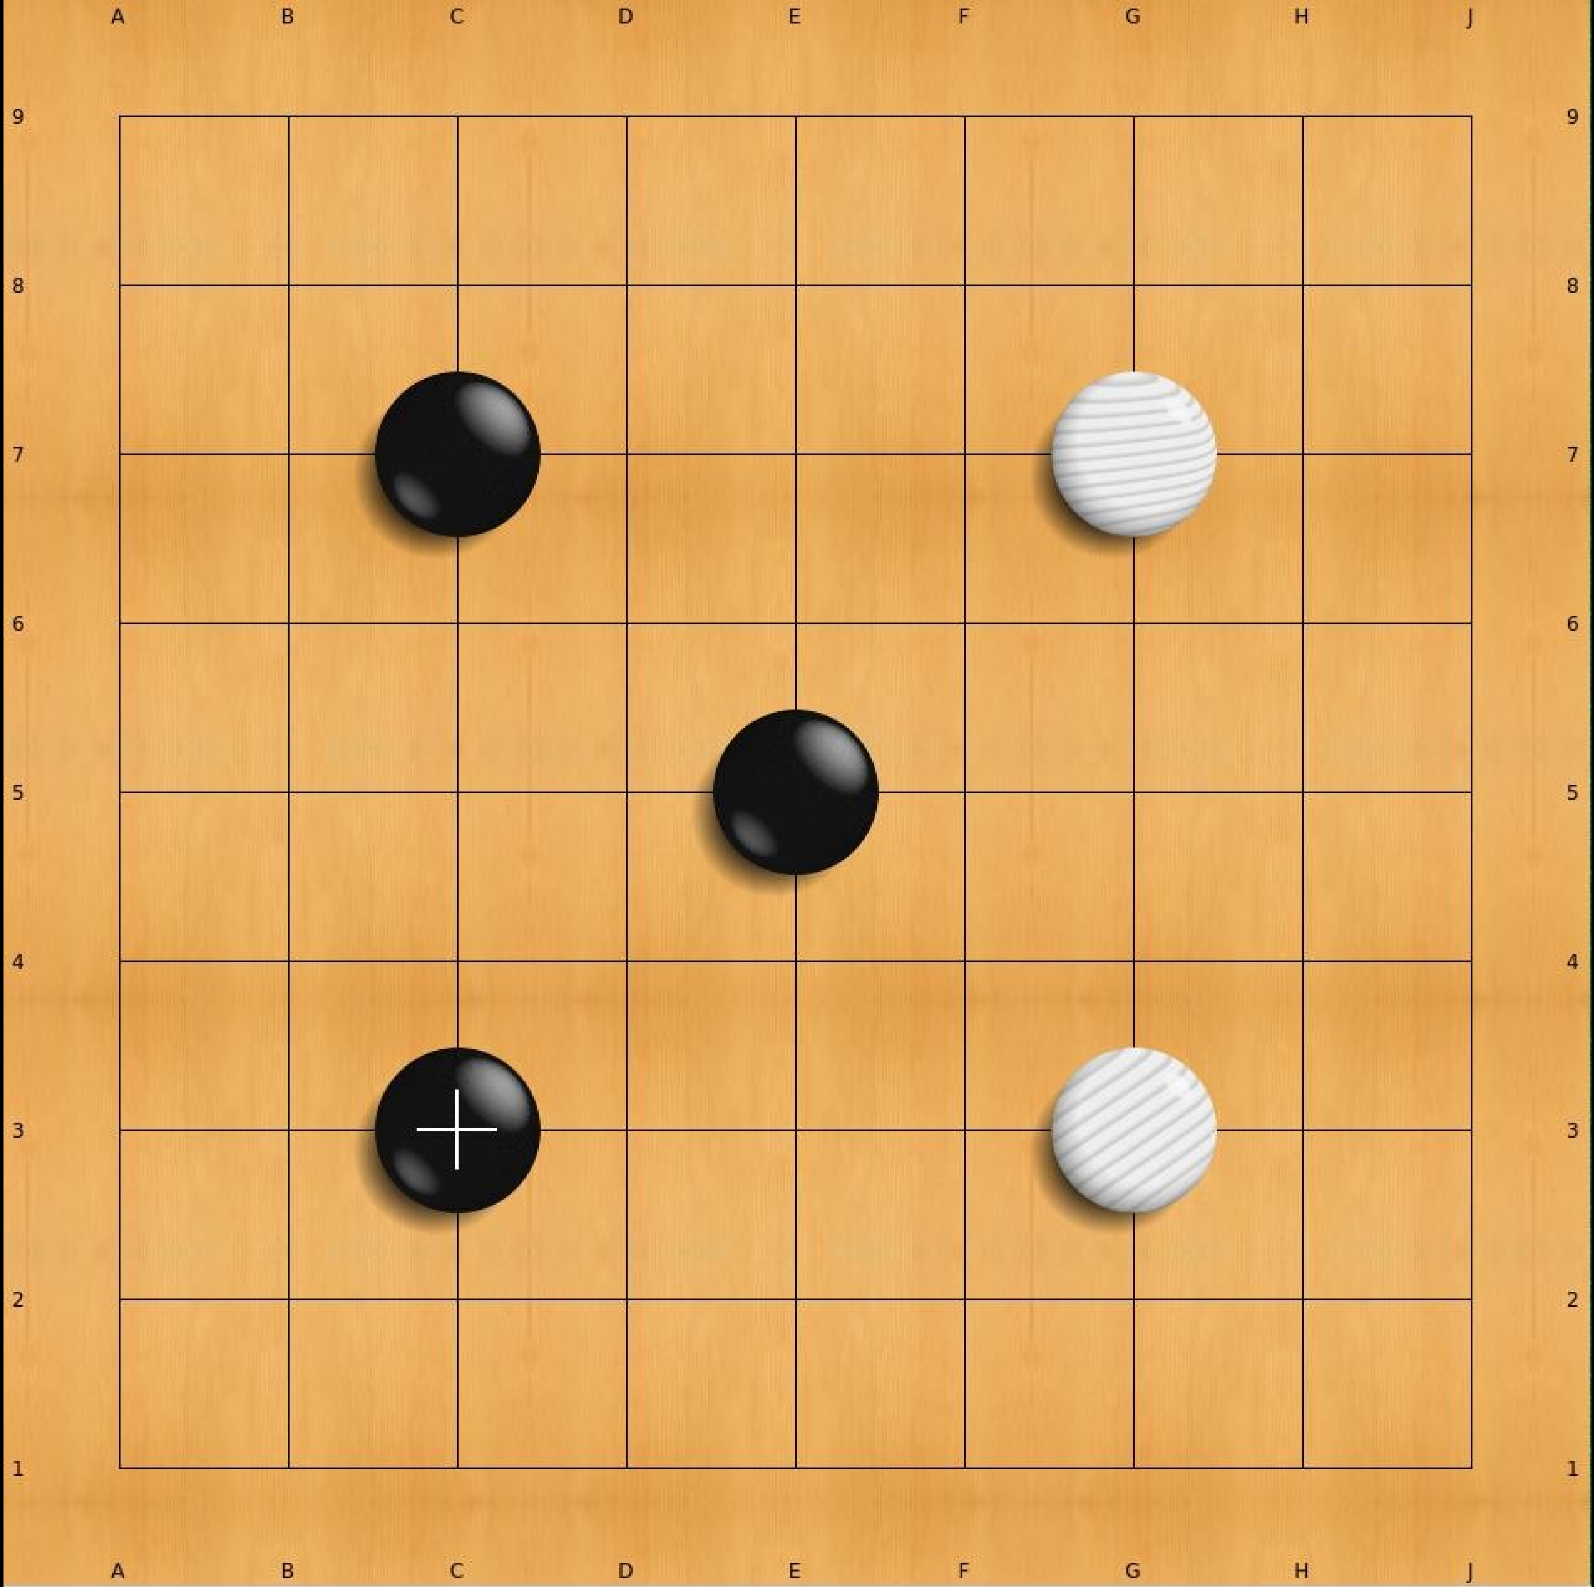
A B C D E F G H J
9 8 7 6 5 4 3 2 1

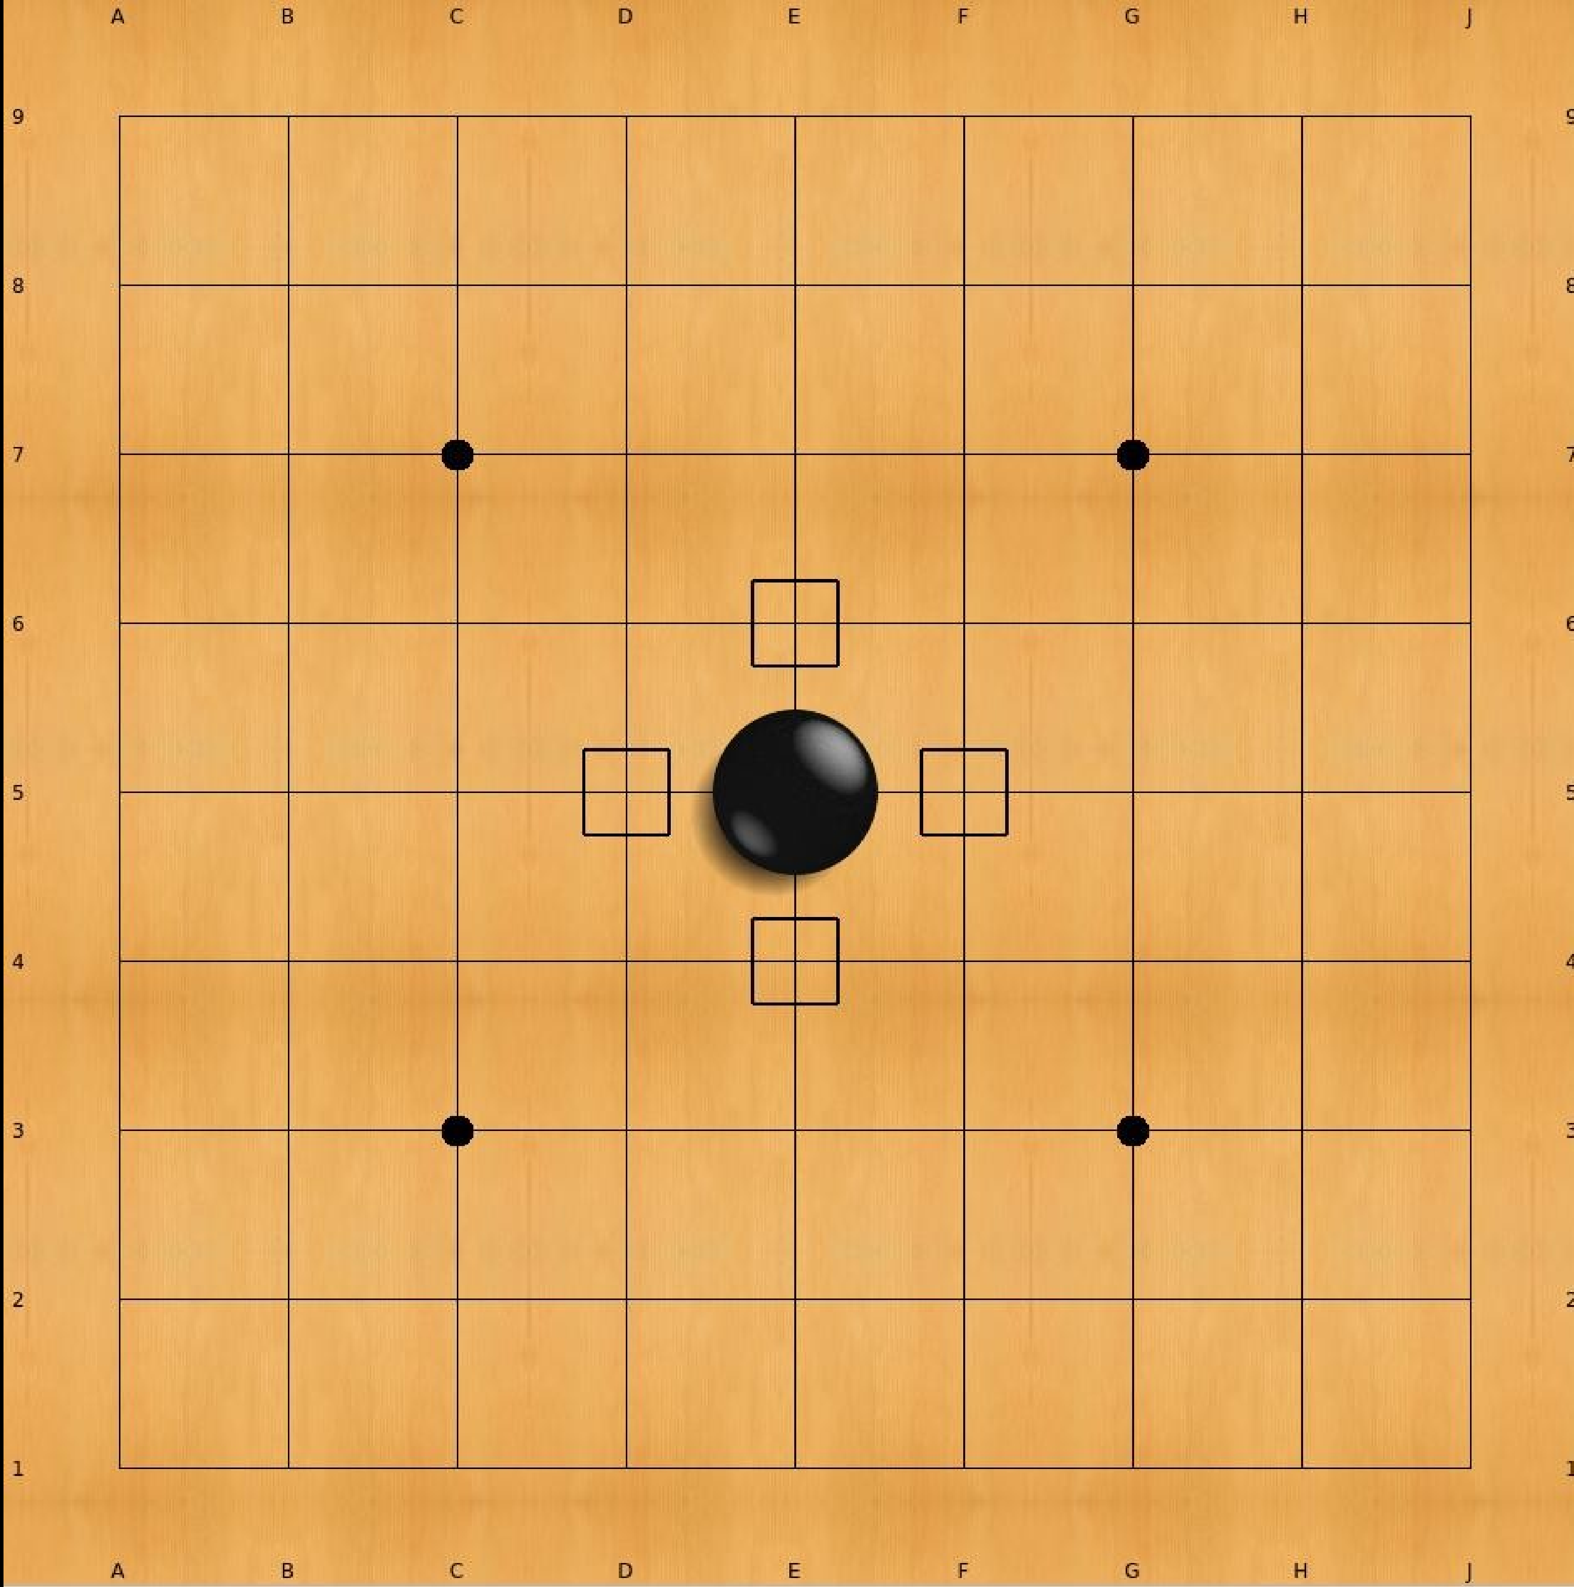
A B C D E F G H J
9 8 7 6 5 4 3 2 1

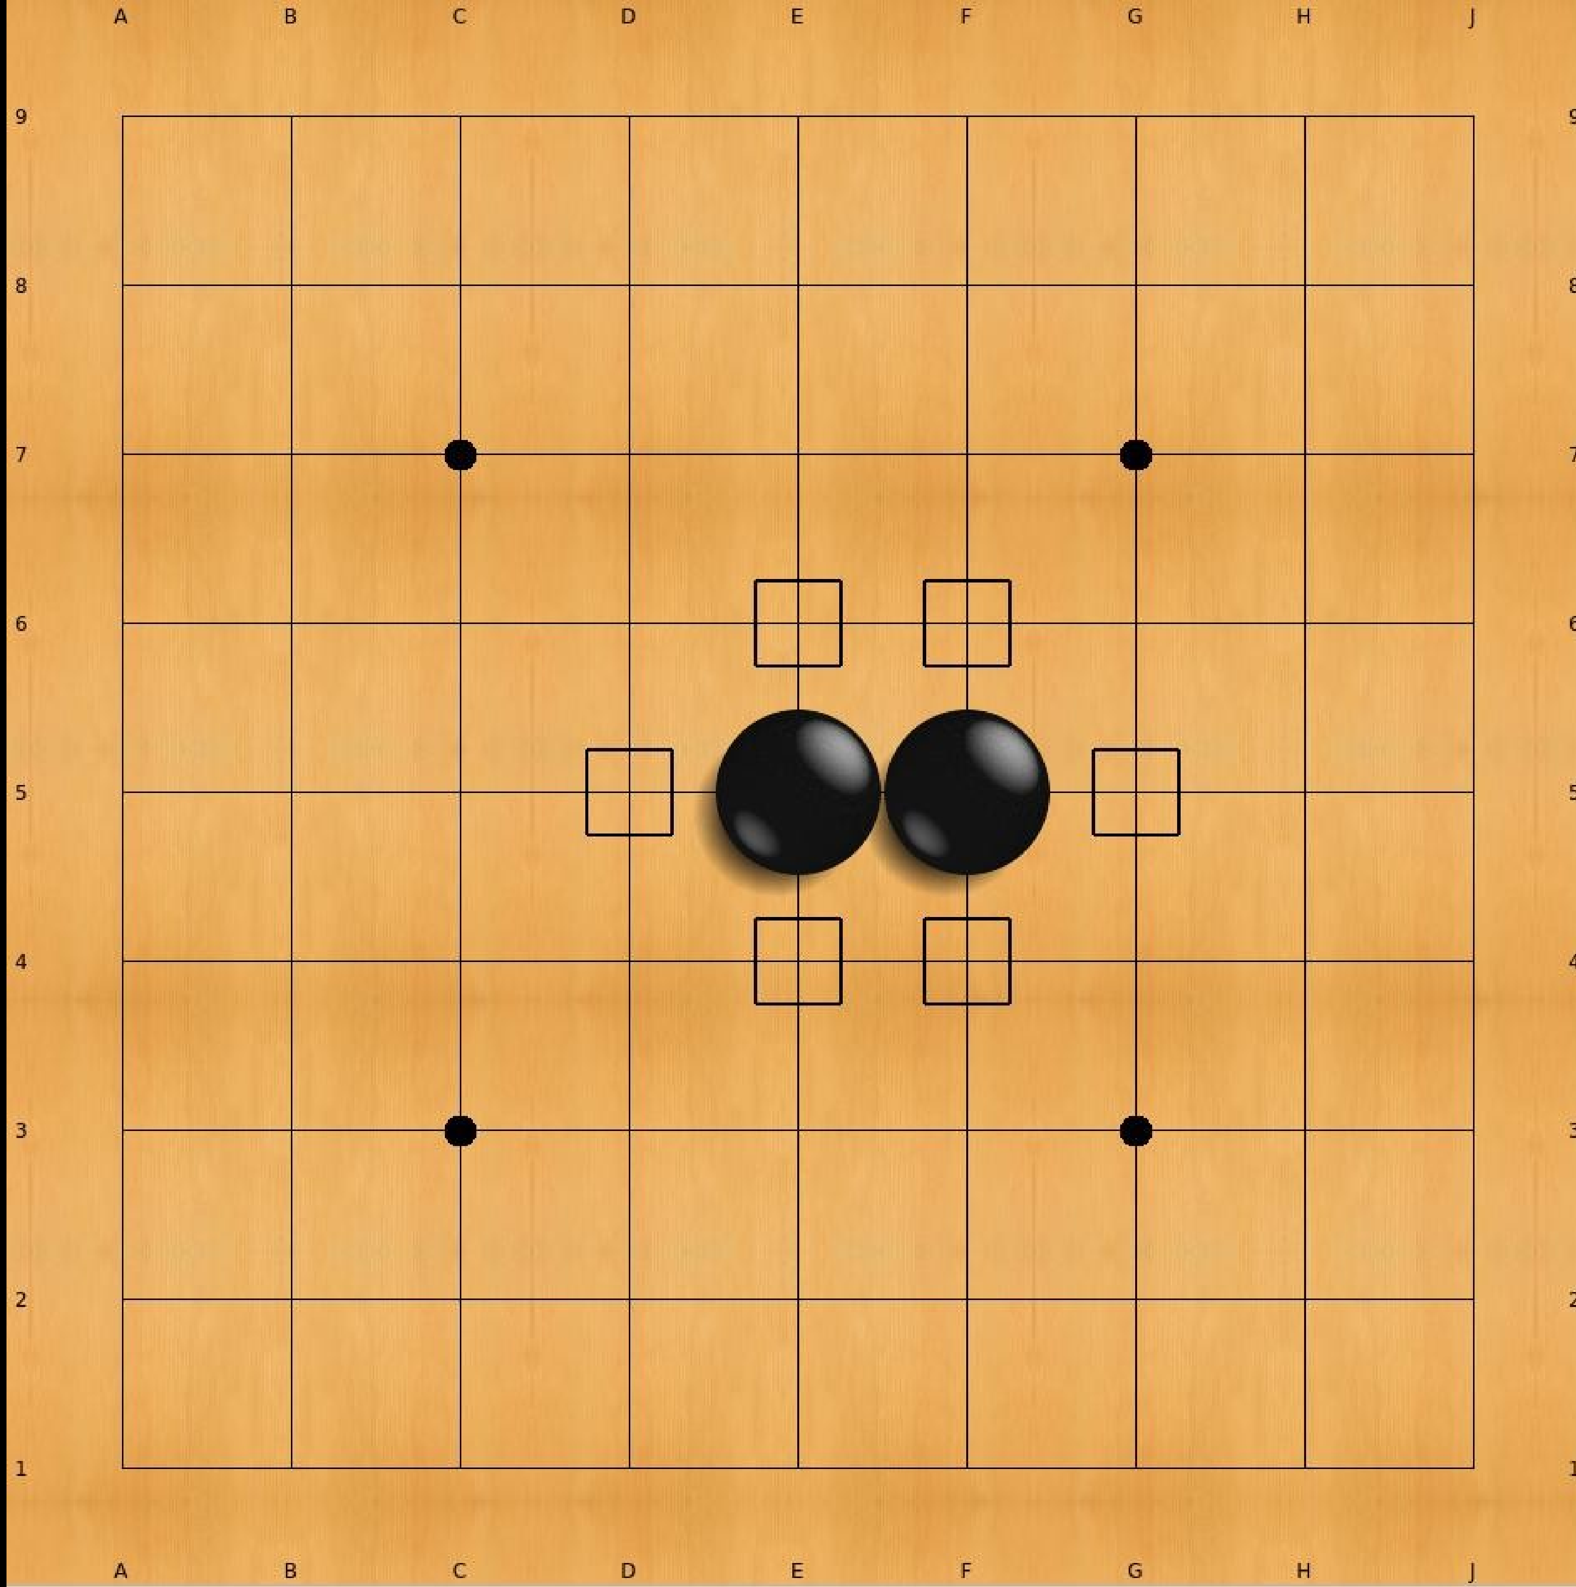
A B C D E F G H J
9 8 7 6 5 4 3 2 1

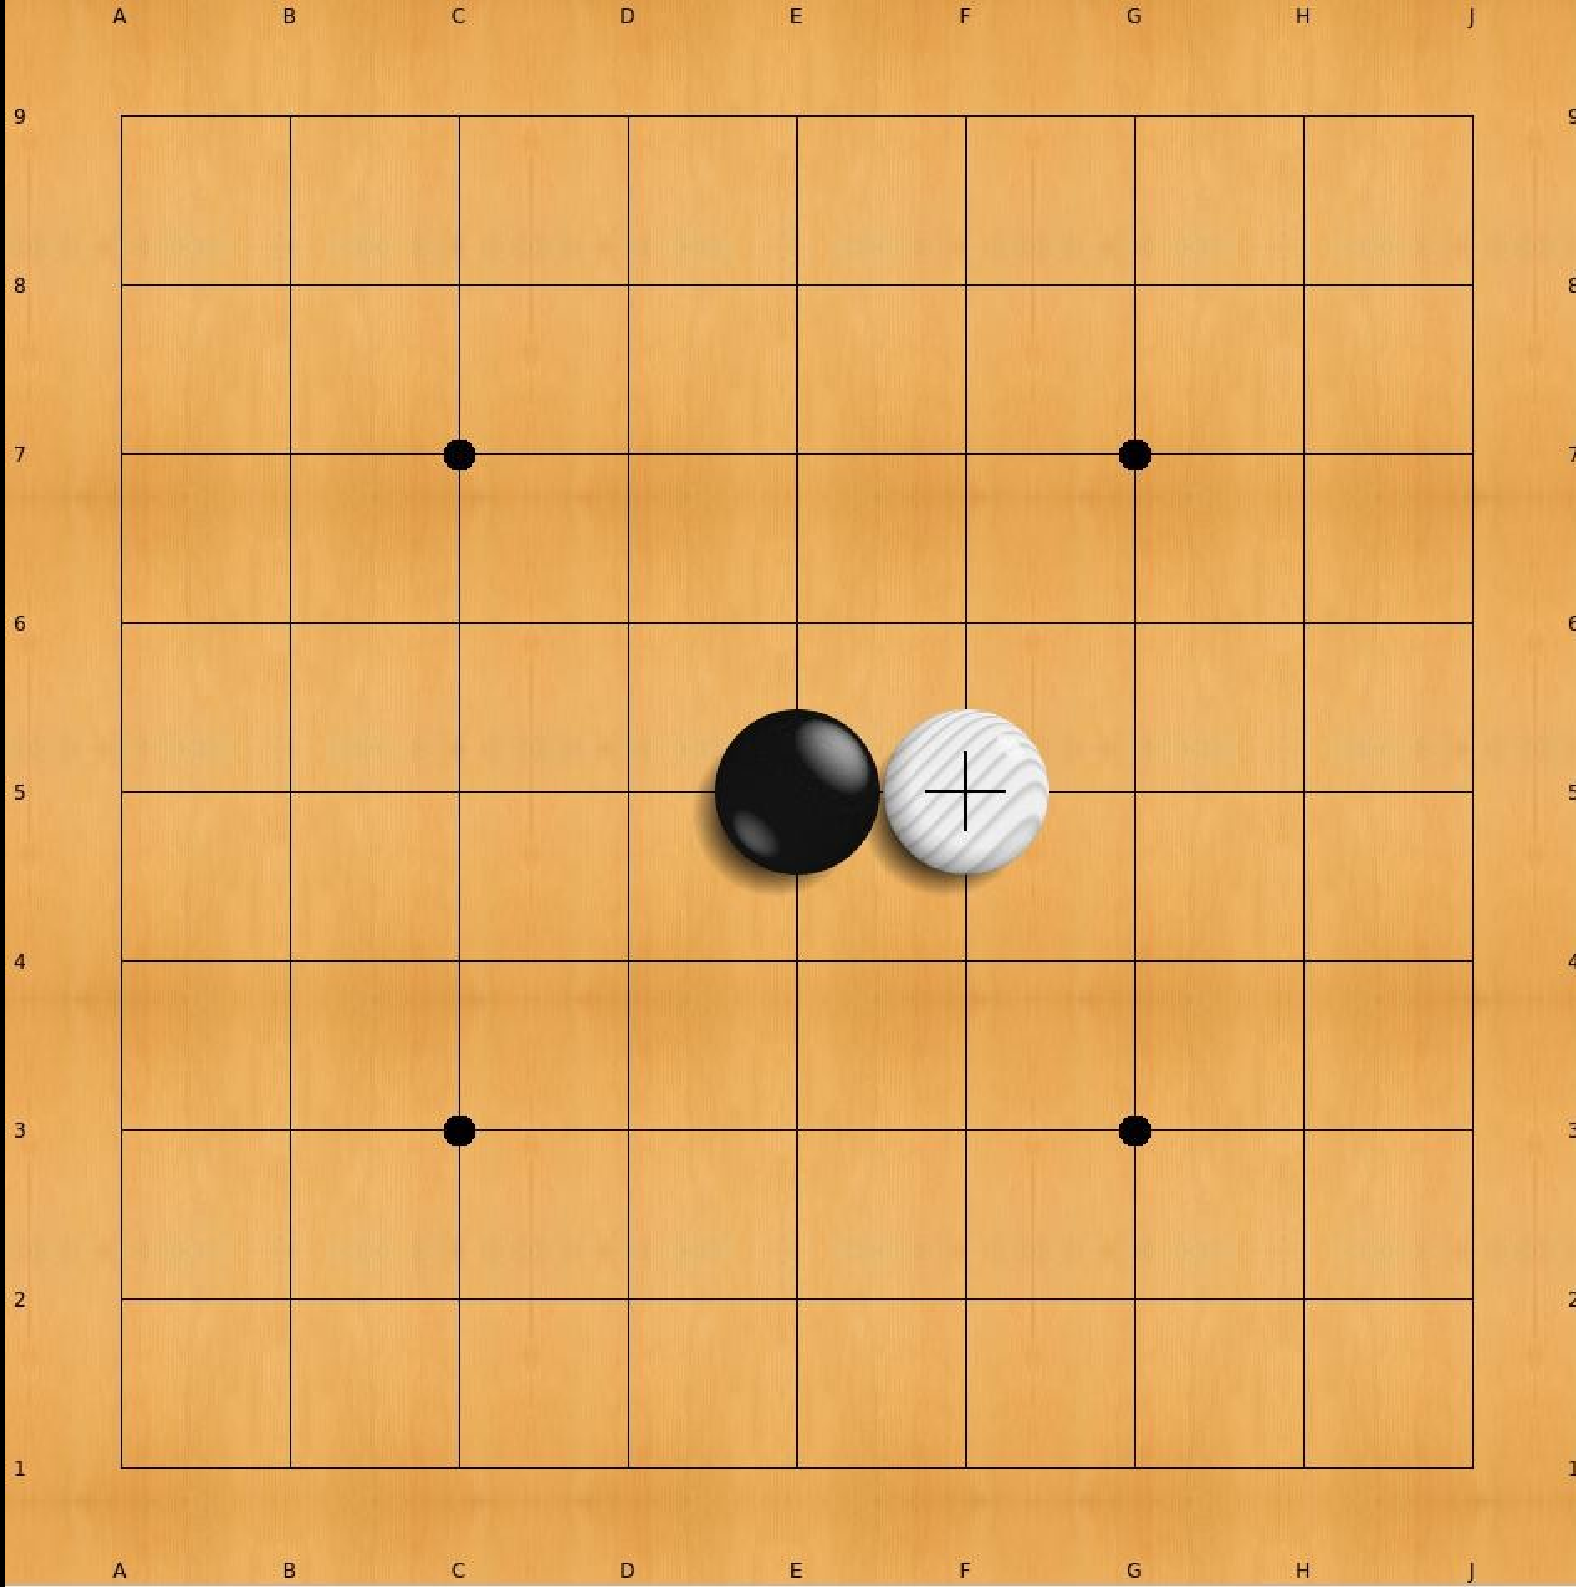
A B C D E F G H J
9 8 7 6 5 4 3 2 1

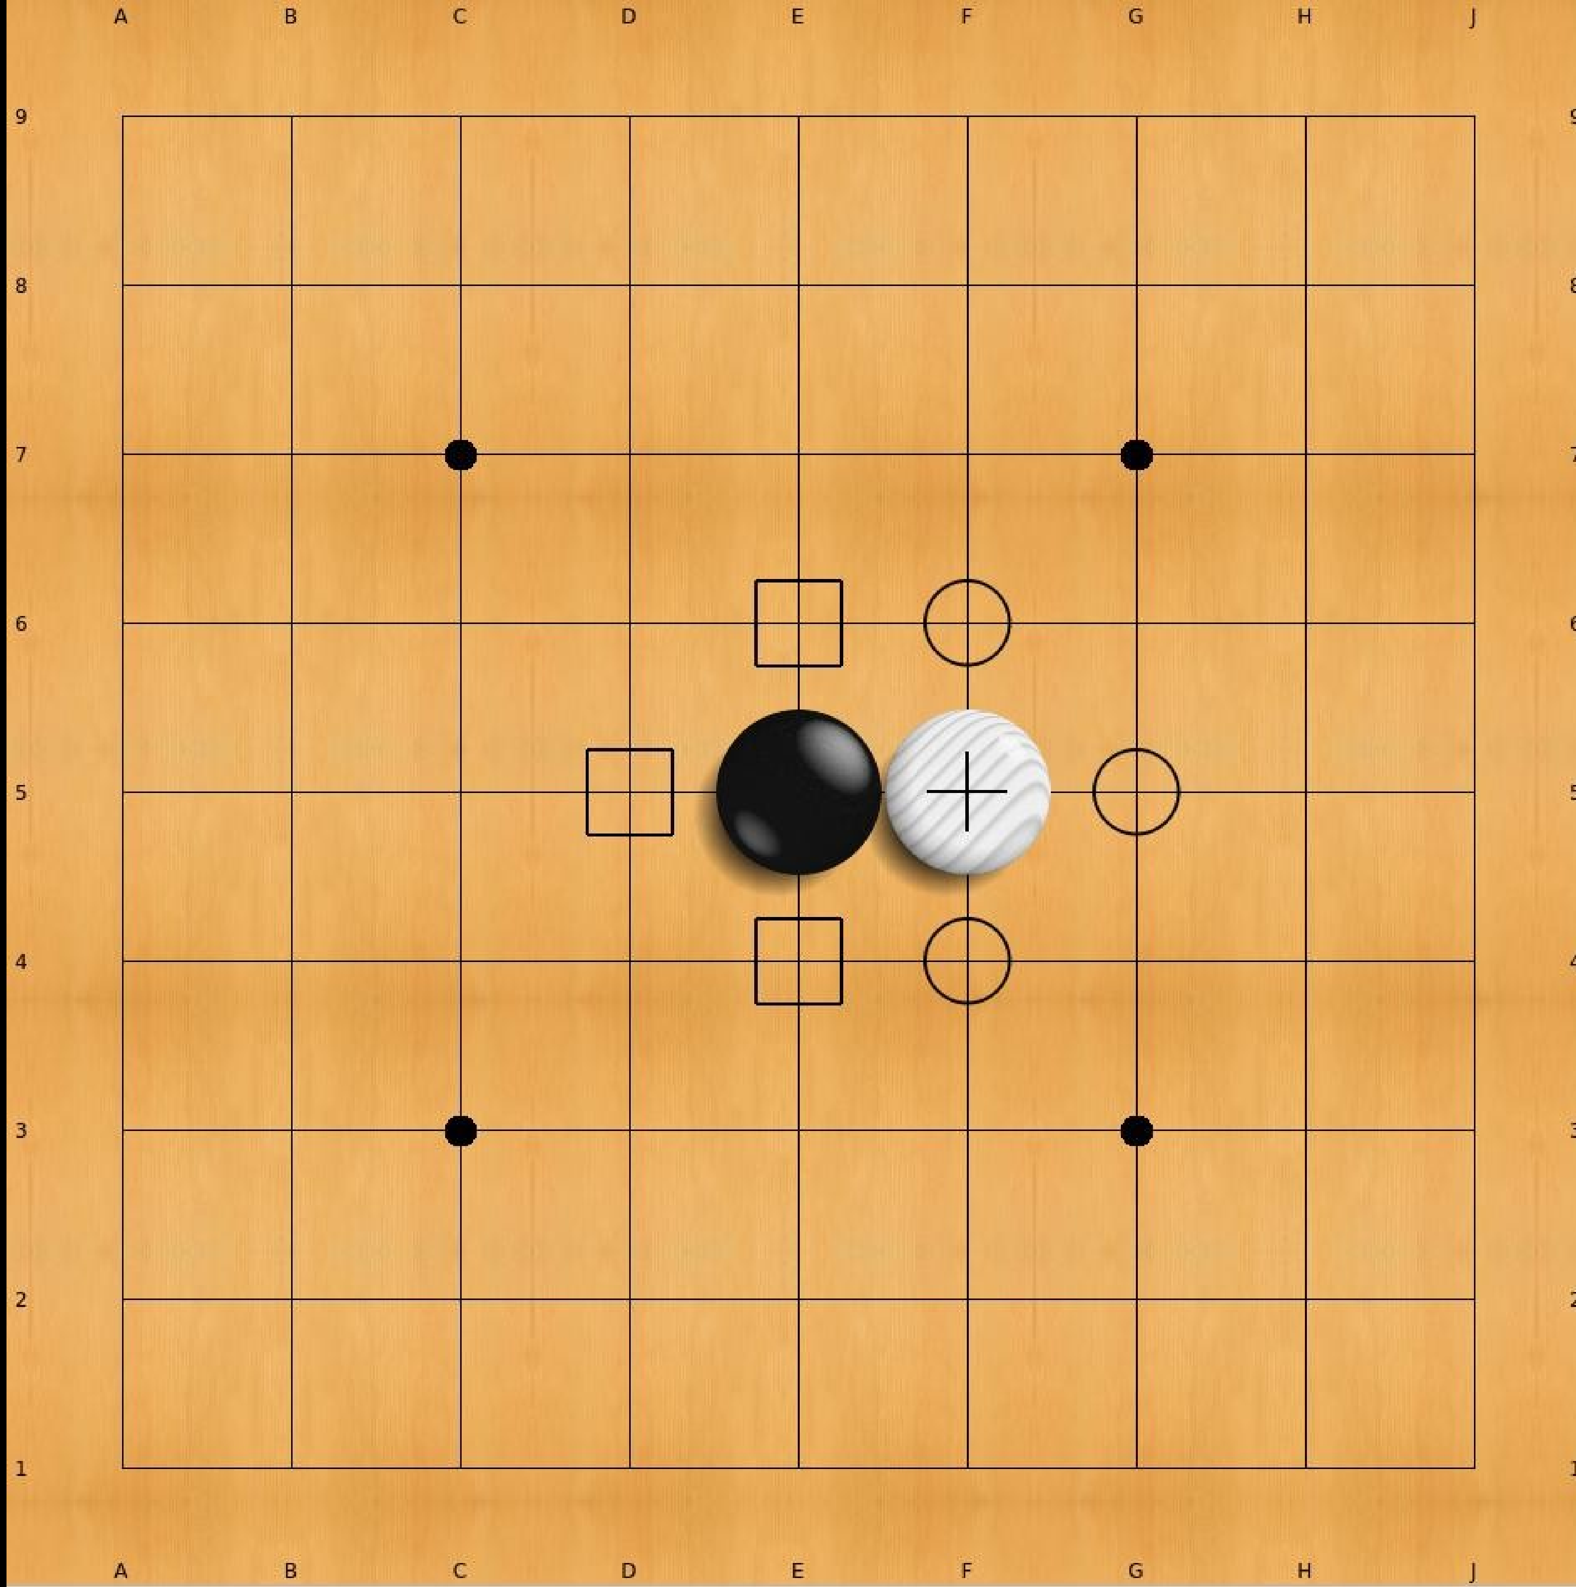
A B C D E F G H J
9 8 7 6 5 4 3 2 1

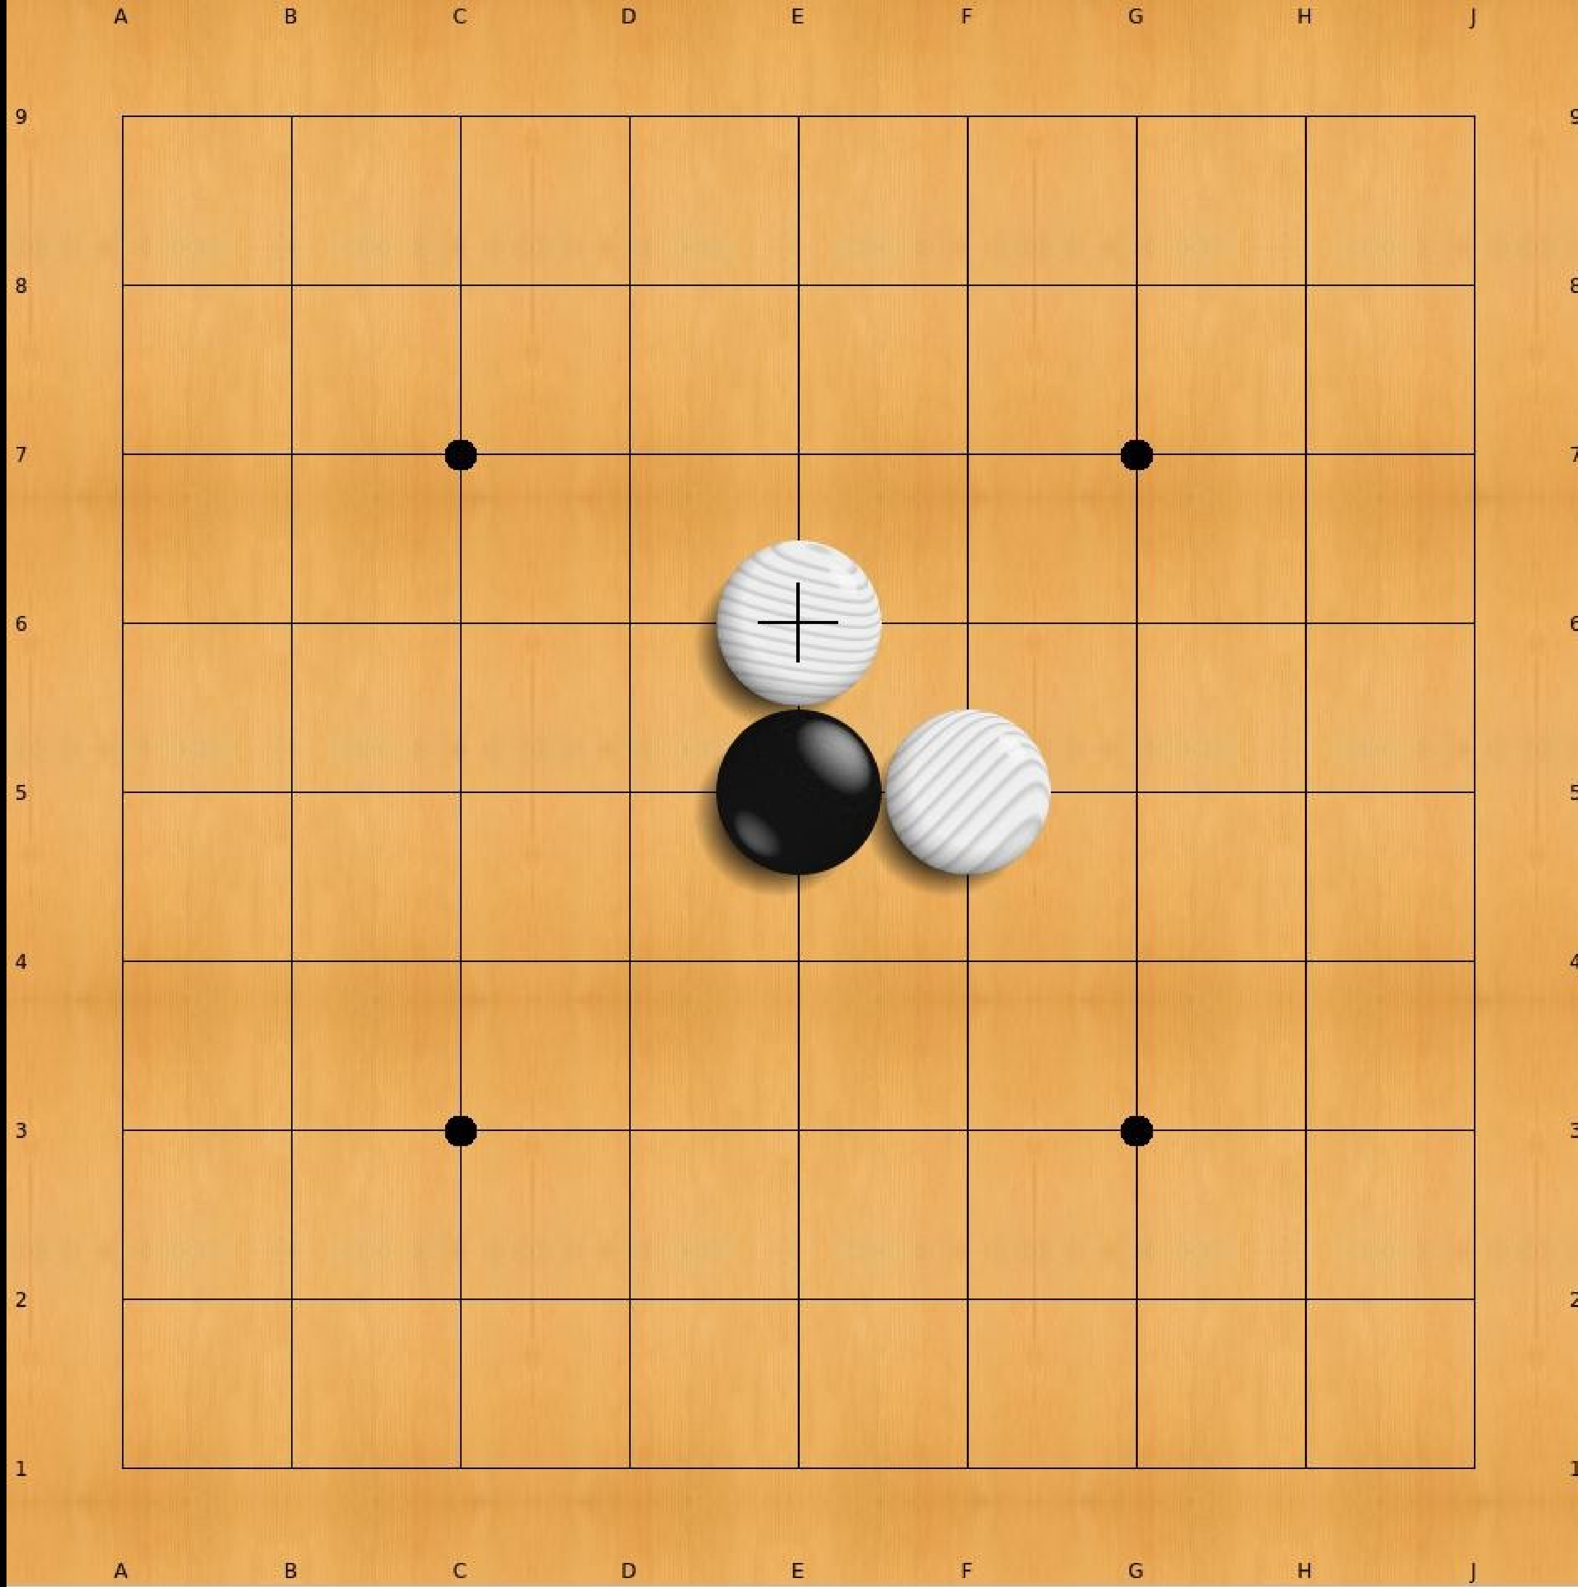
A B C D E F G H J
9
8
7
6
5
4
3
2
1
A B C D E F G H J

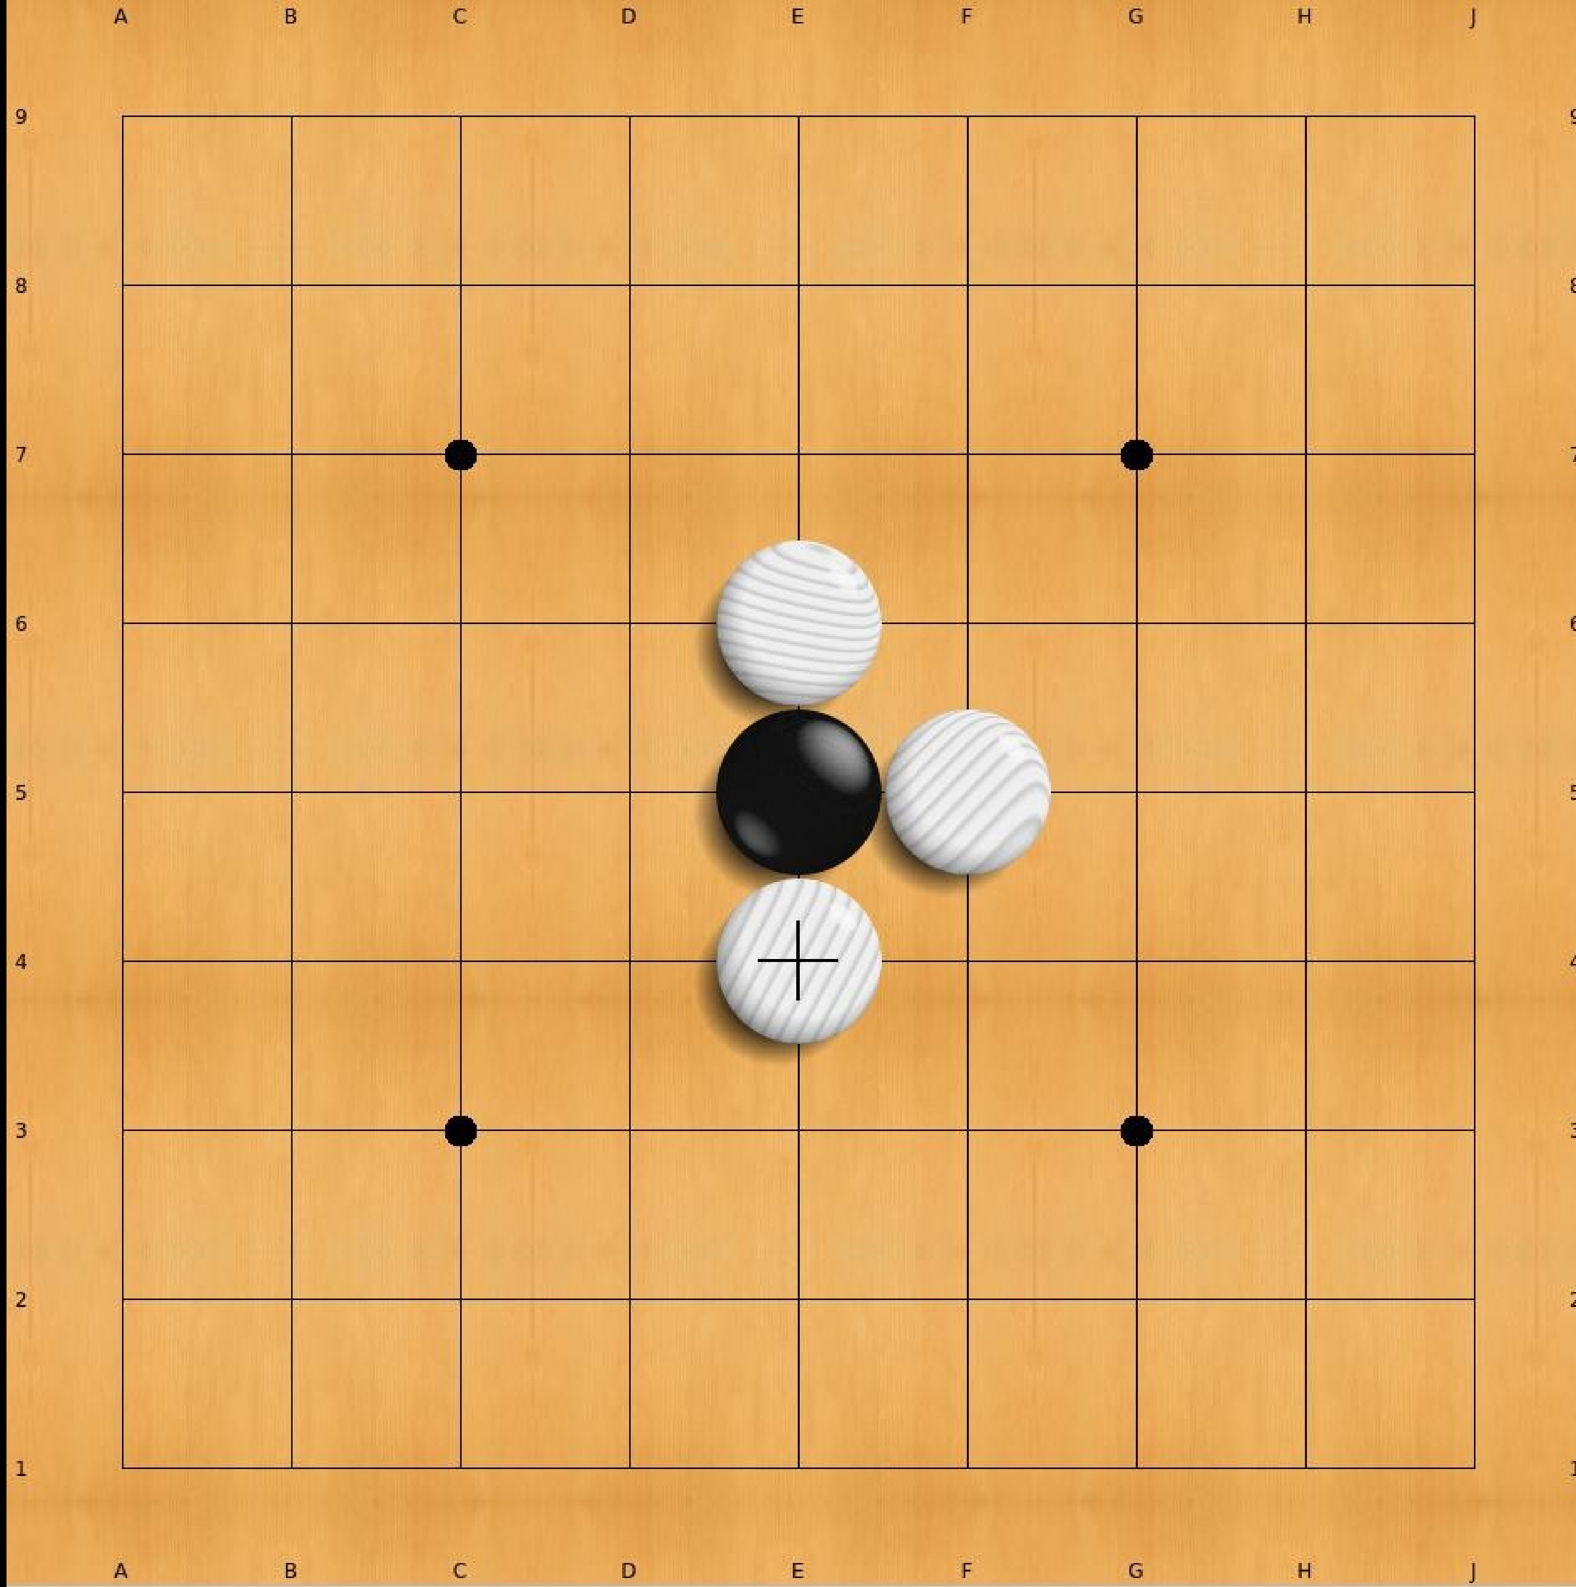
A B C D E F G H J
9 8 7 6 5 4 3 2 1

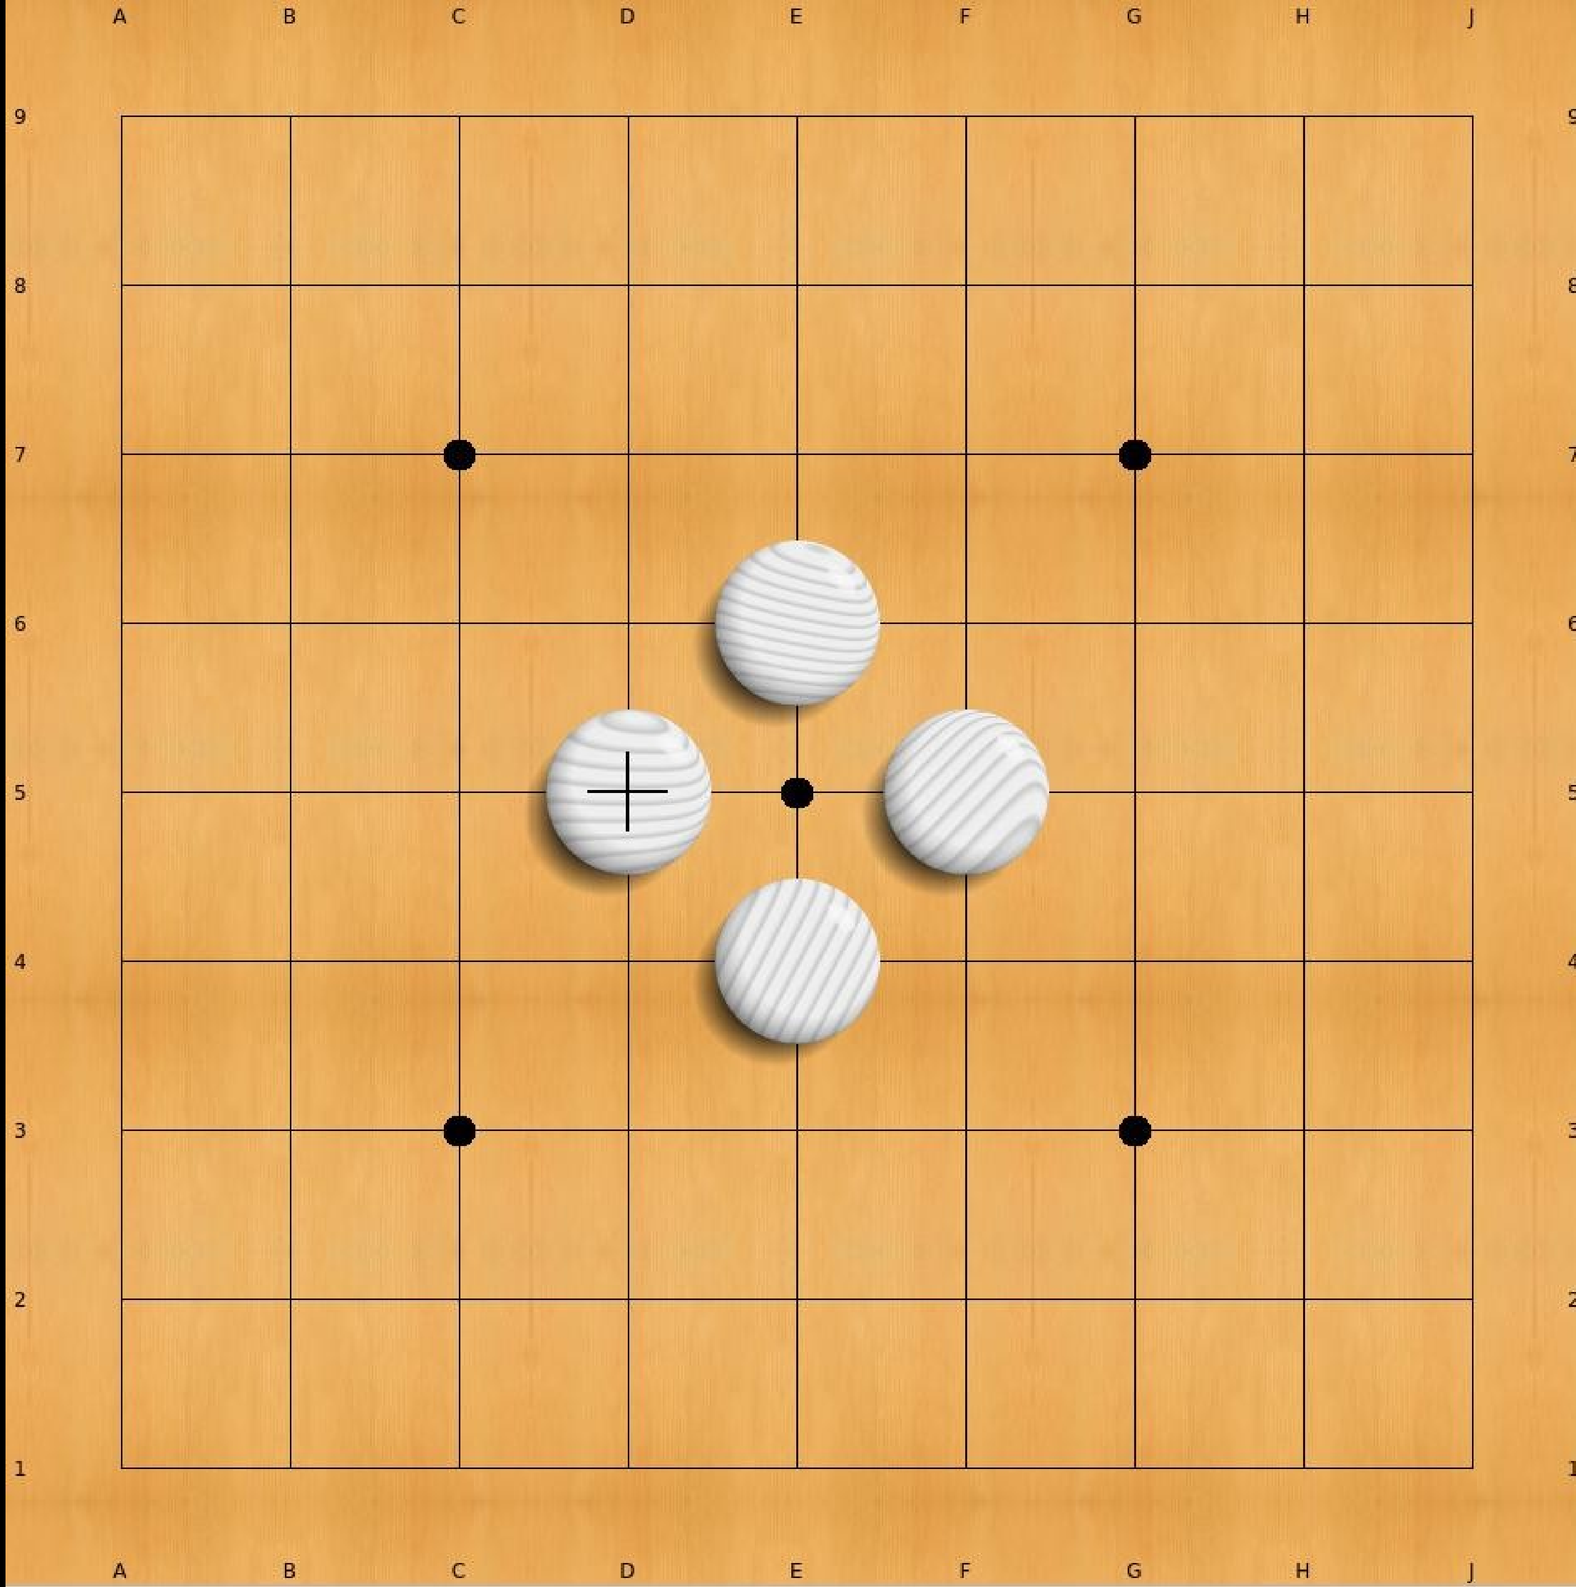
A B C D E F G H J
9 8 7 6 5 4 3 2 1

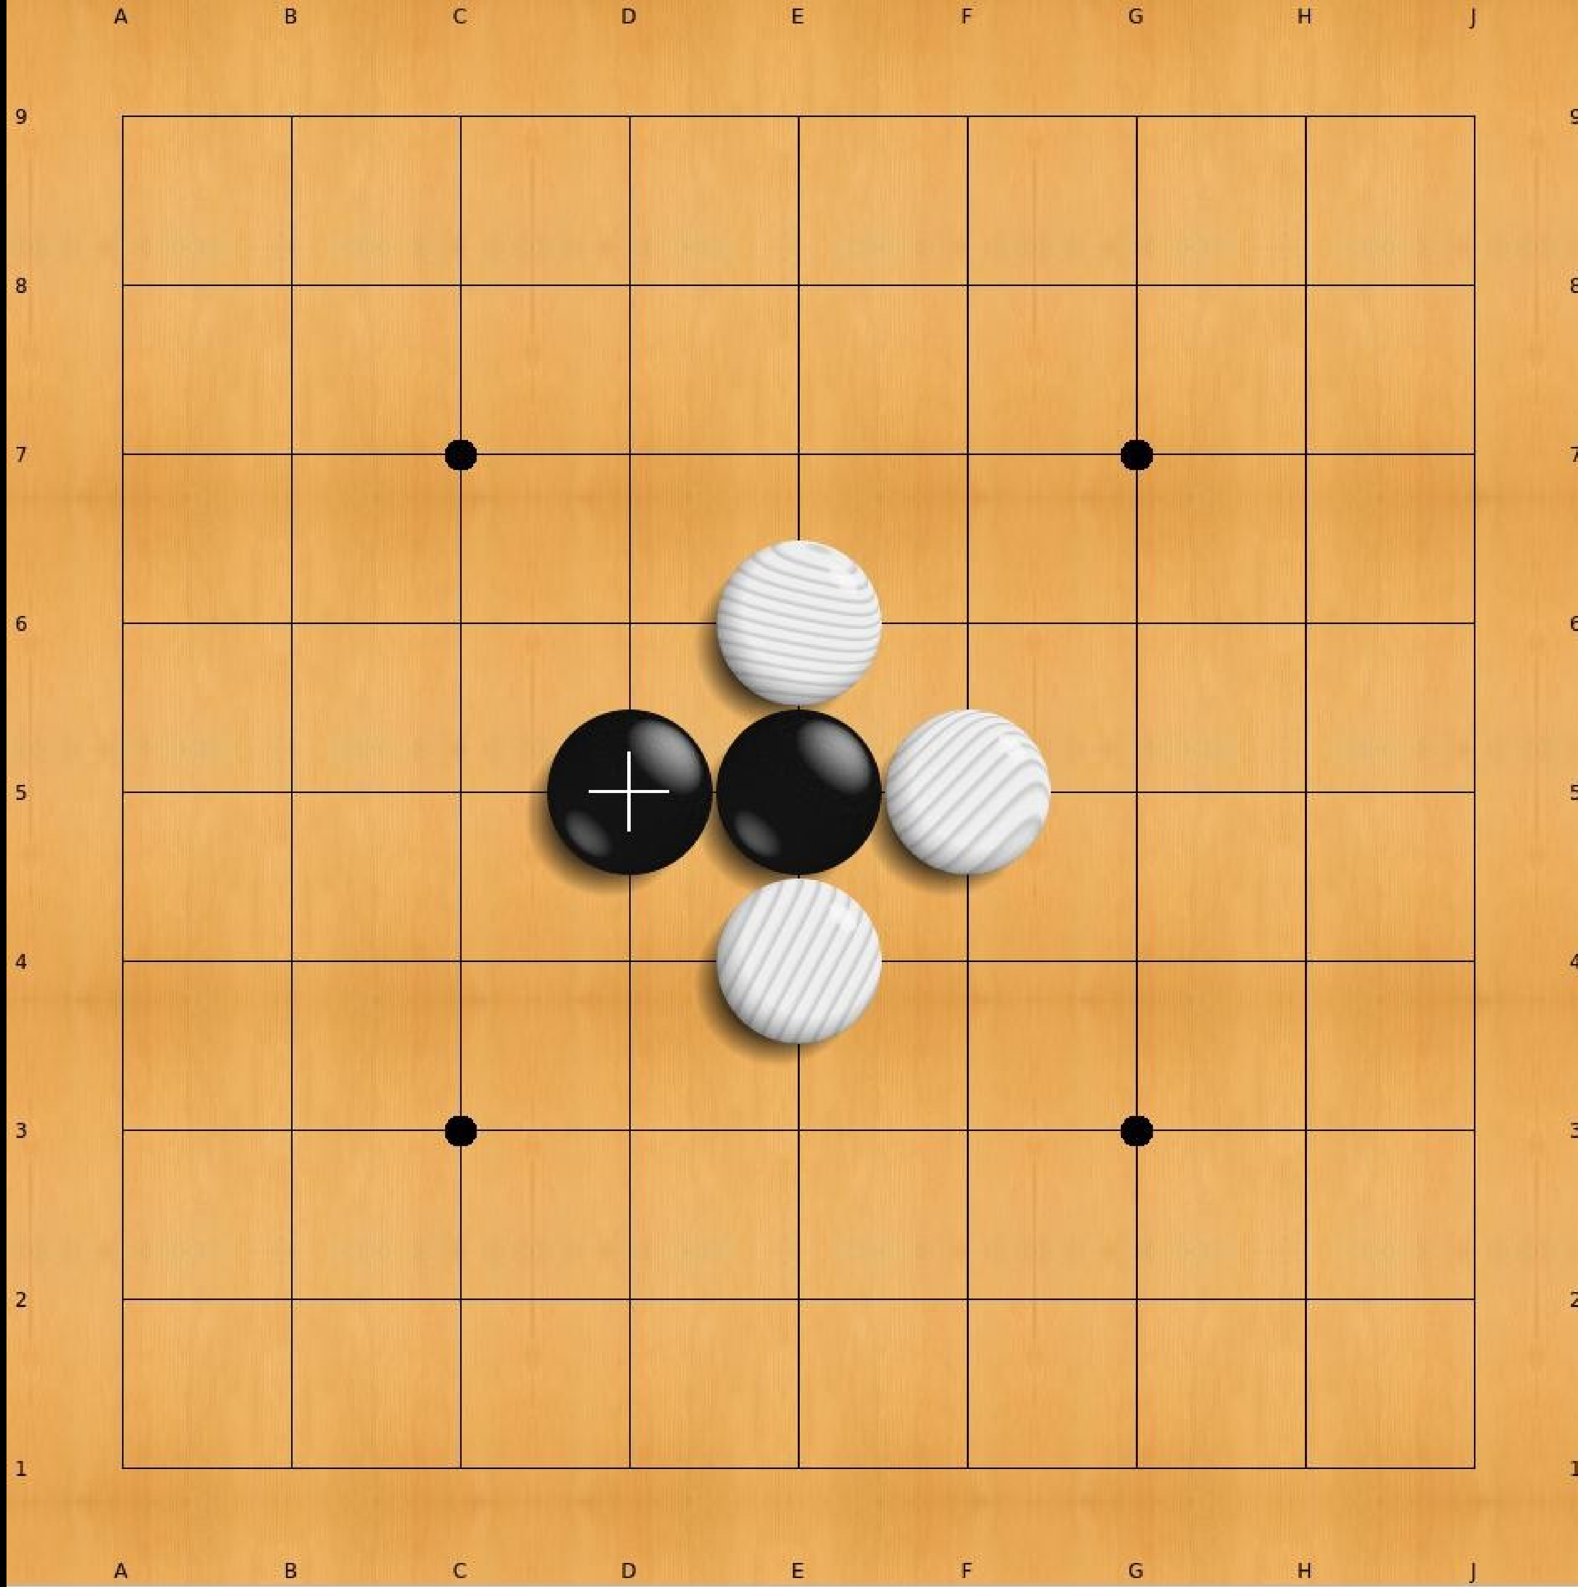
A B C D E F G H J
9 8 7 6 5 4 3 2 1

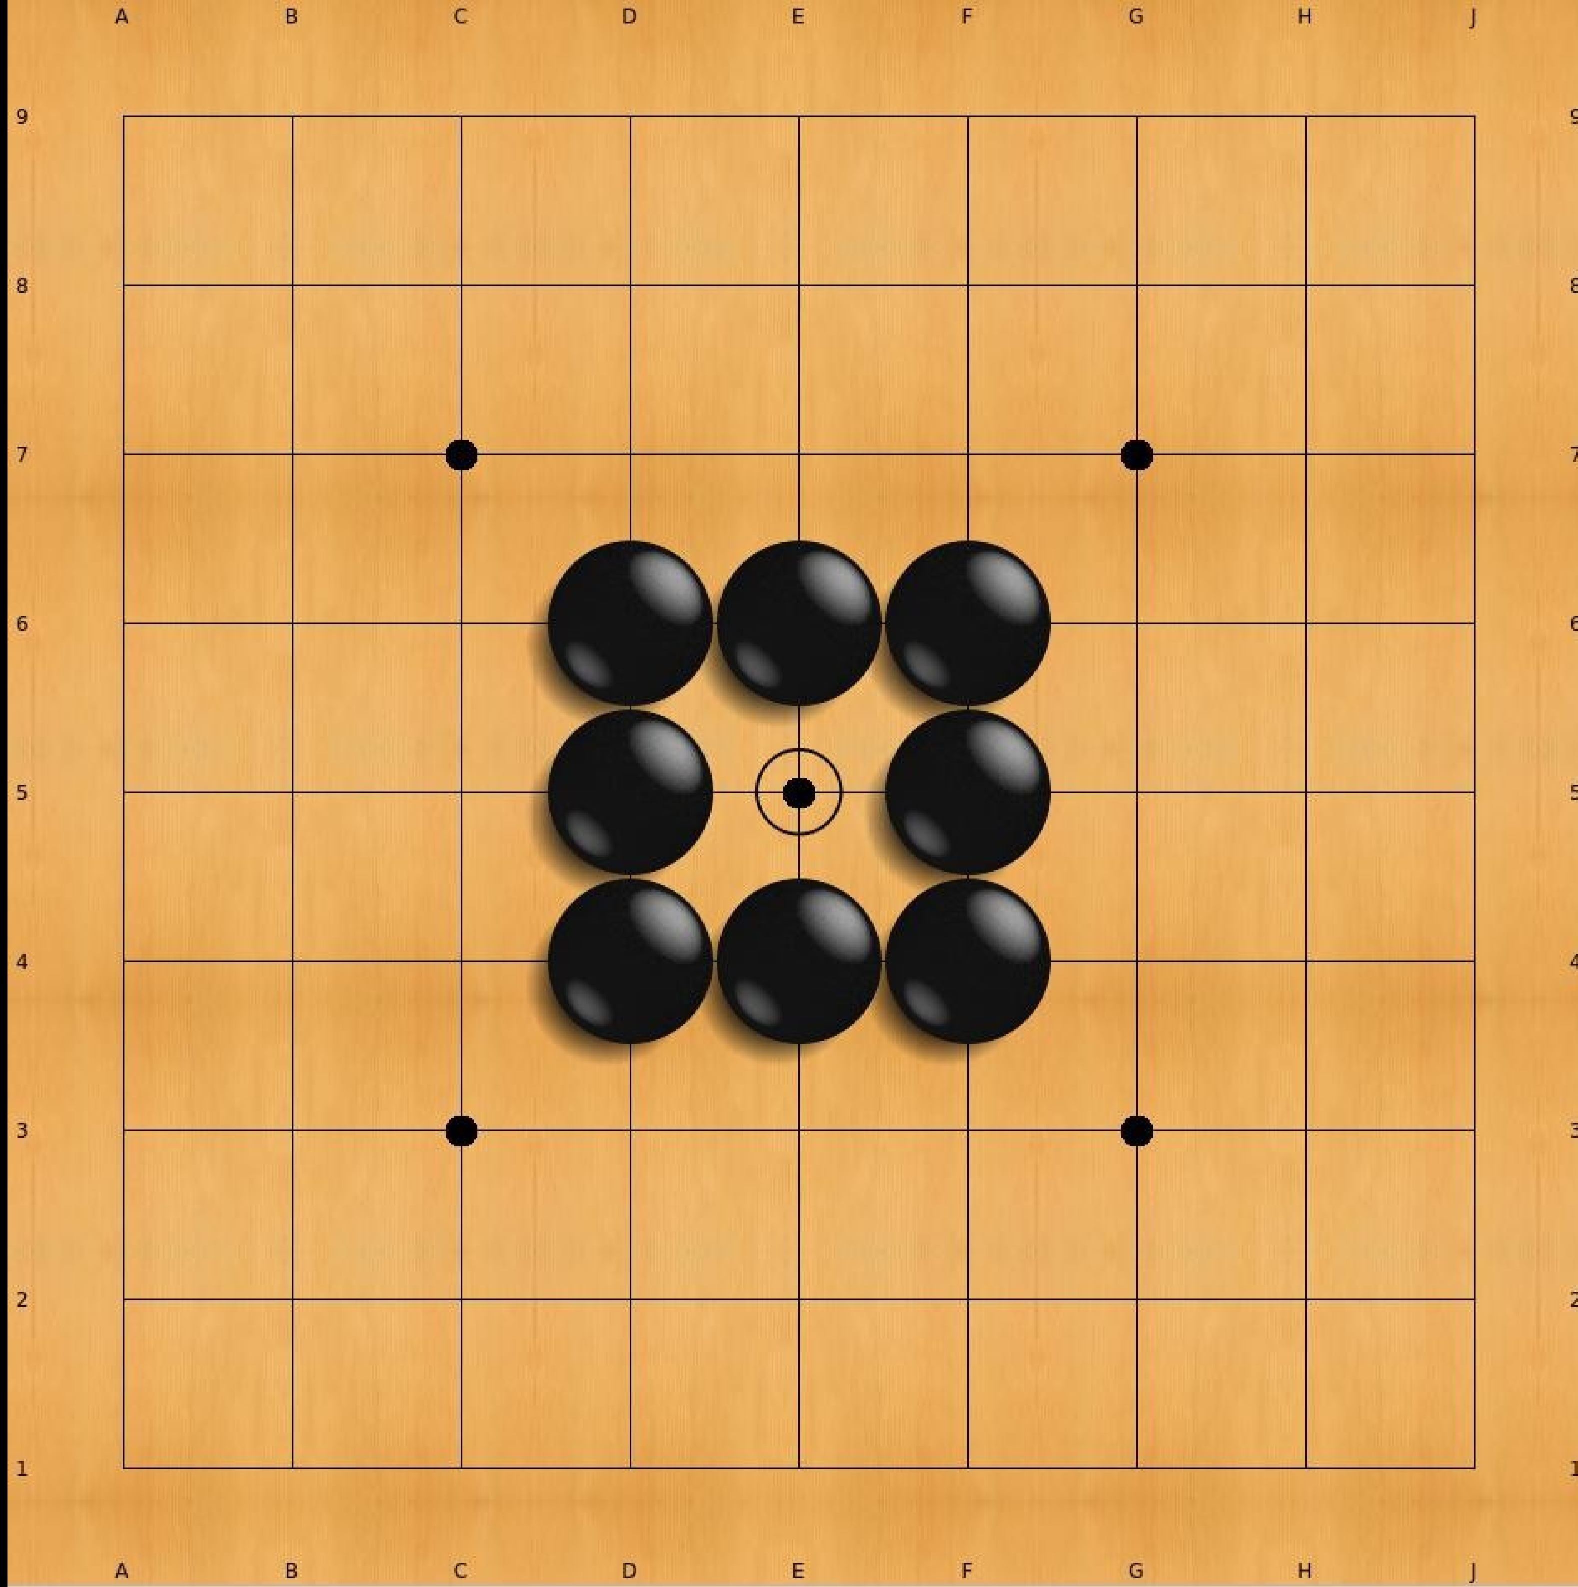
A B C D E F G H J
9 8 7 6 5 4 3 2 1

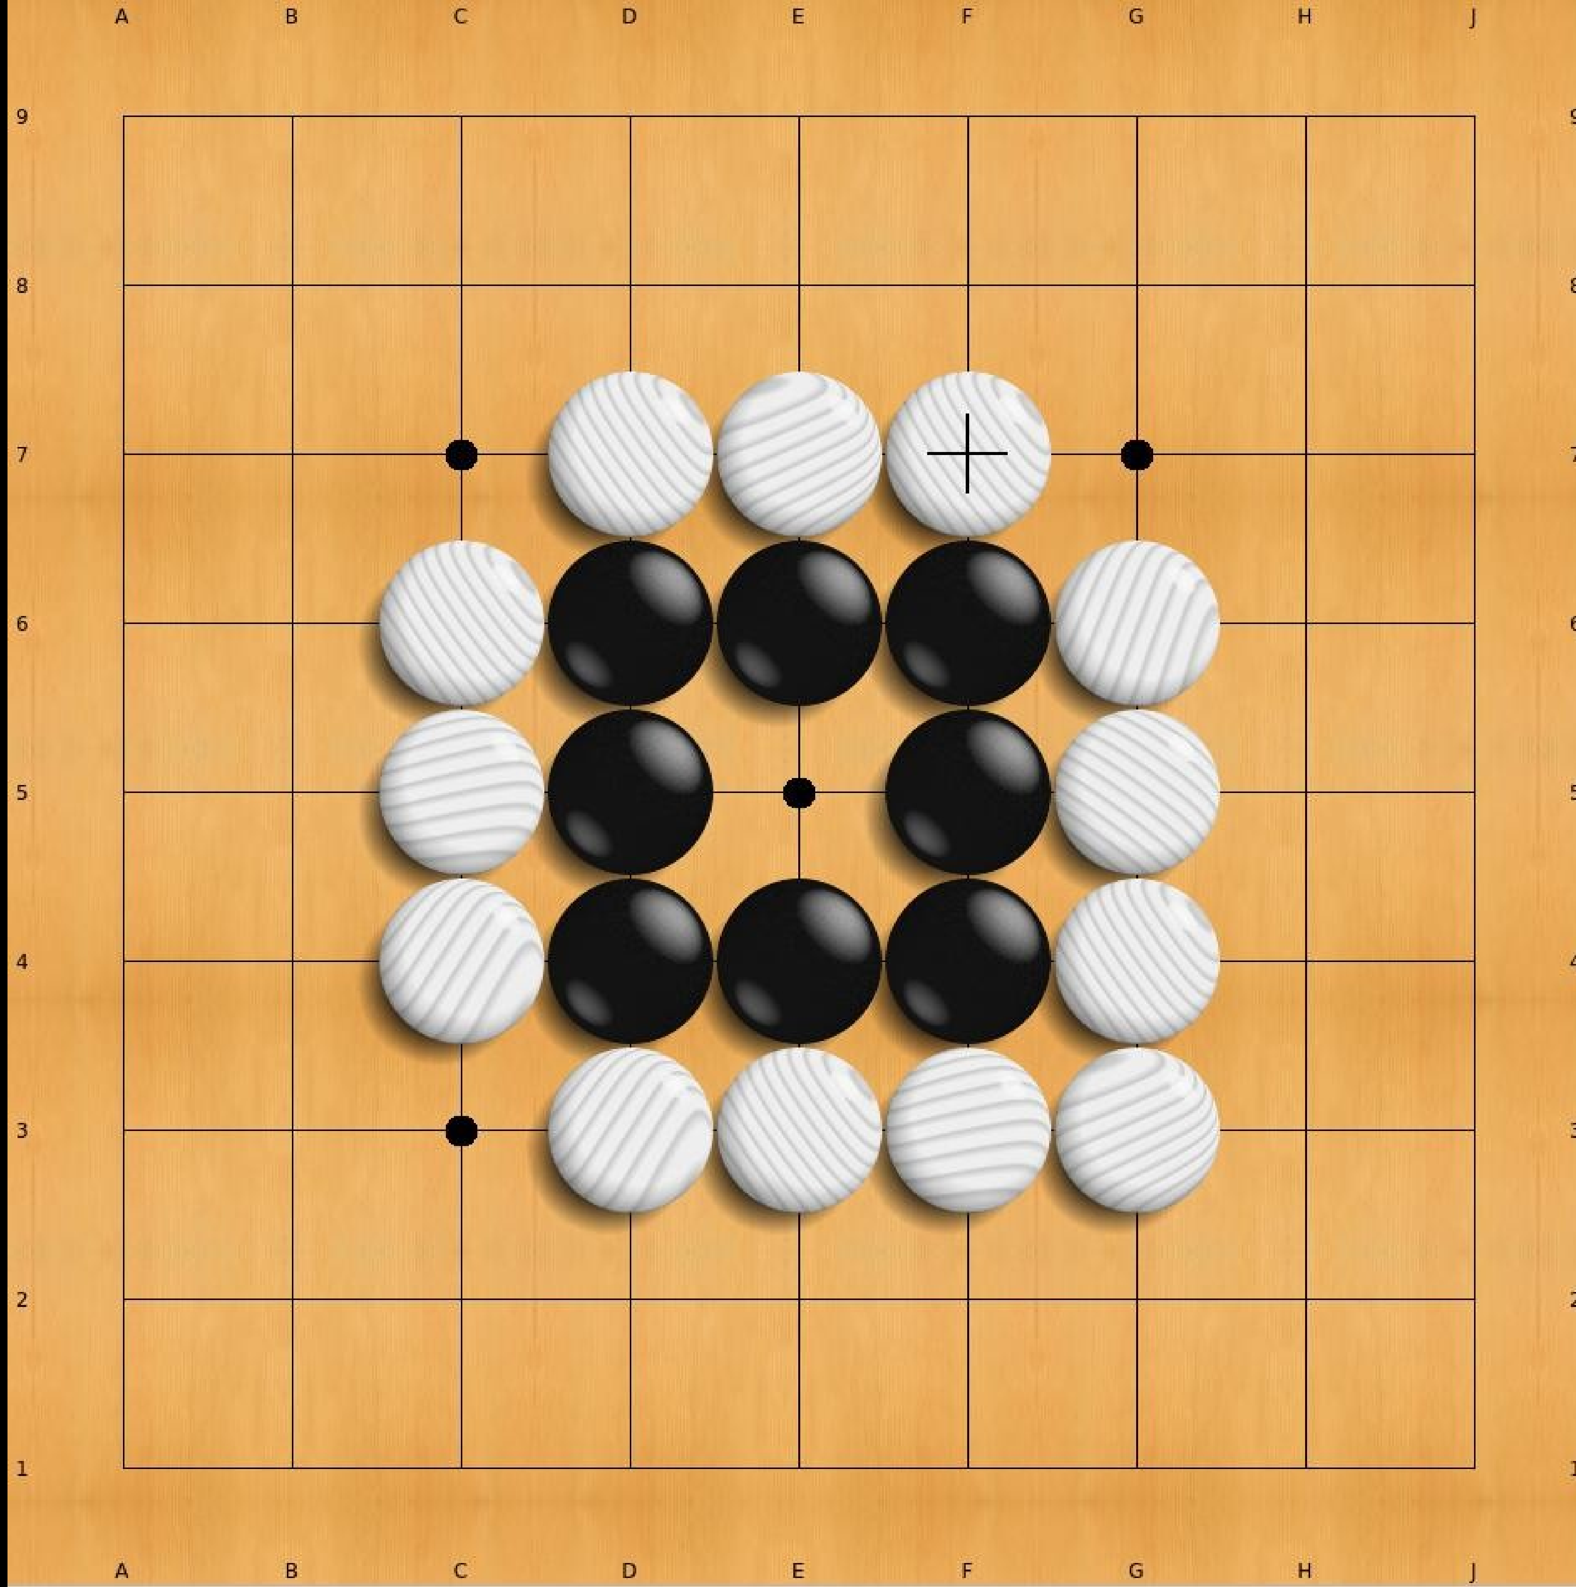
A B C D E F G H J
9 8 7 6 5 4 3 2 1

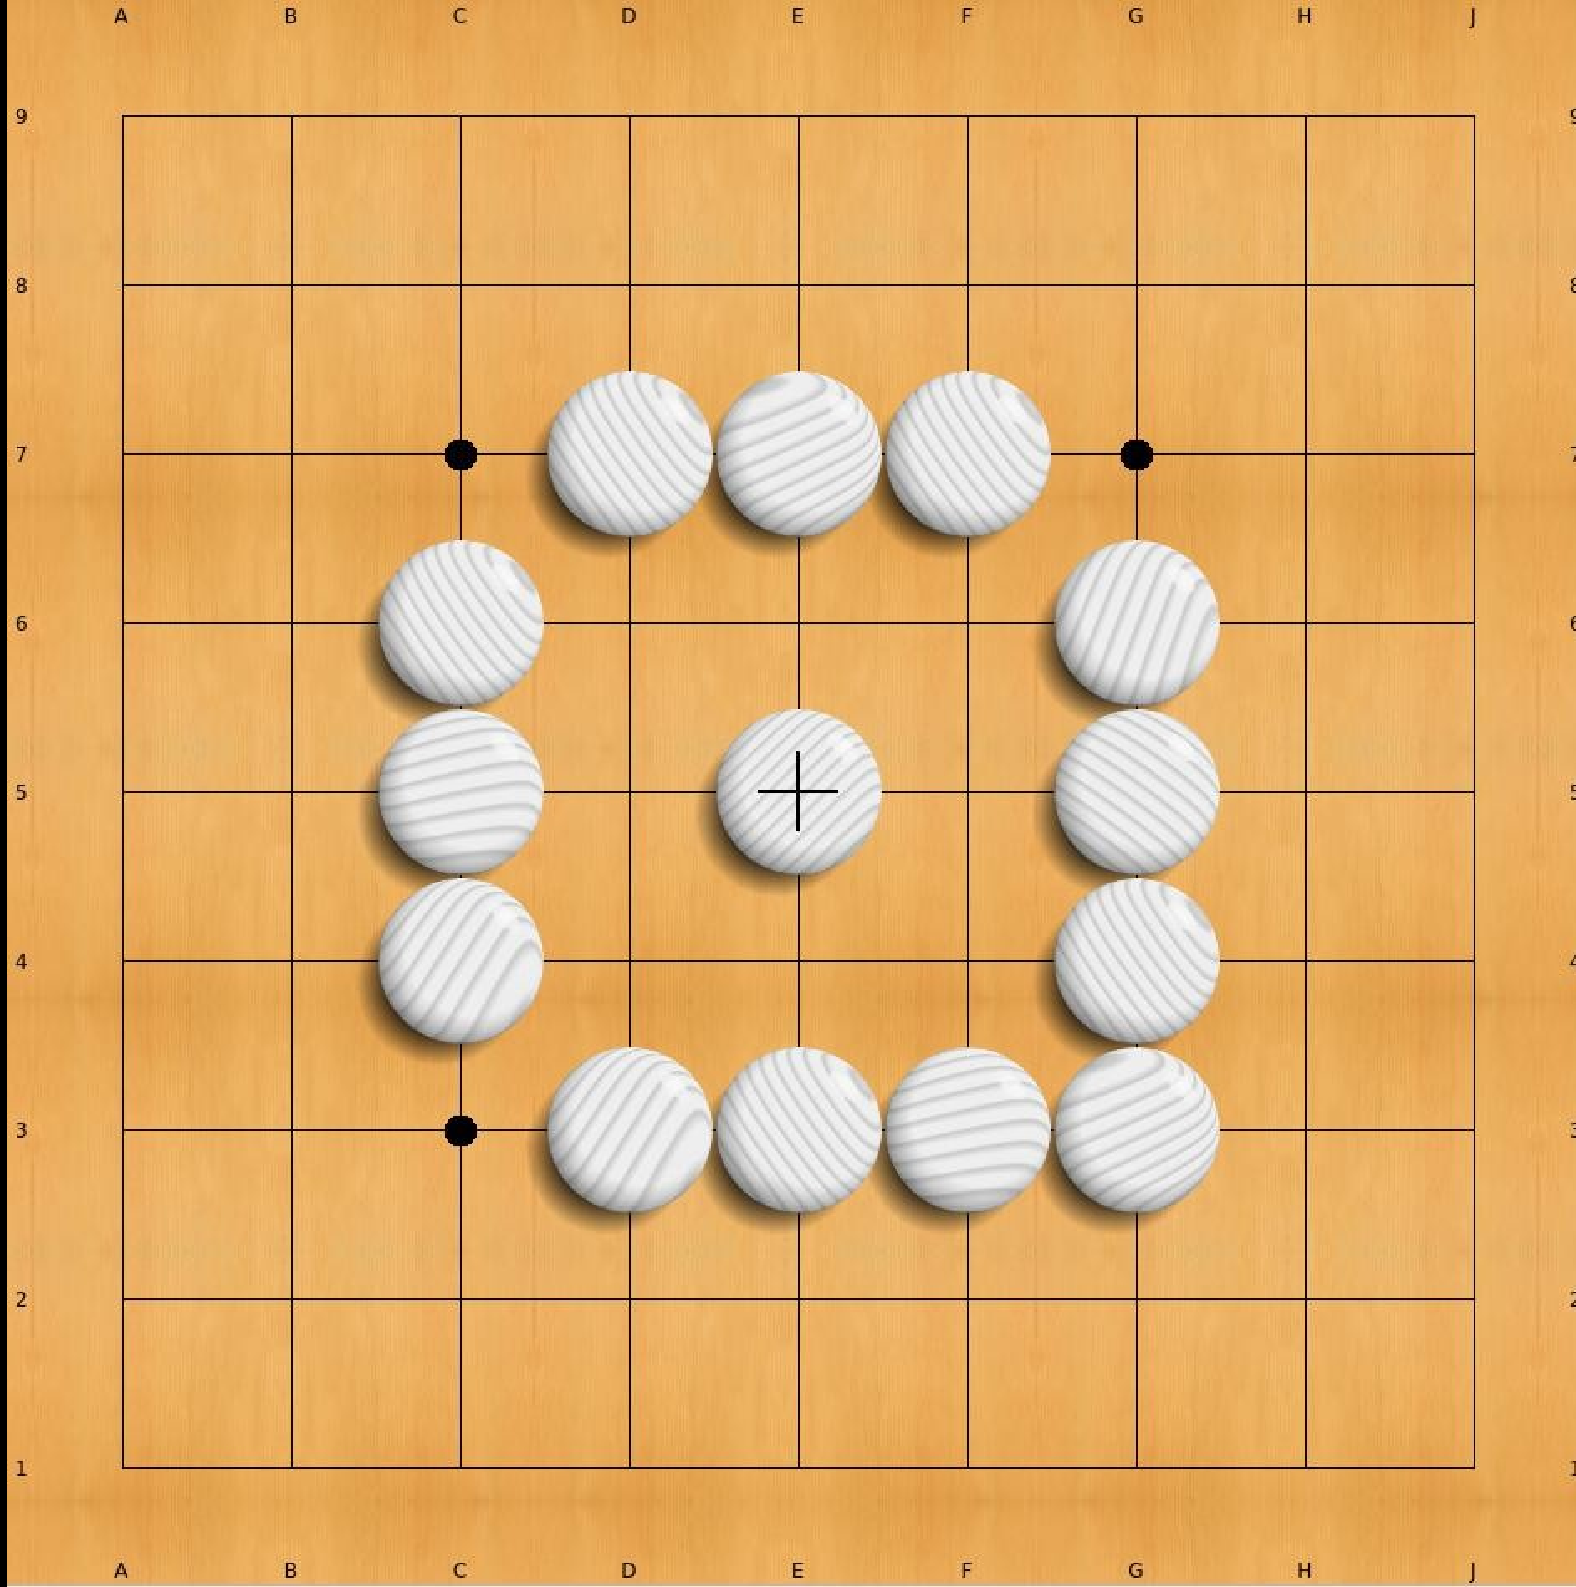
A B C D E F G H J
9 8 7 6 5 4 3 2 1

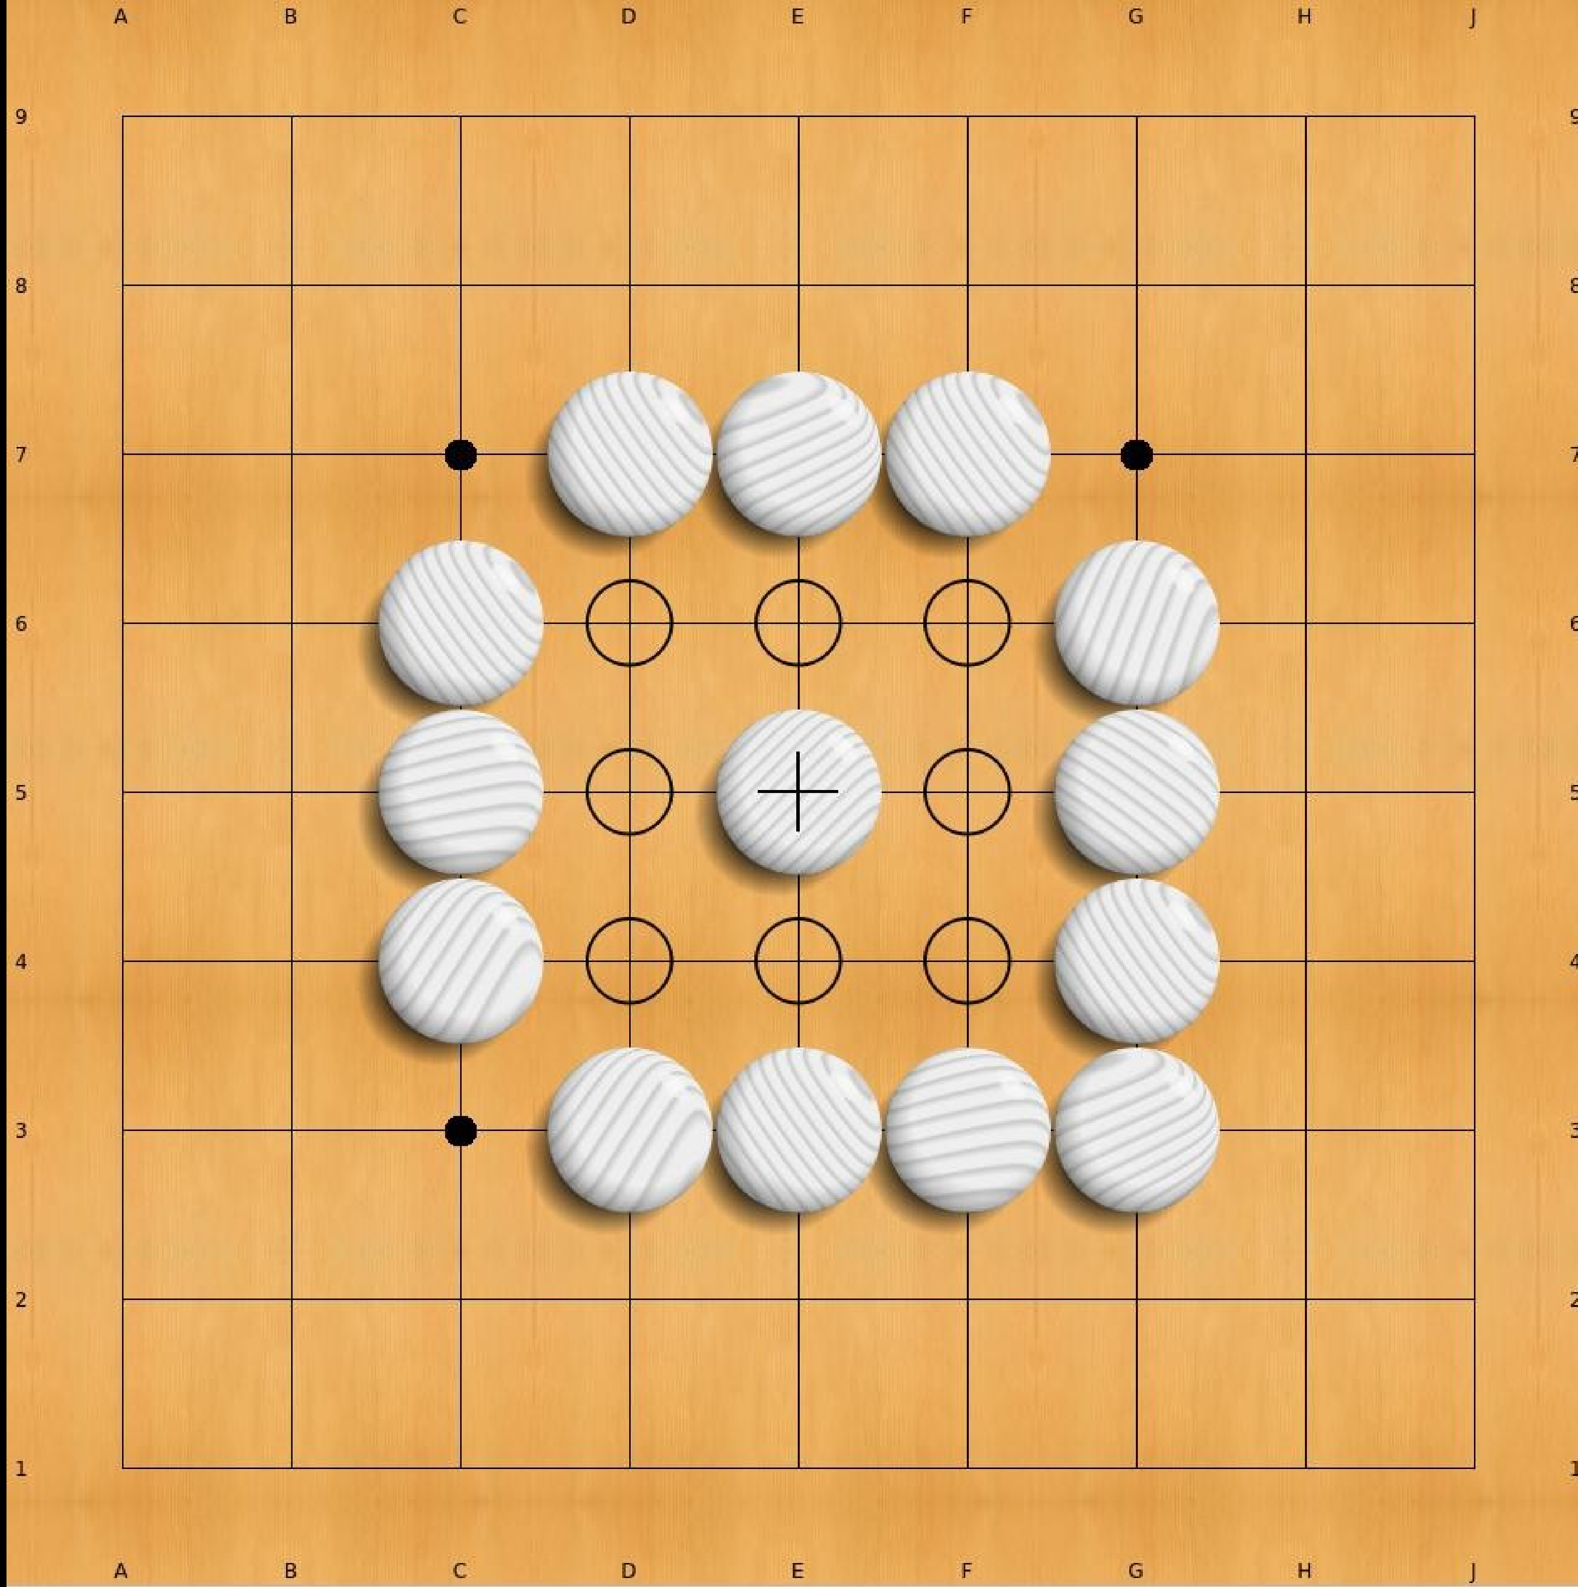
A B C D E F G H J
9 8 7 6 5 4 3 2 1

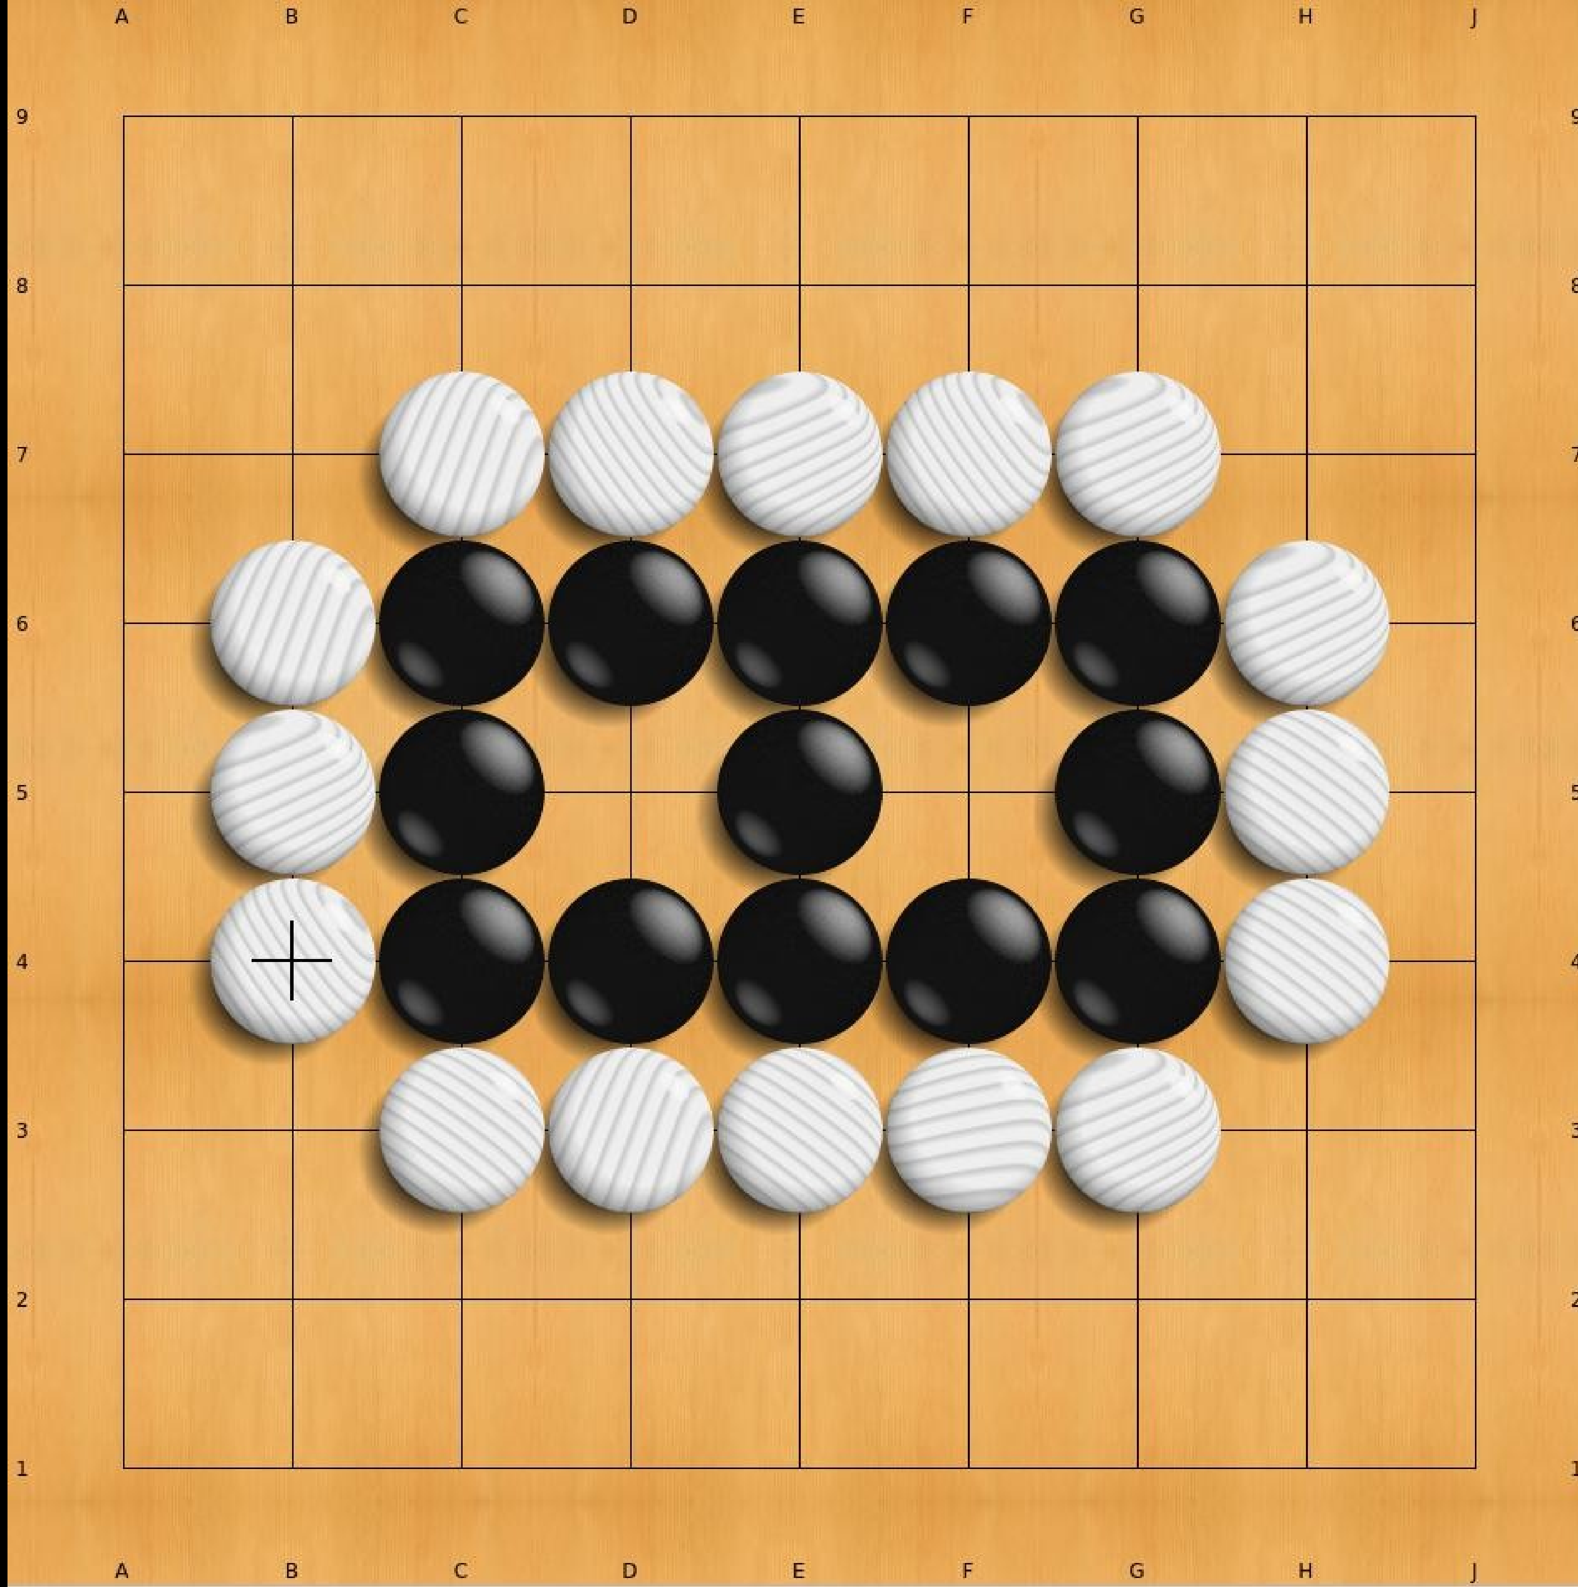
A B C D E F G H J
9 8 7 6 5 4 3 2 1

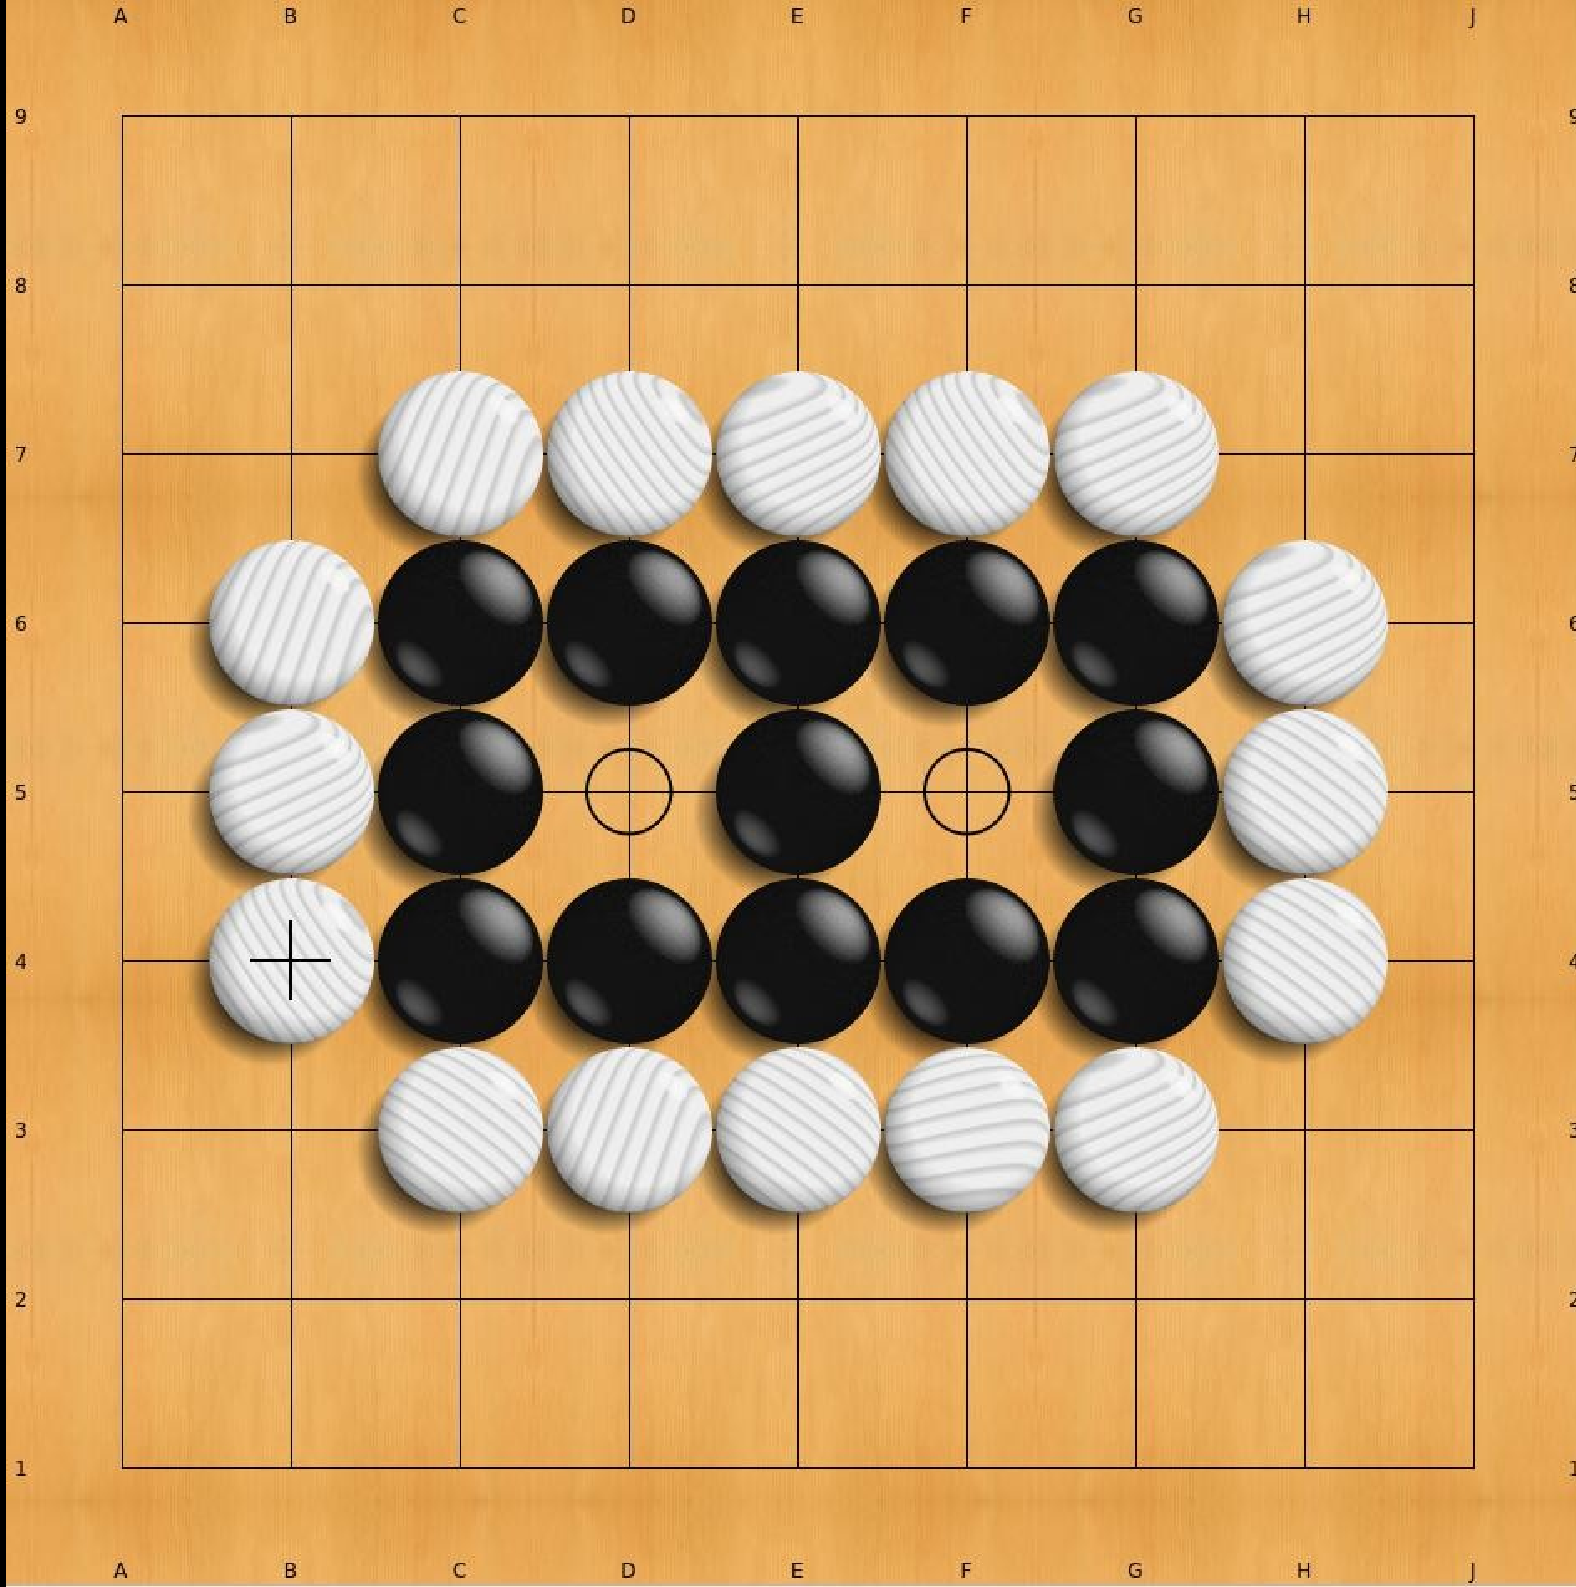
A B C D E F G H J
9 8 7 6 5 4 3 2 1


A B C D E F G H J
9 8 7 6 5 4 3 2 1


A B C D E F G H J
9 8 7 6 5 4 3 2 1

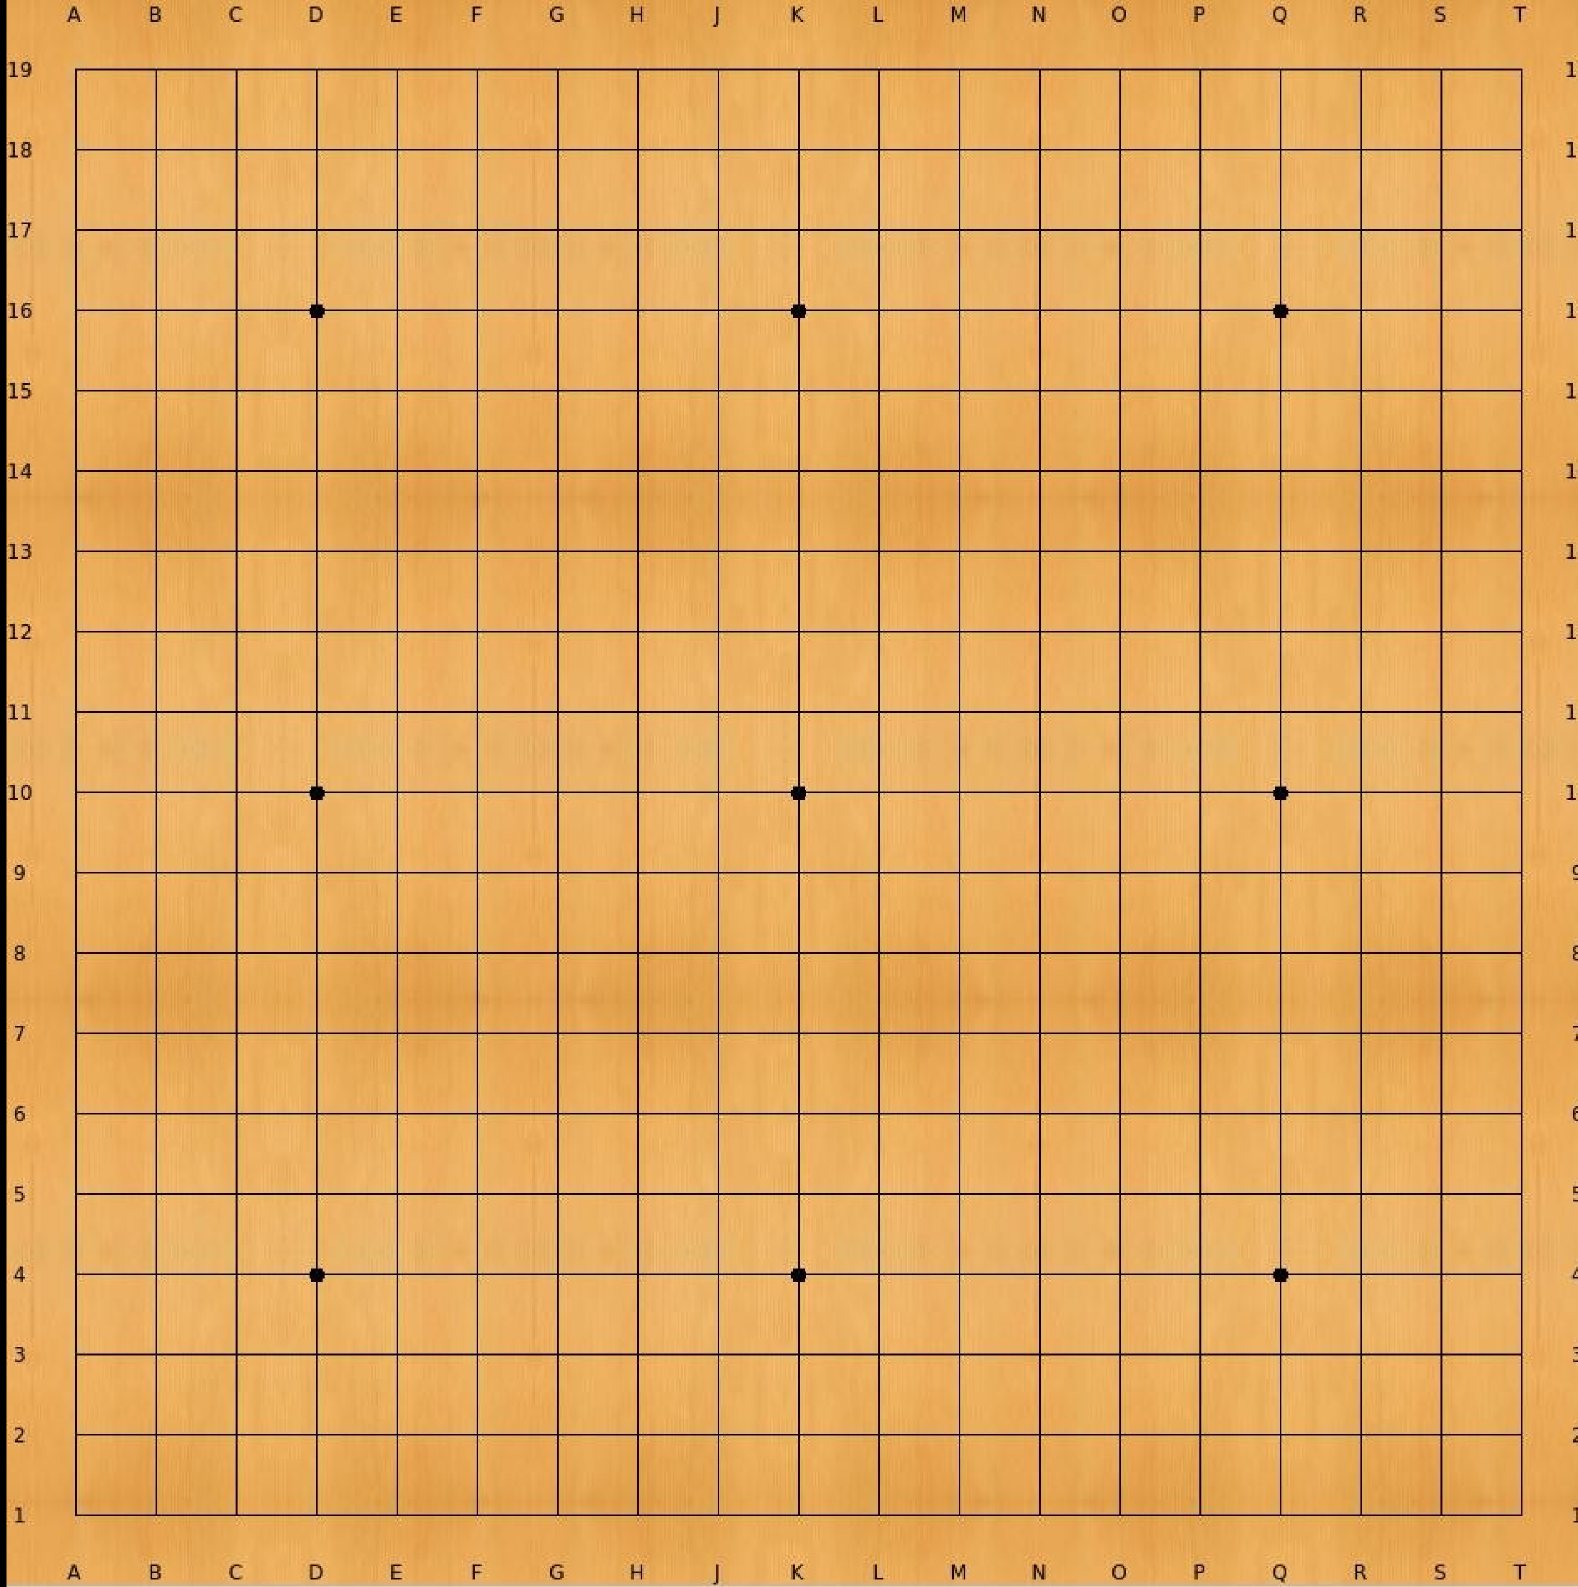
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

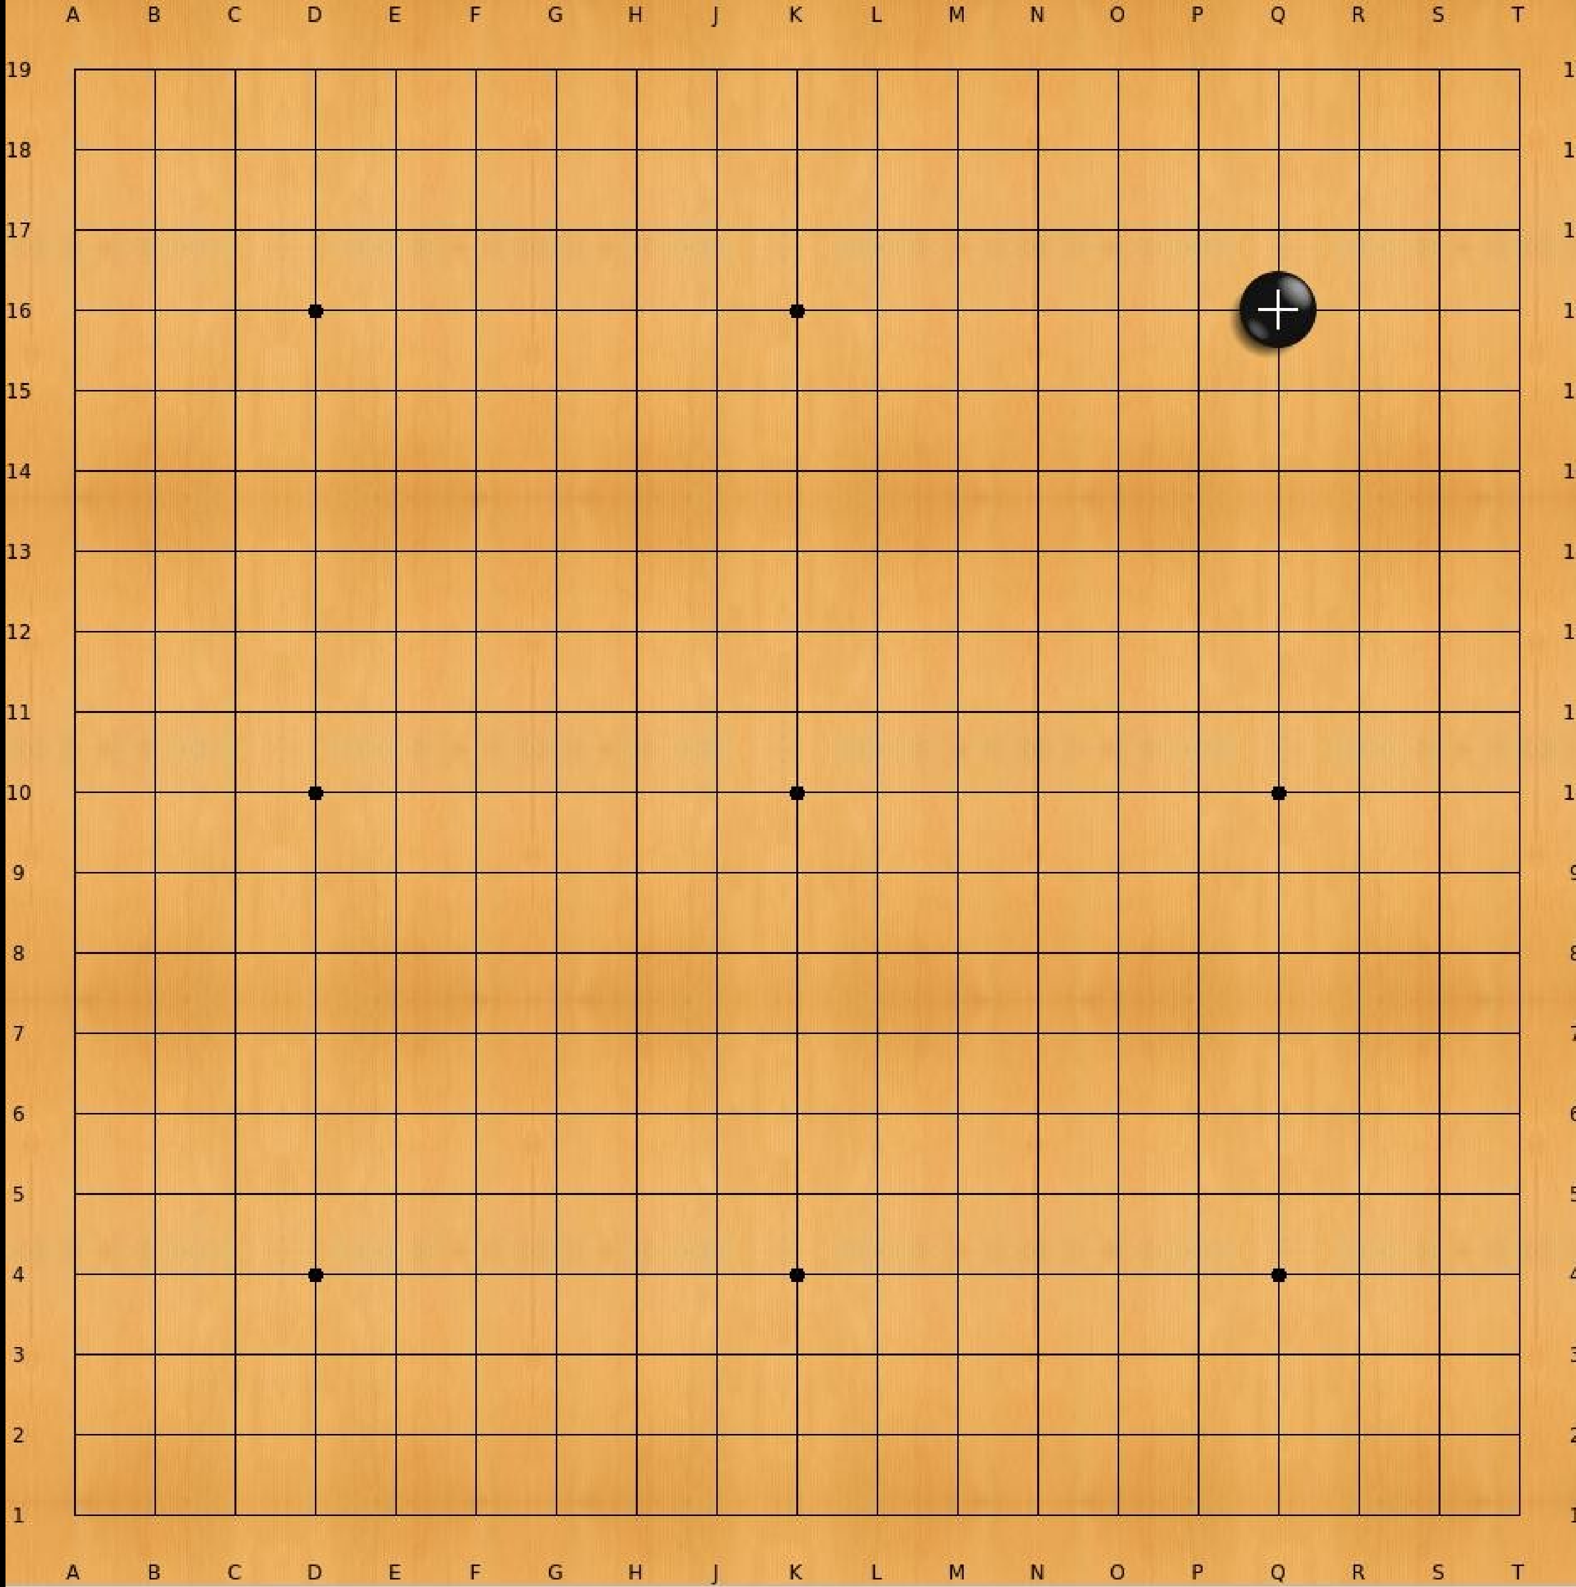
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

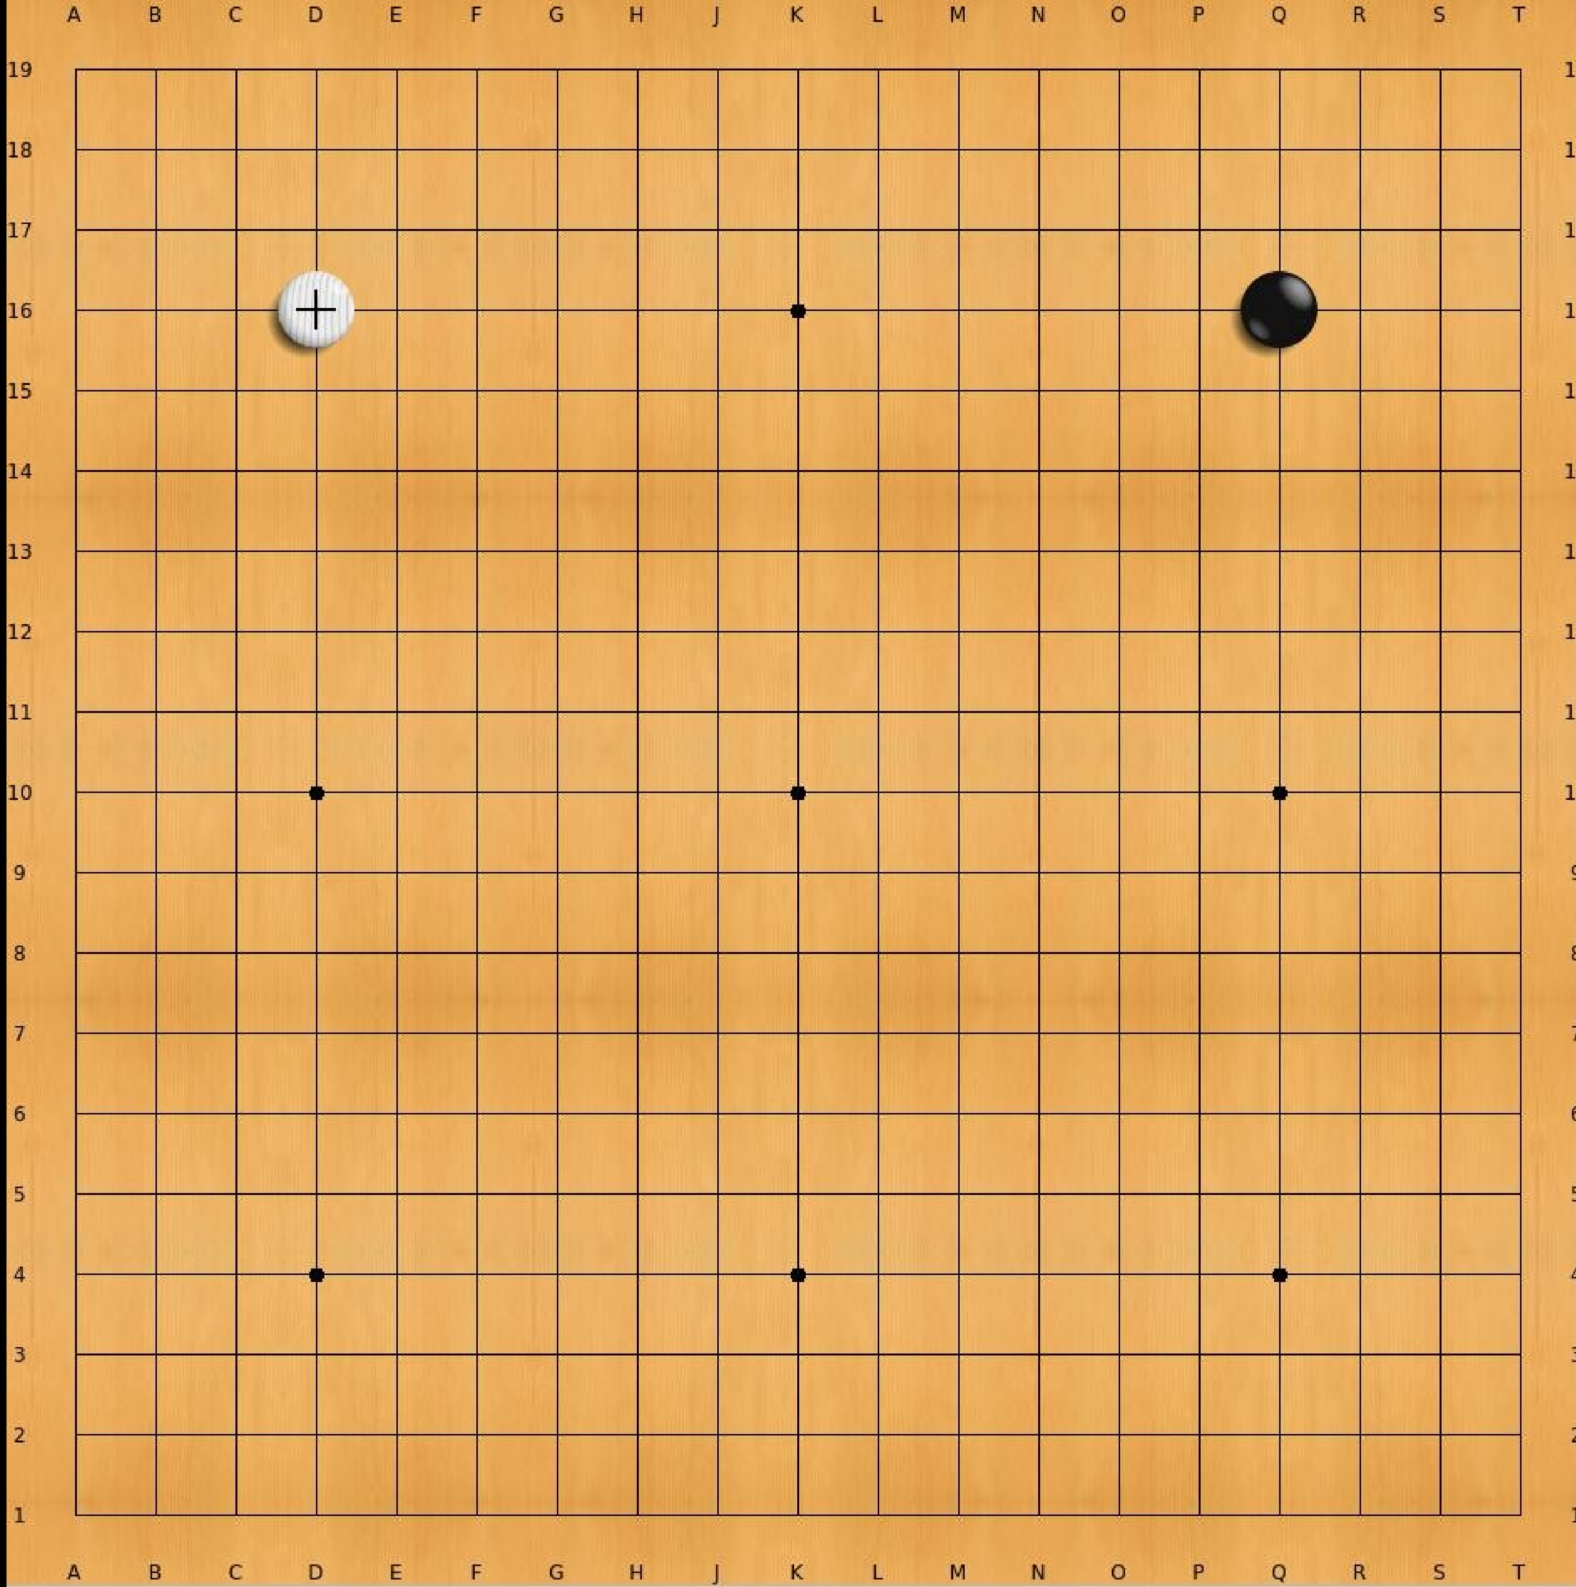
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

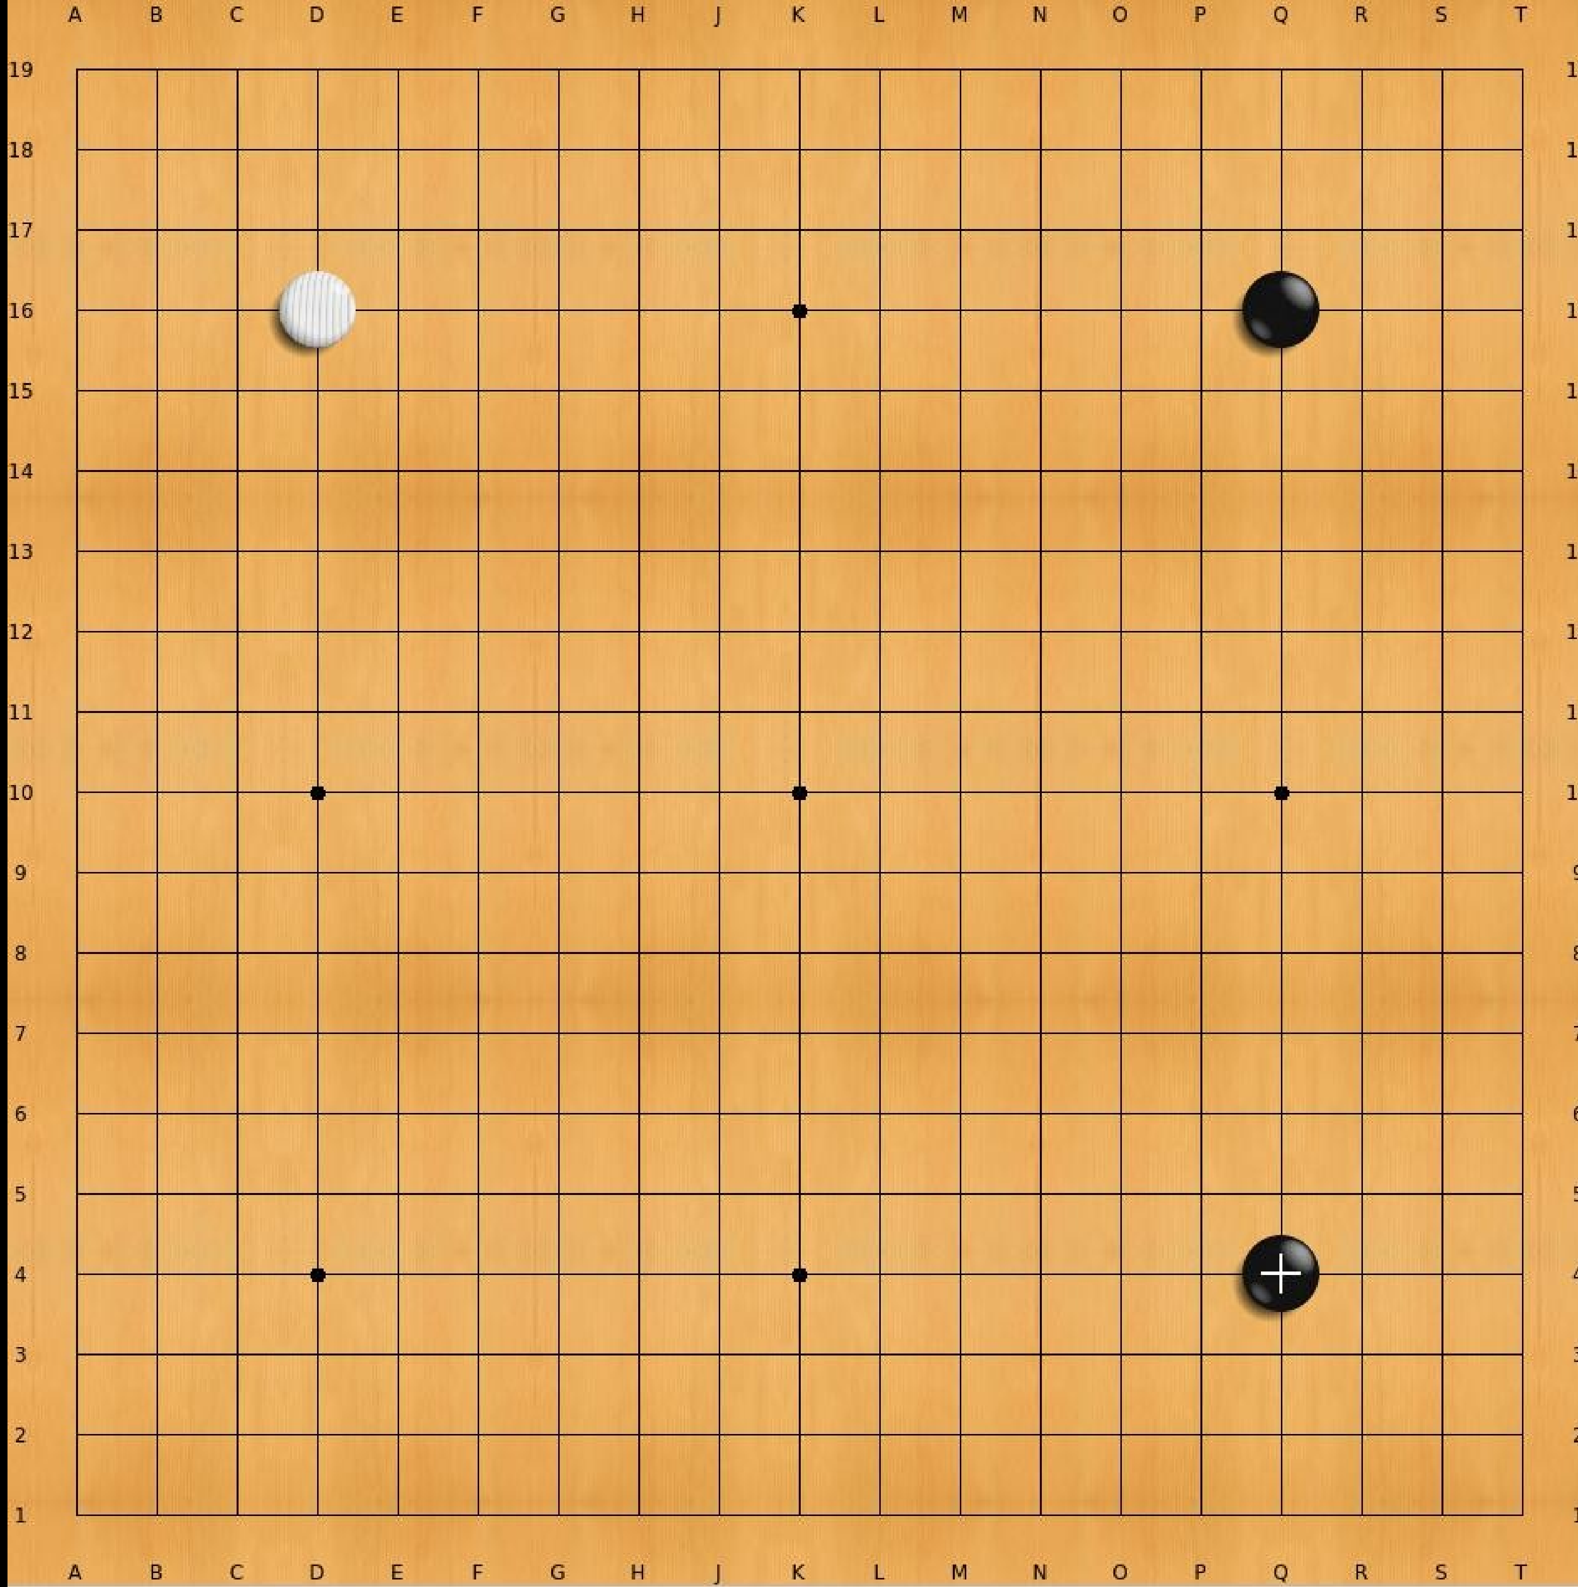
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

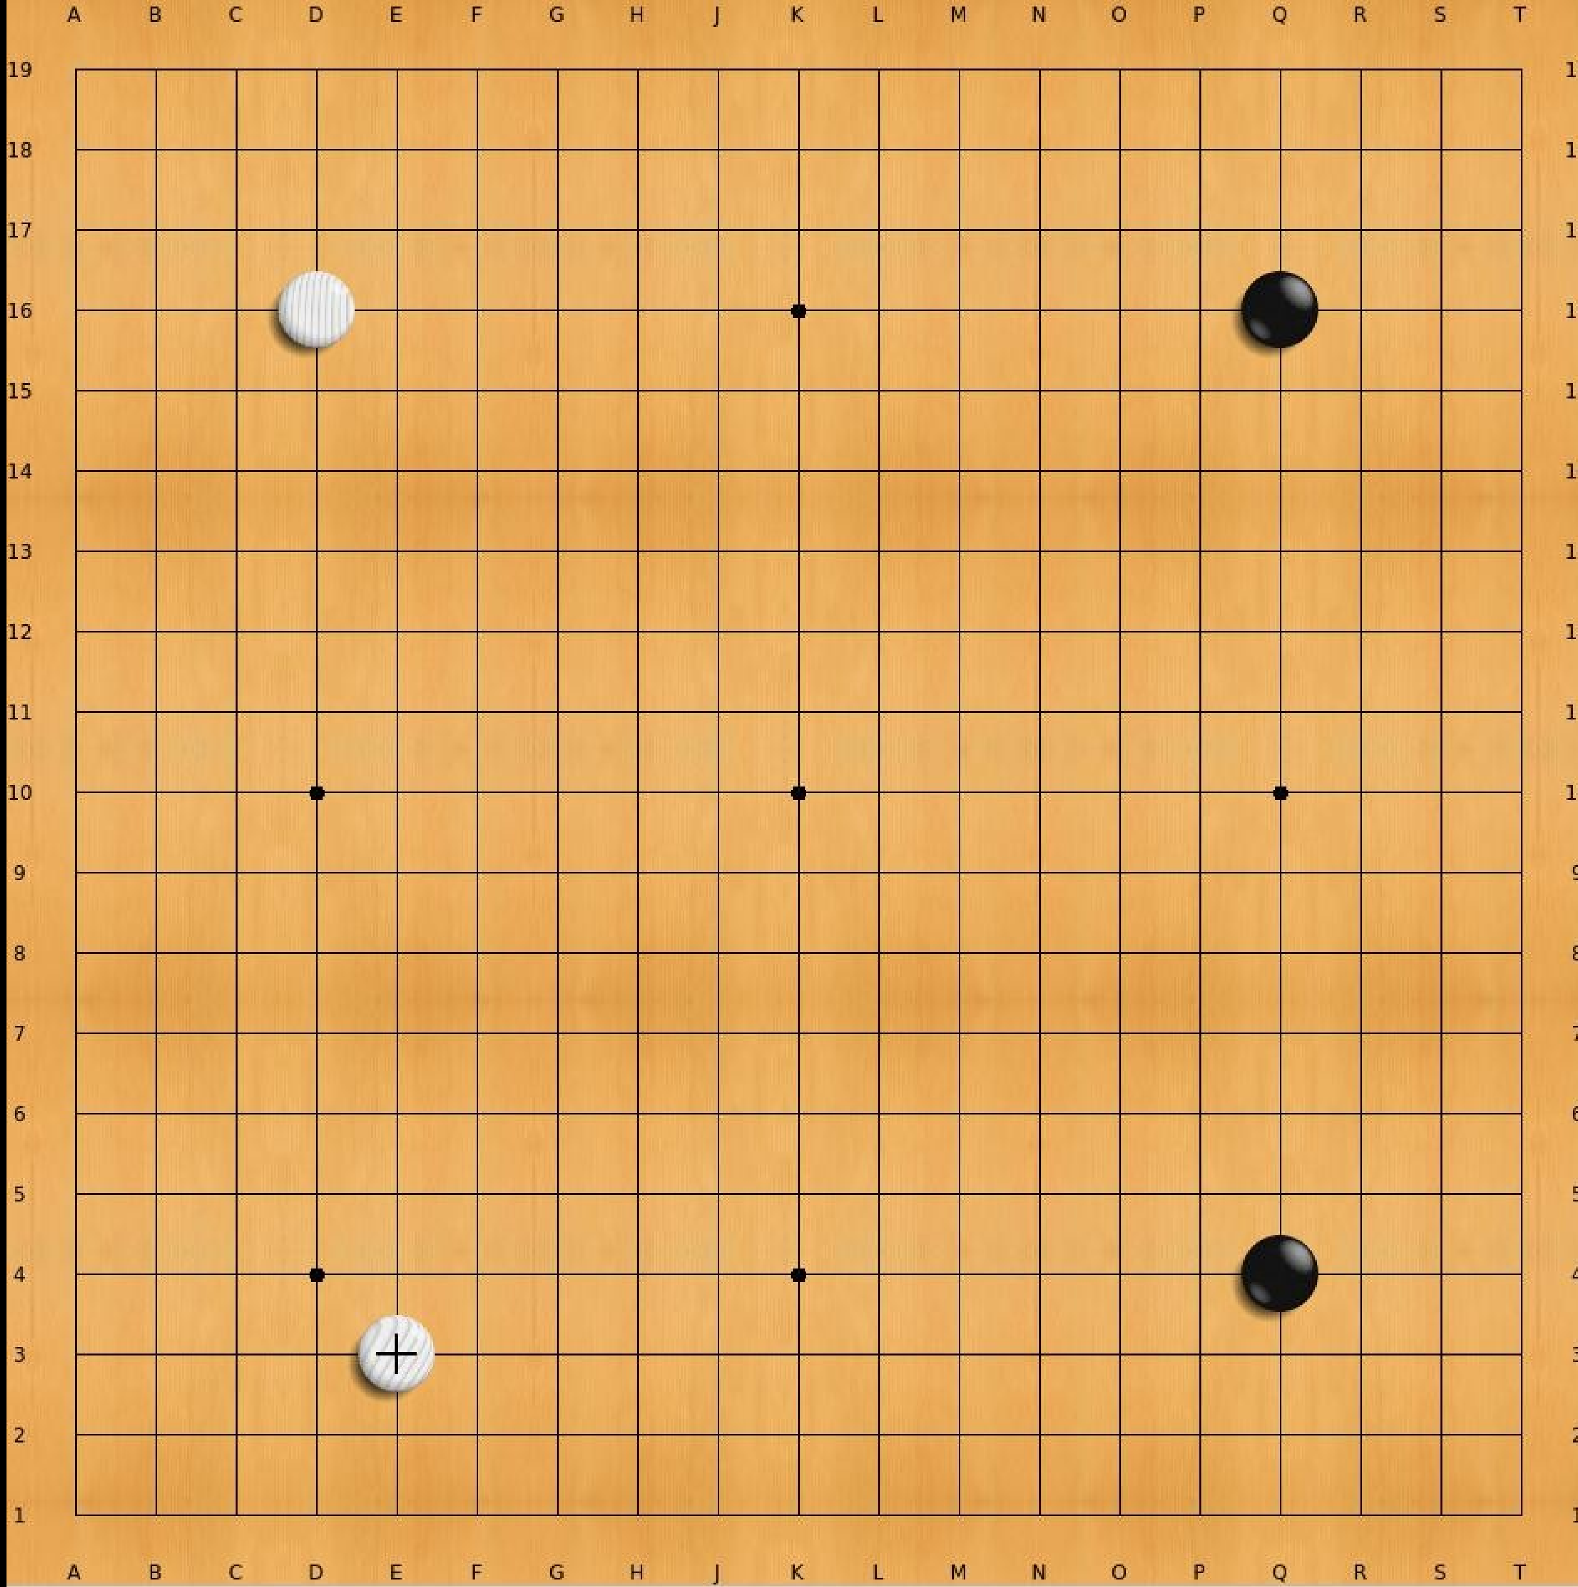
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

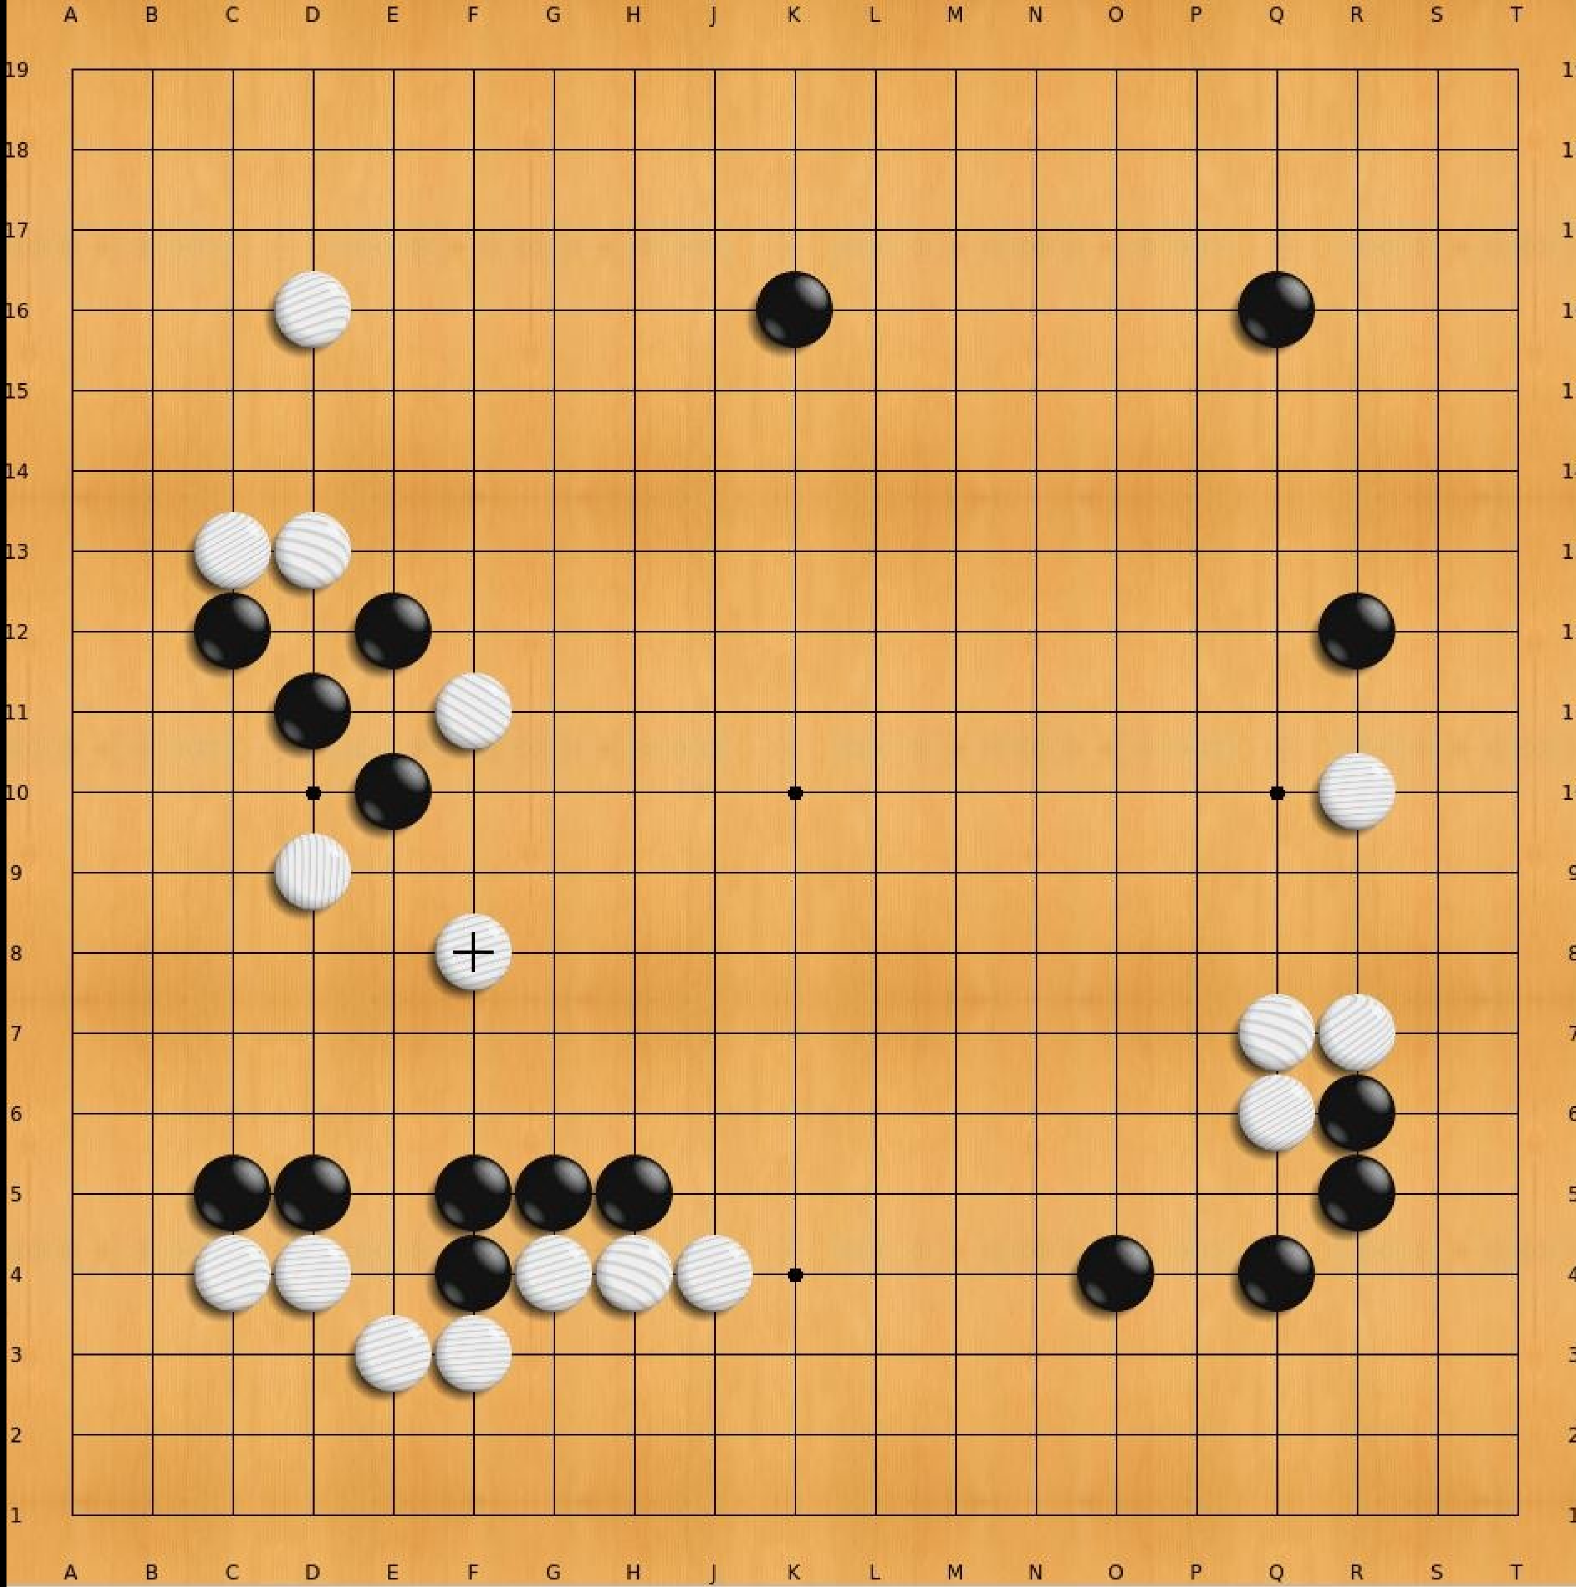
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

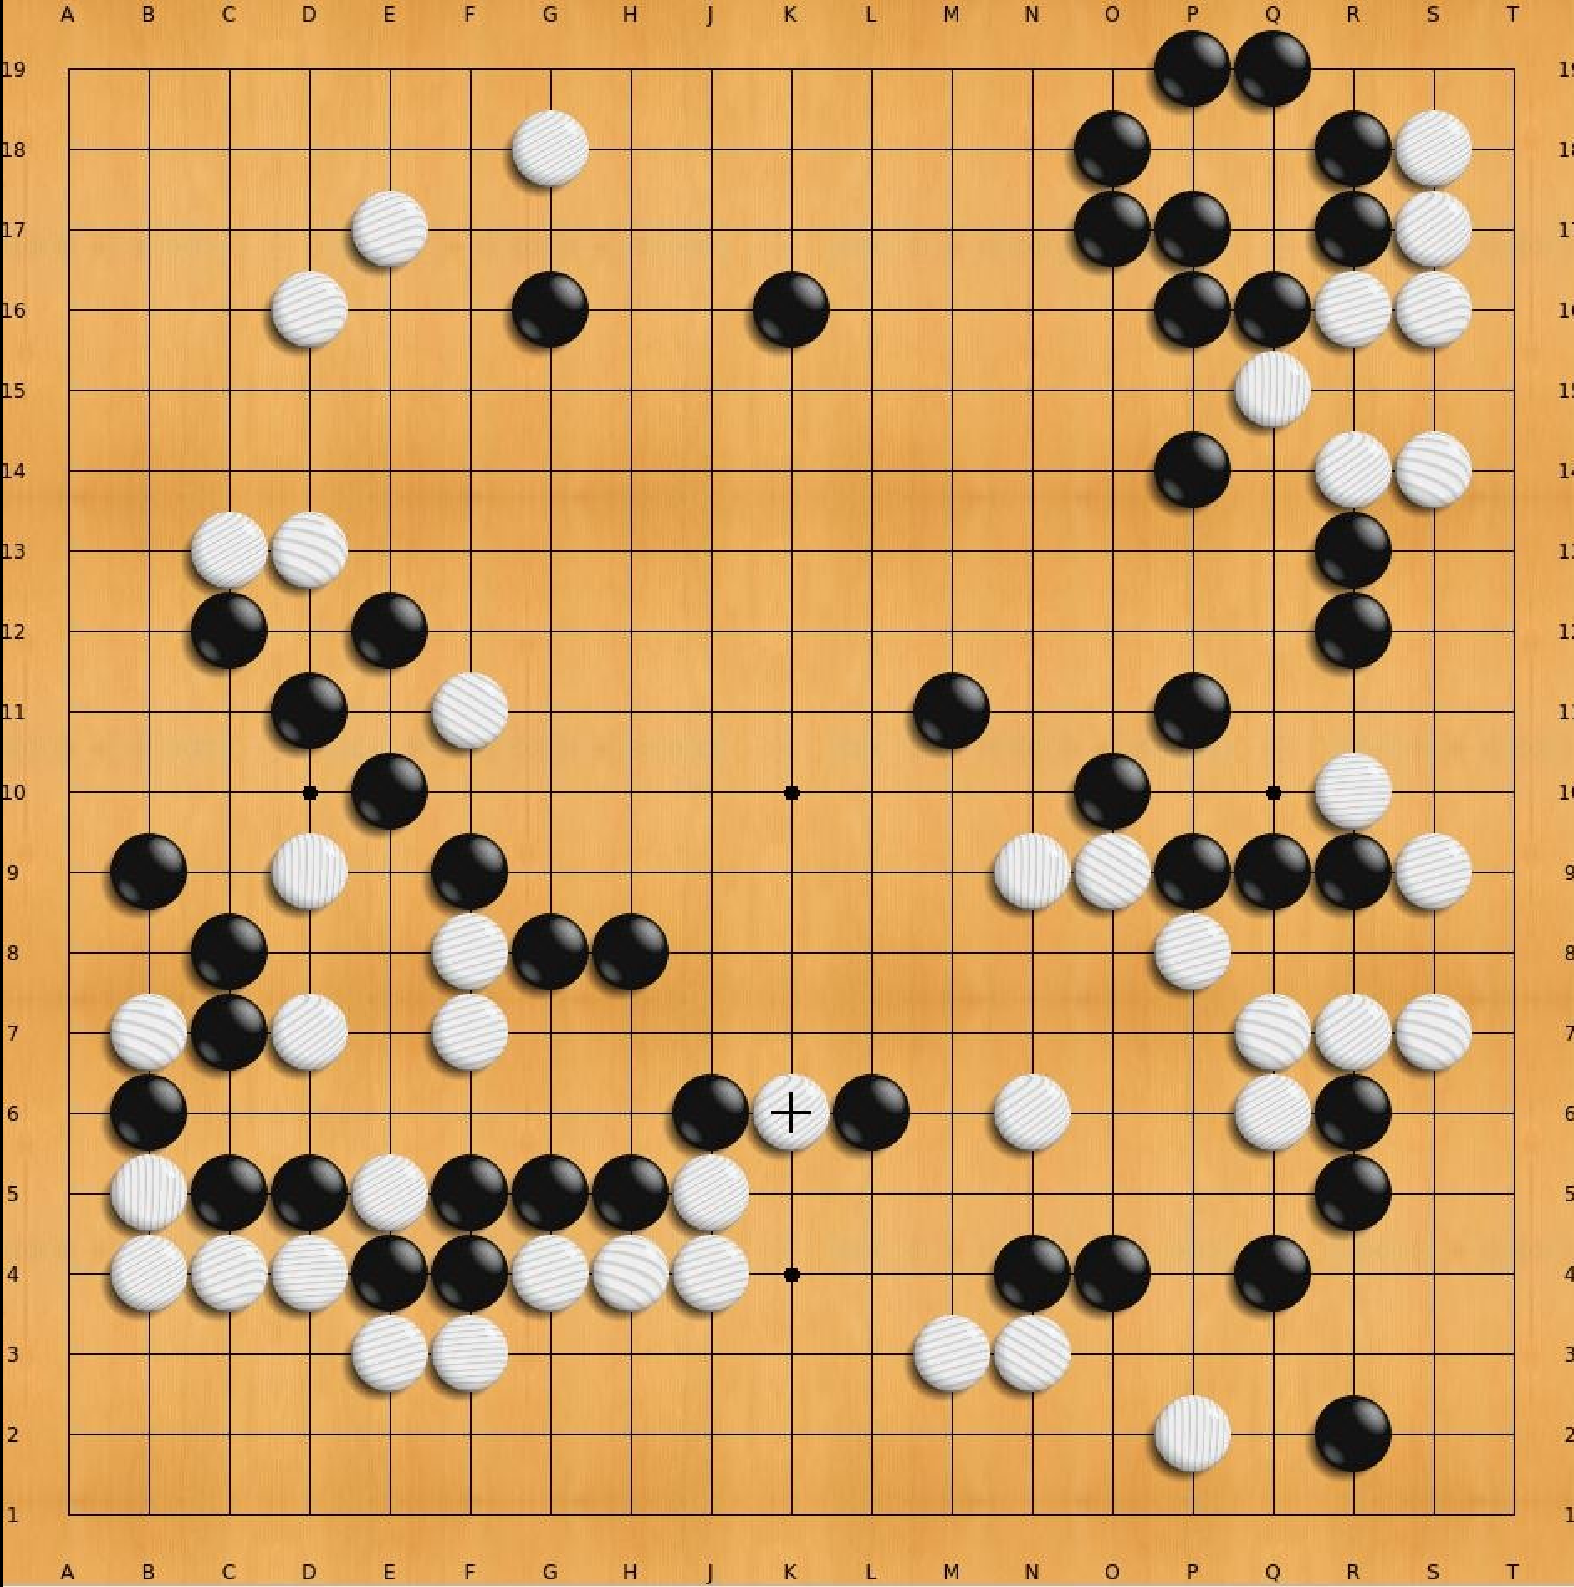
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

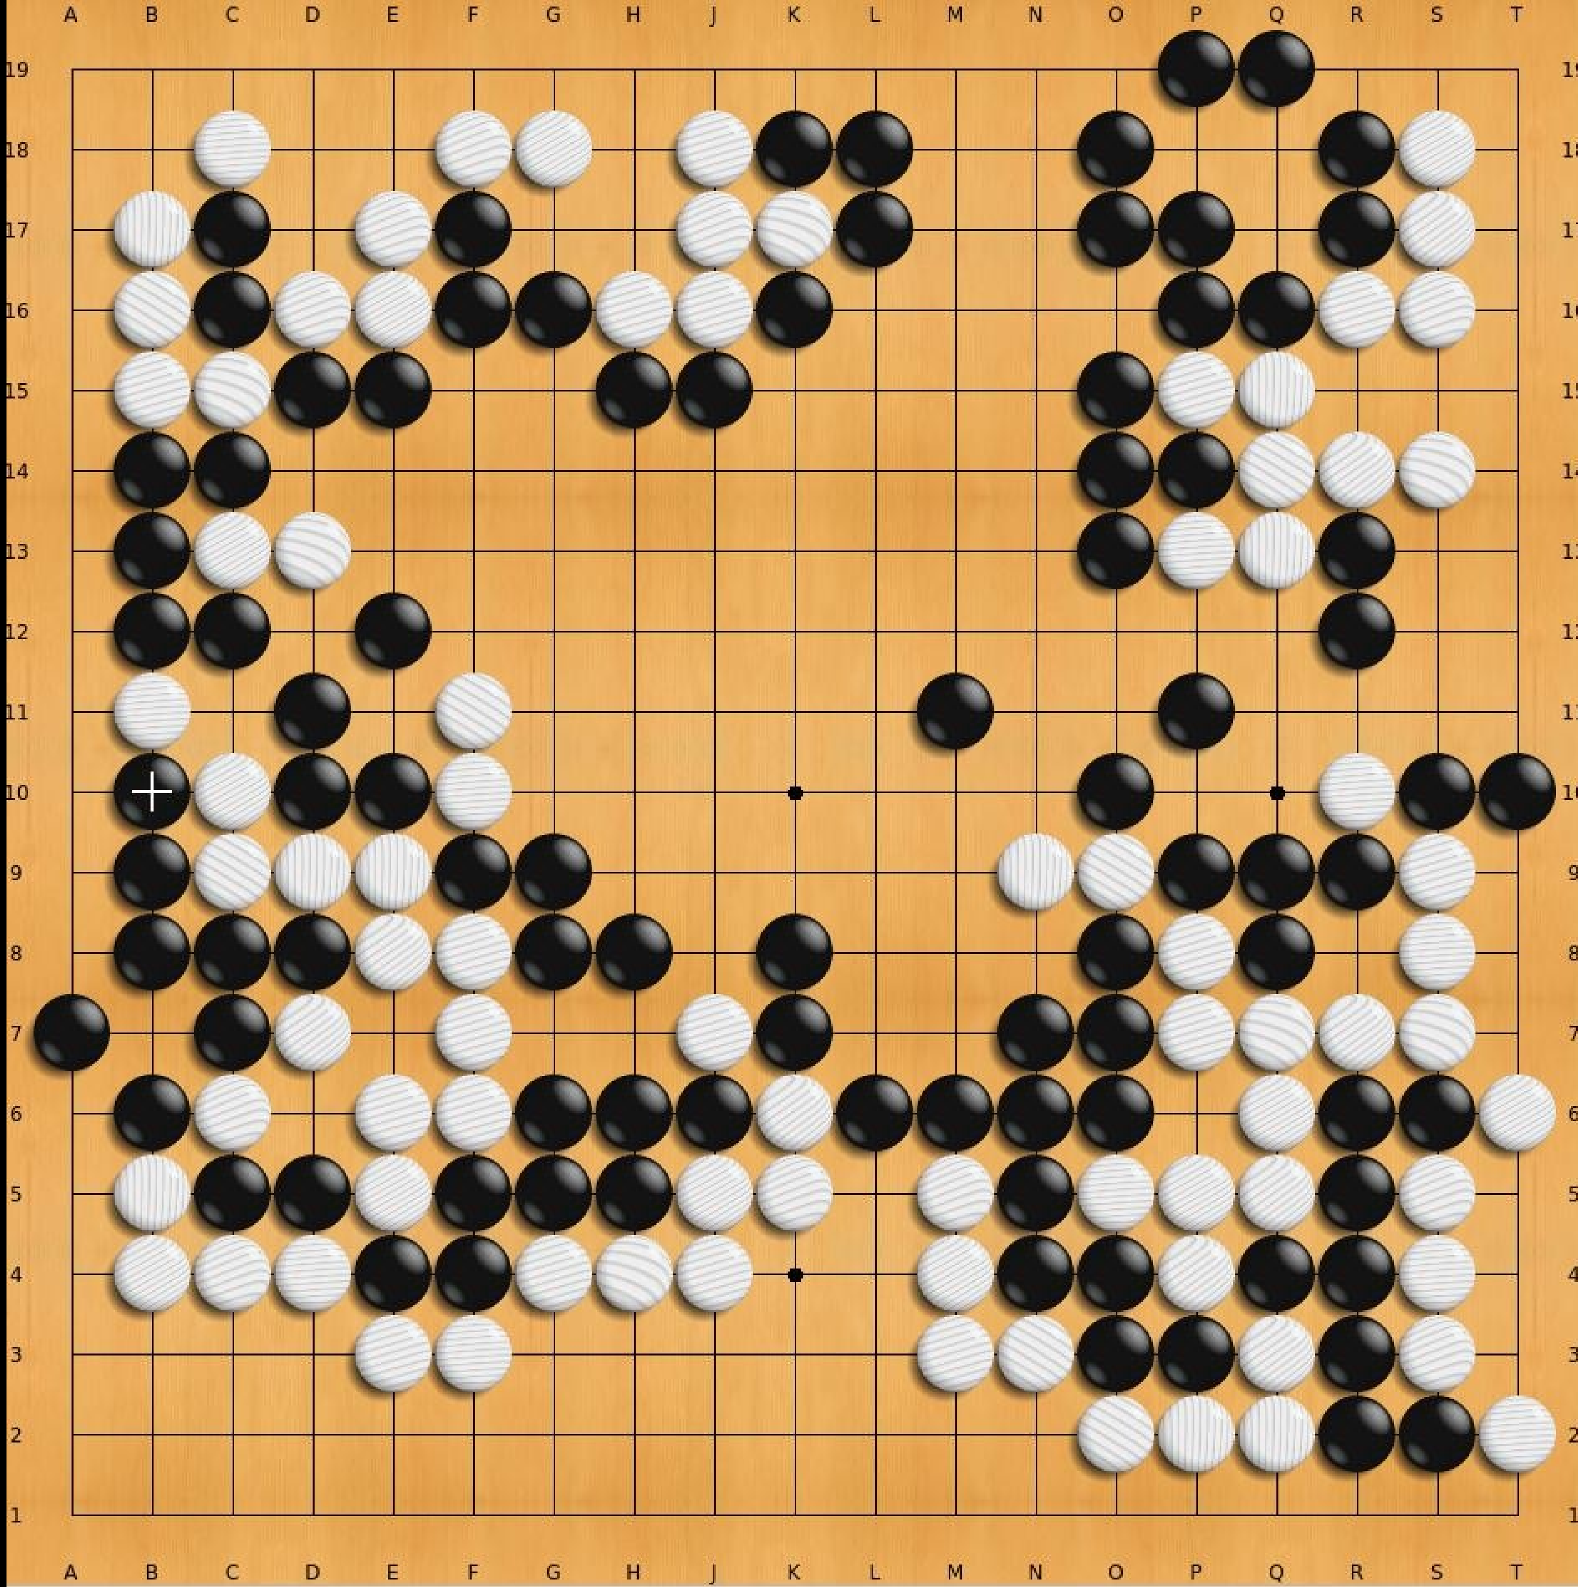
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

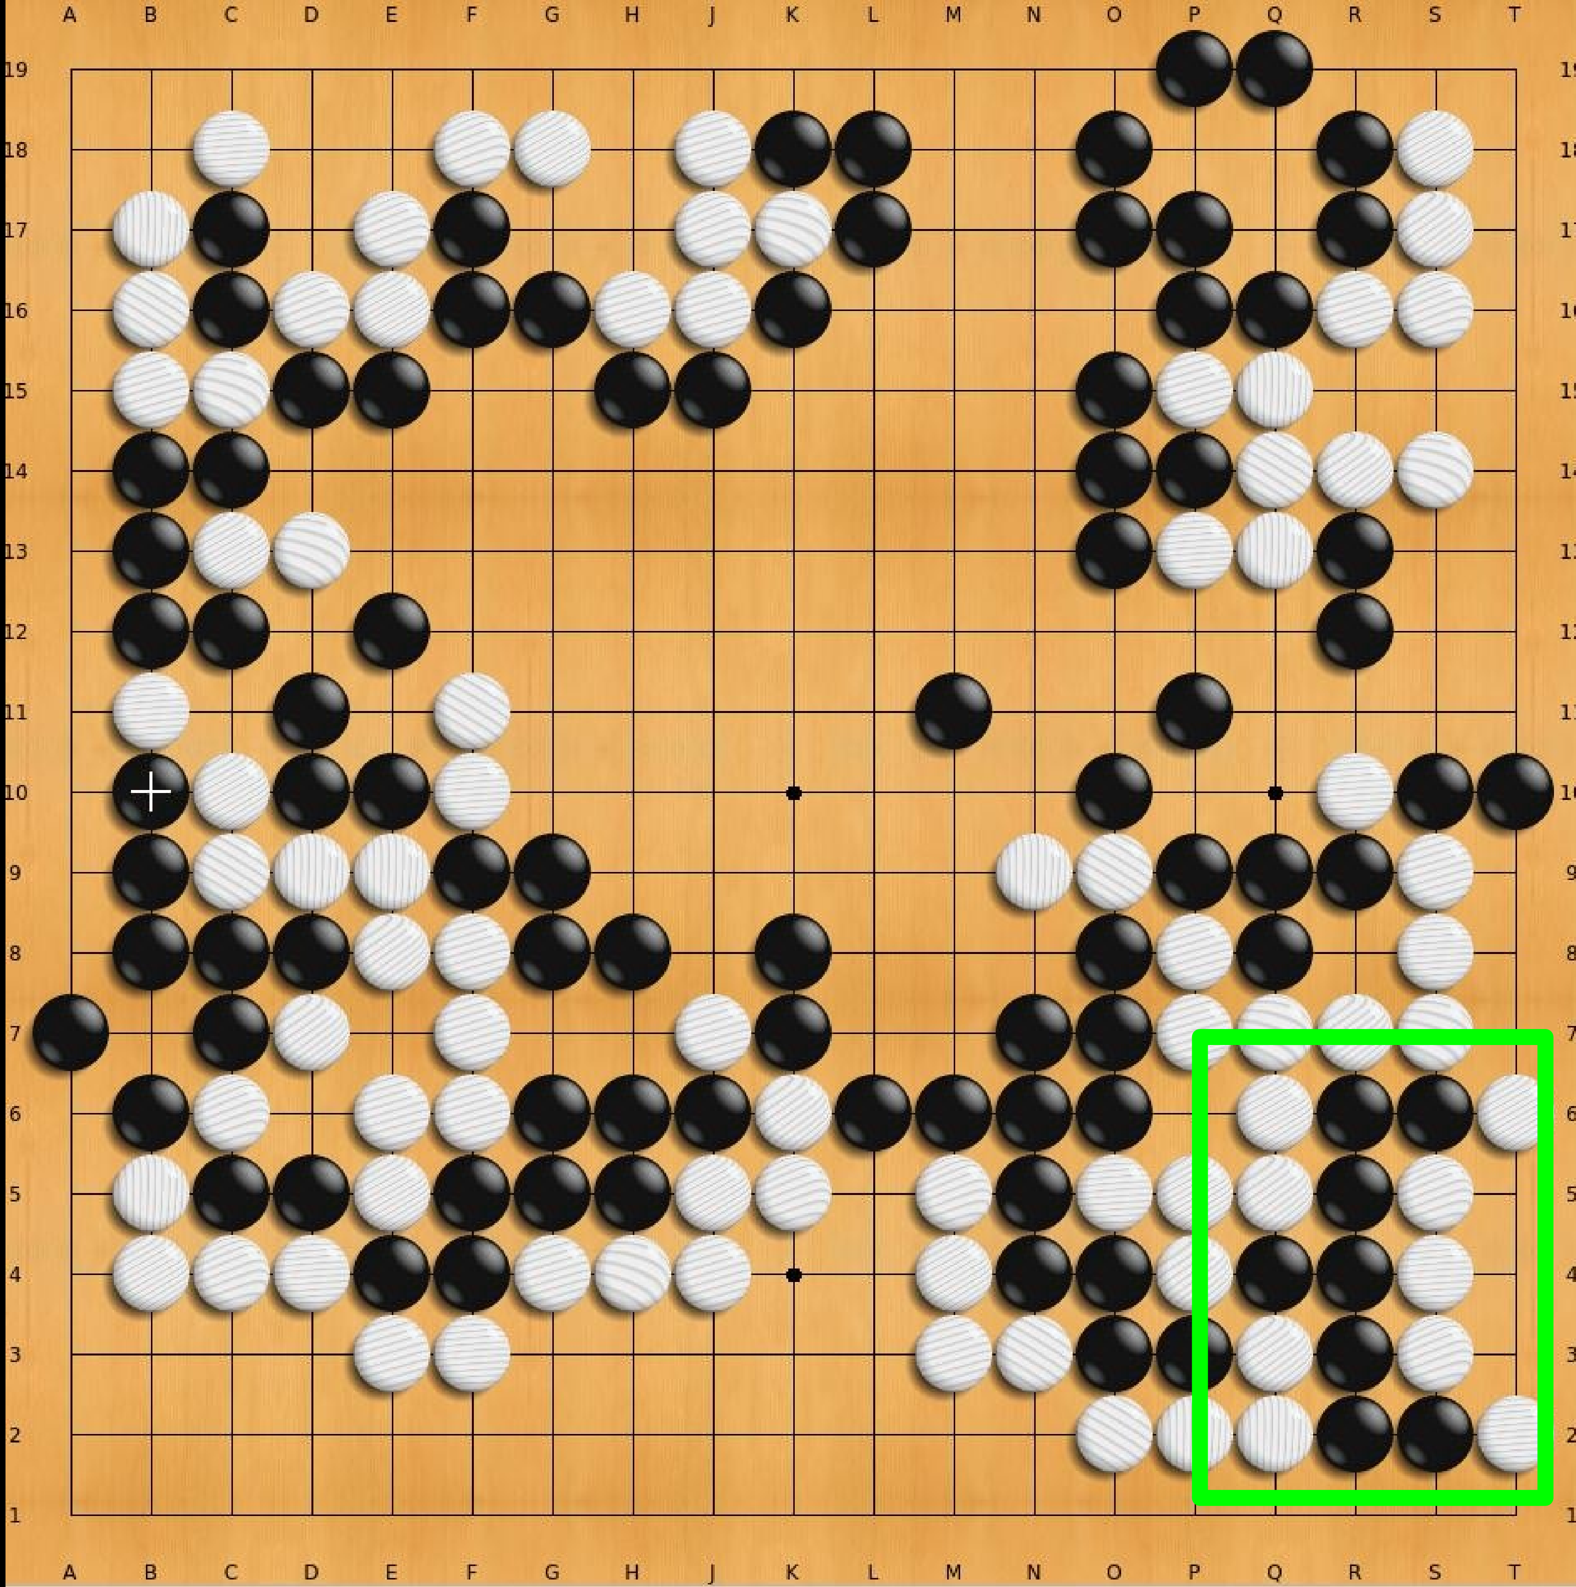
A B C D E F G H J K L M N O P Q R S T
19
18
17
16
15
14
13
12
11
10
9
8
7
6
5
4
3
2
1

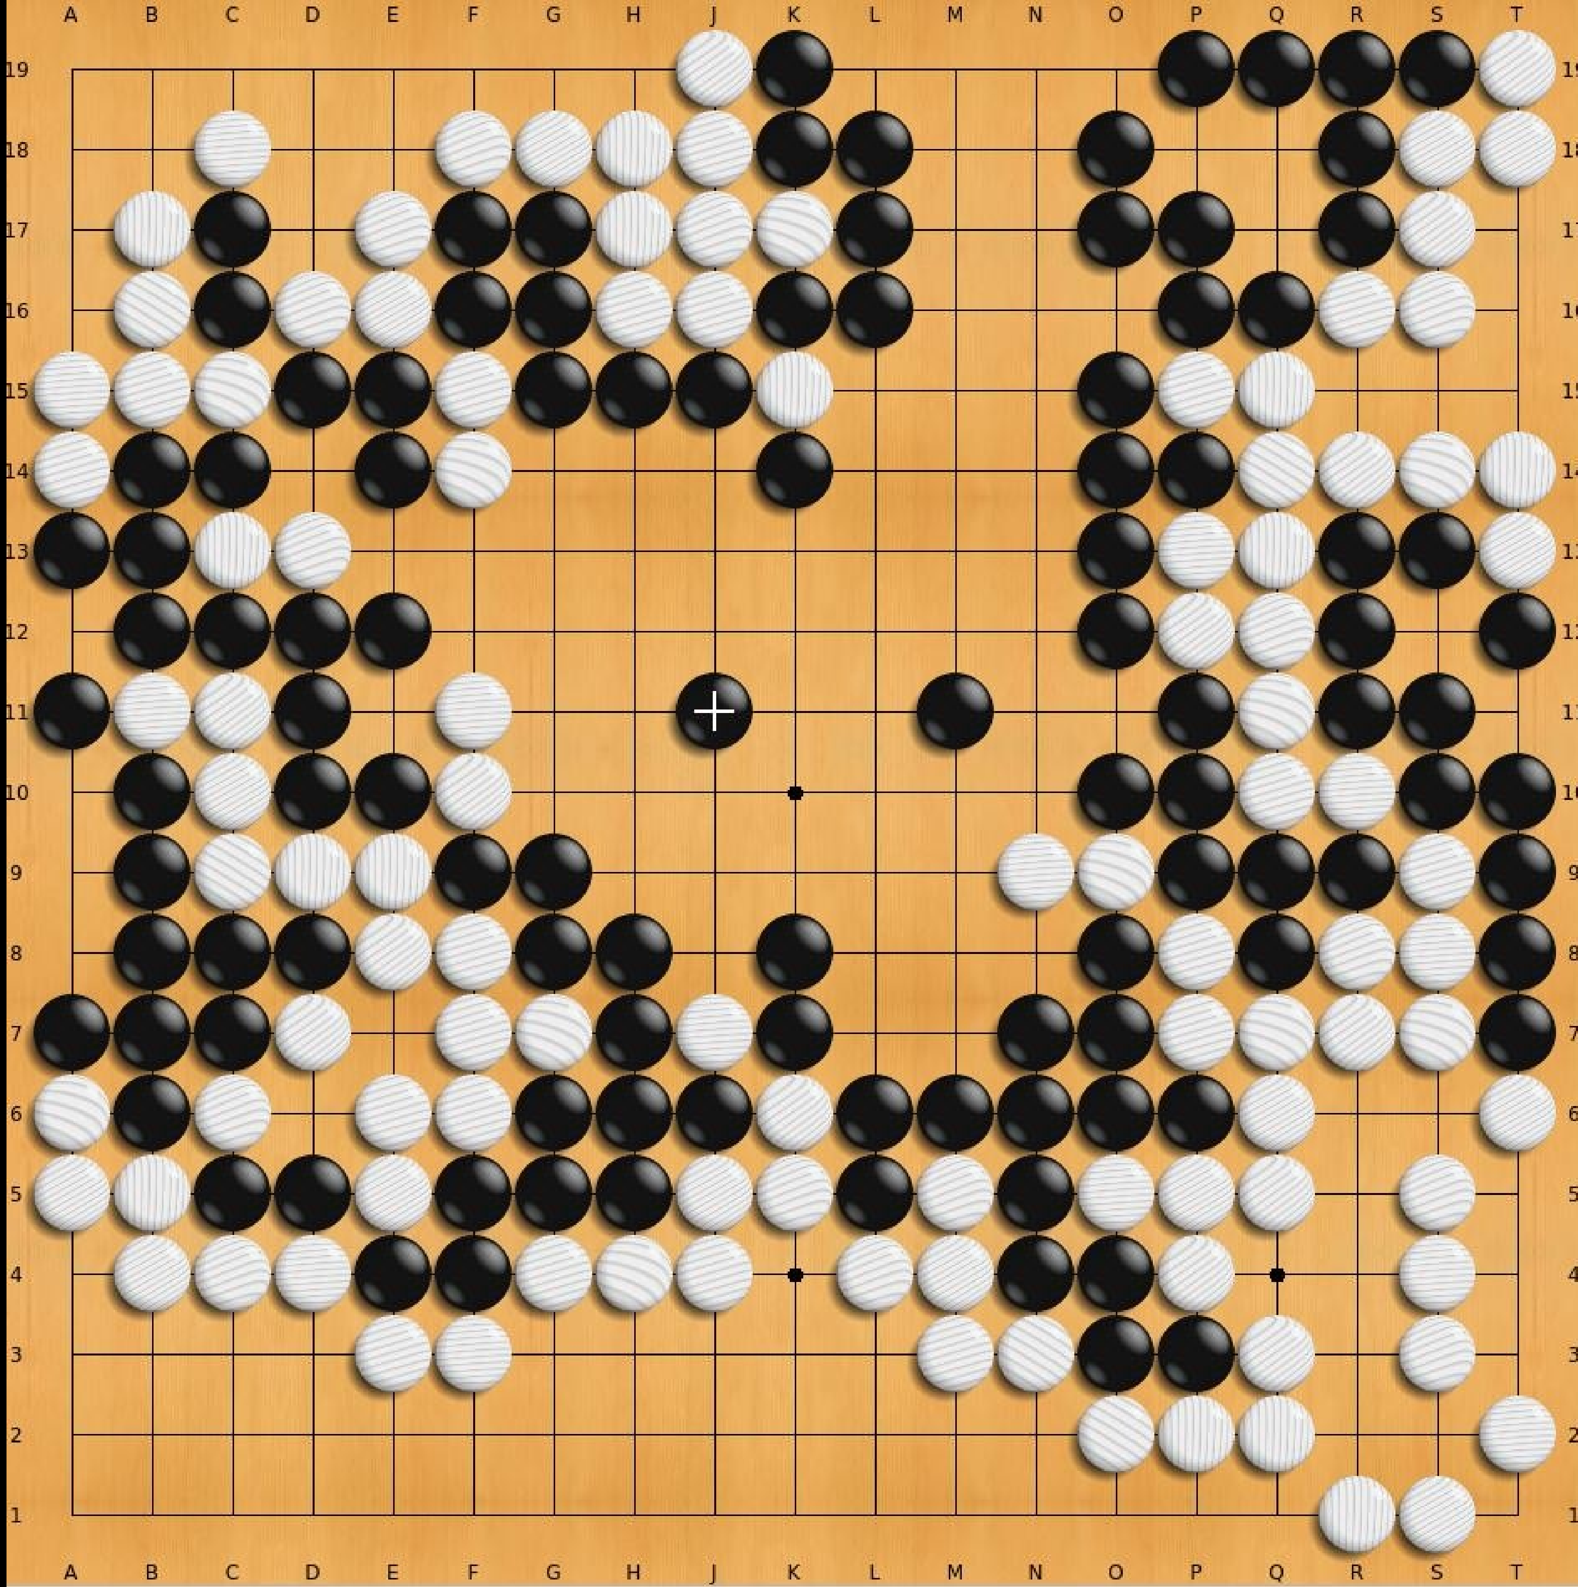
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

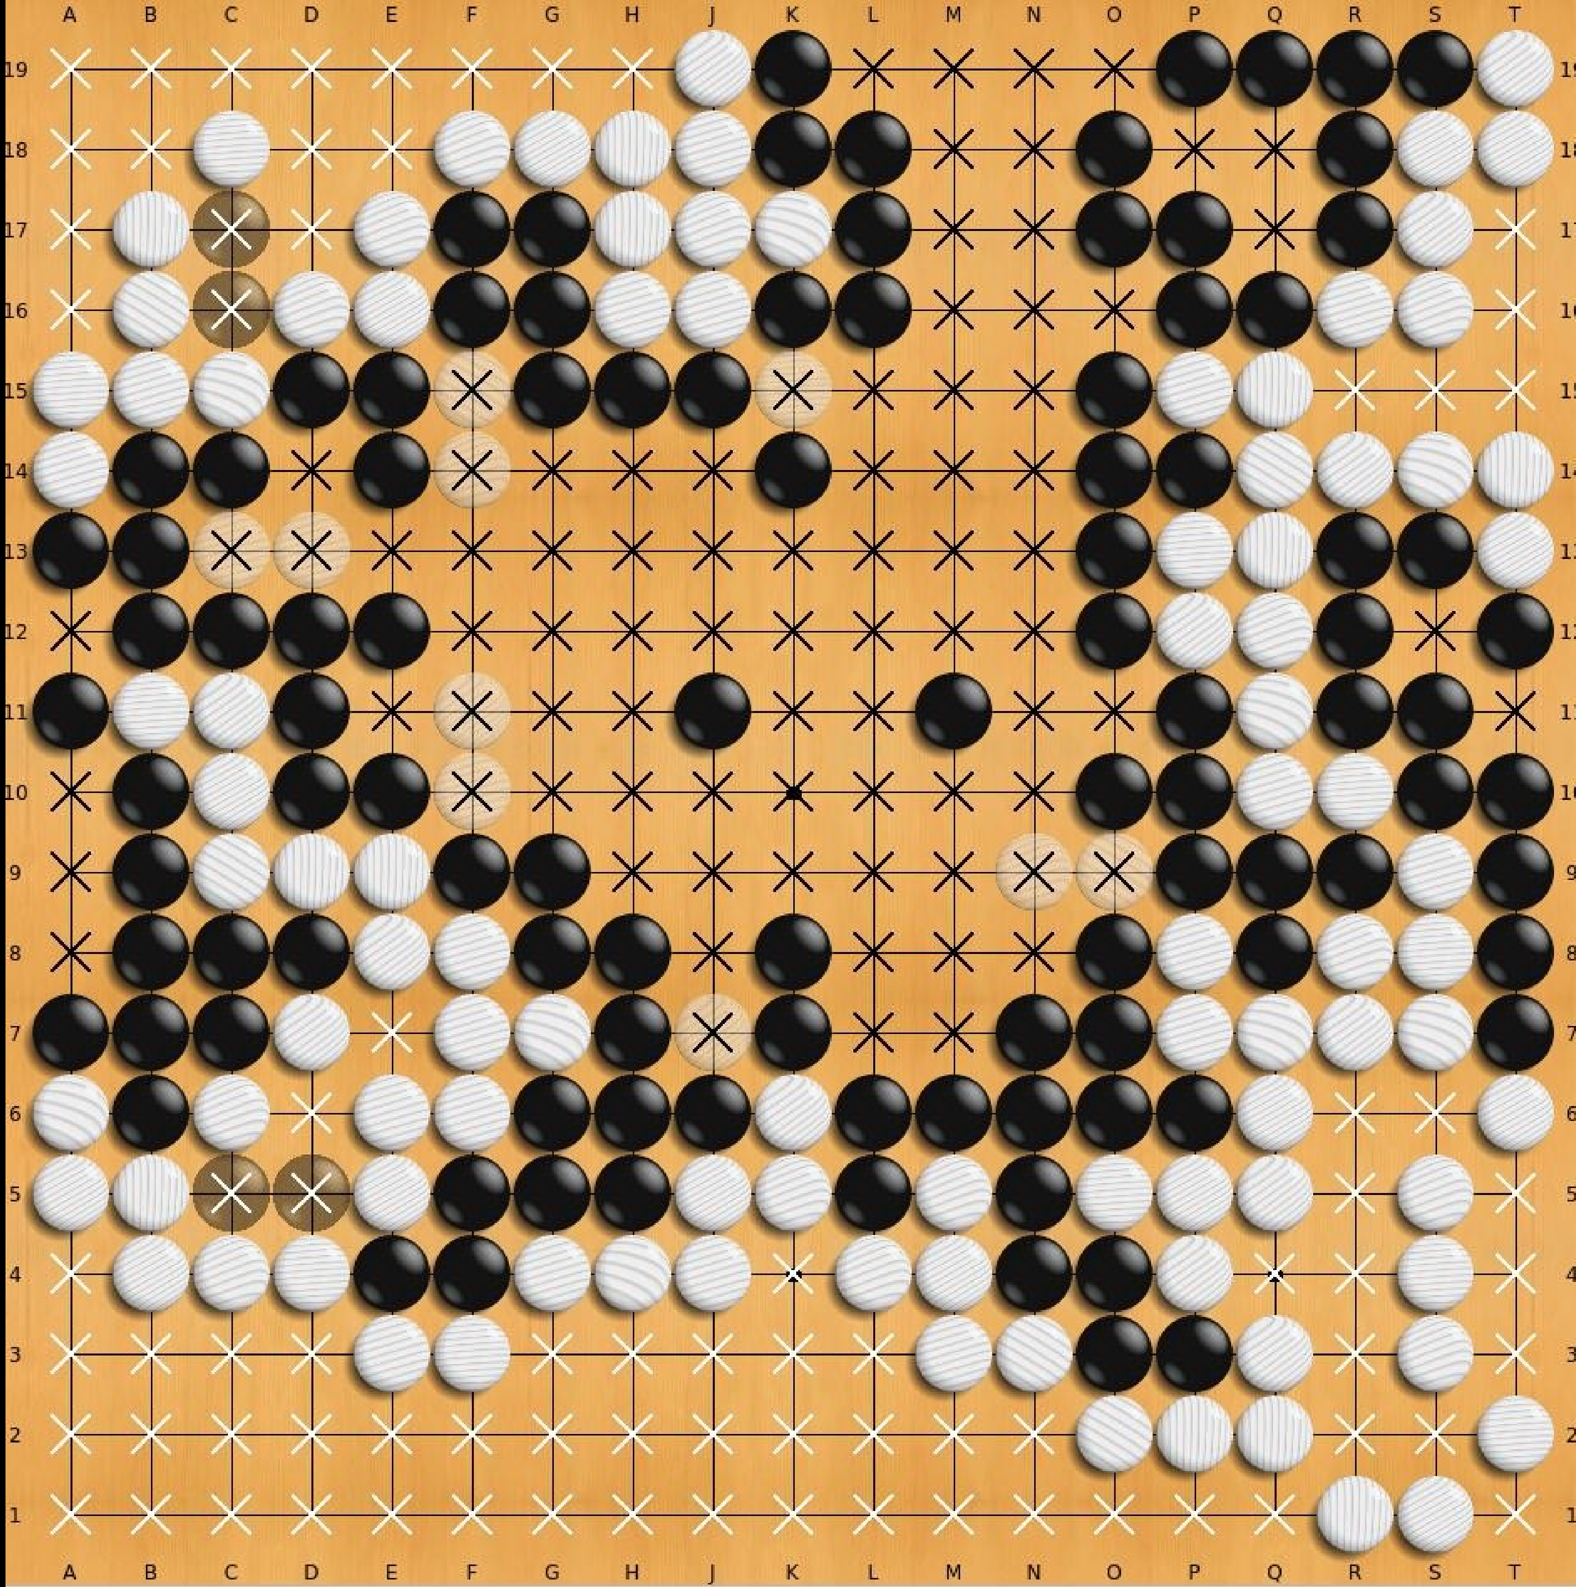
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1


Computational Challenge

Go vs. Chess

Complex vs. Complicated

„While the Baroque rules of ***chess could only have been created by humans****, the rules of* ***go*** *are so elegant, organic, and rigorously logical that* ***if intelligent life forms exist*** *elsewhere in the universe, they almost* ***certainly play go****.“*

Edward Lasker (chess grandmaster)

Larger board

19x19 vs. 8x8

Almost every move is legal

Average branching factor:

250 vs 35

State Space Complexity:

10^{171} vs 10^{47}


10^{80}

Global impact of moves

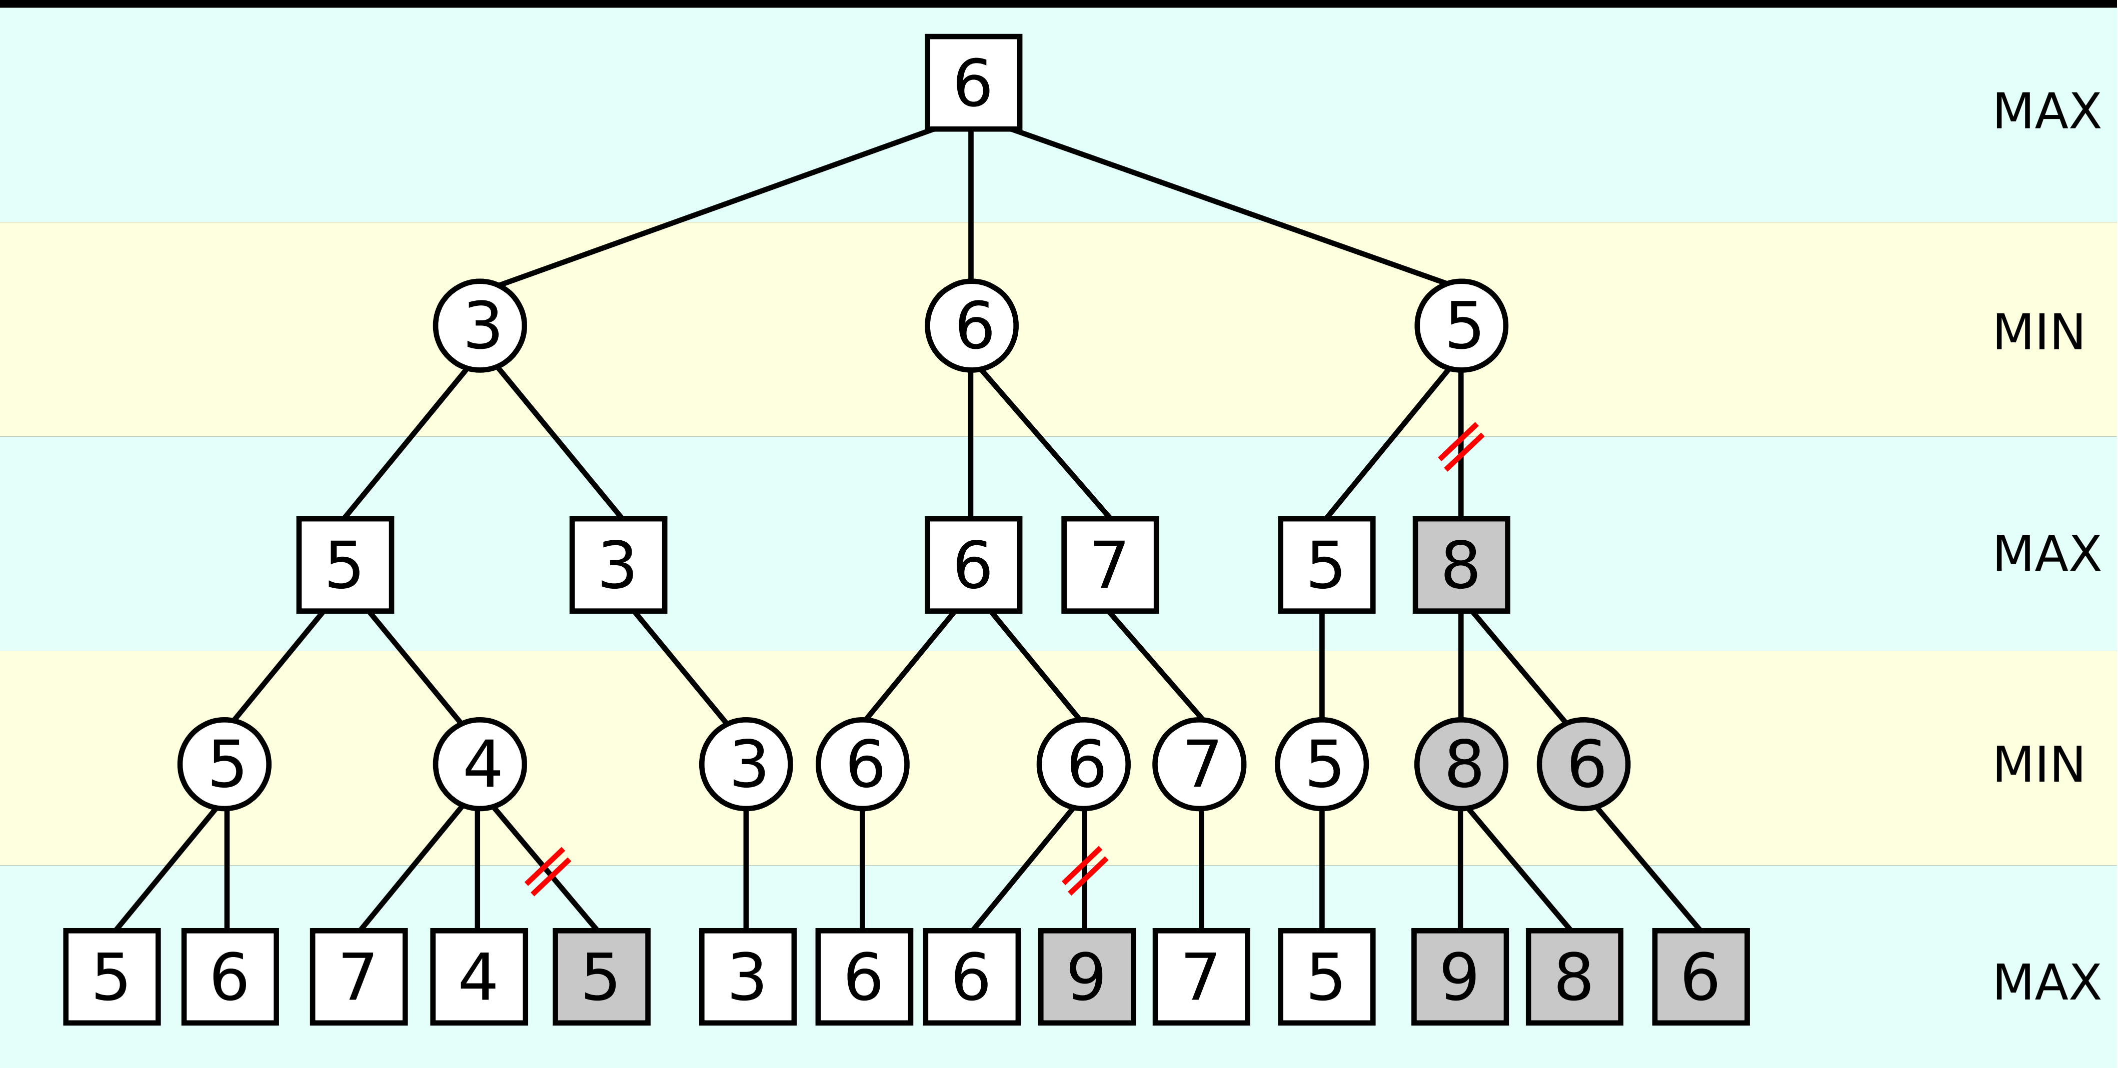
MAX
MIN
MAX
MIN
MAX

Evaluation function

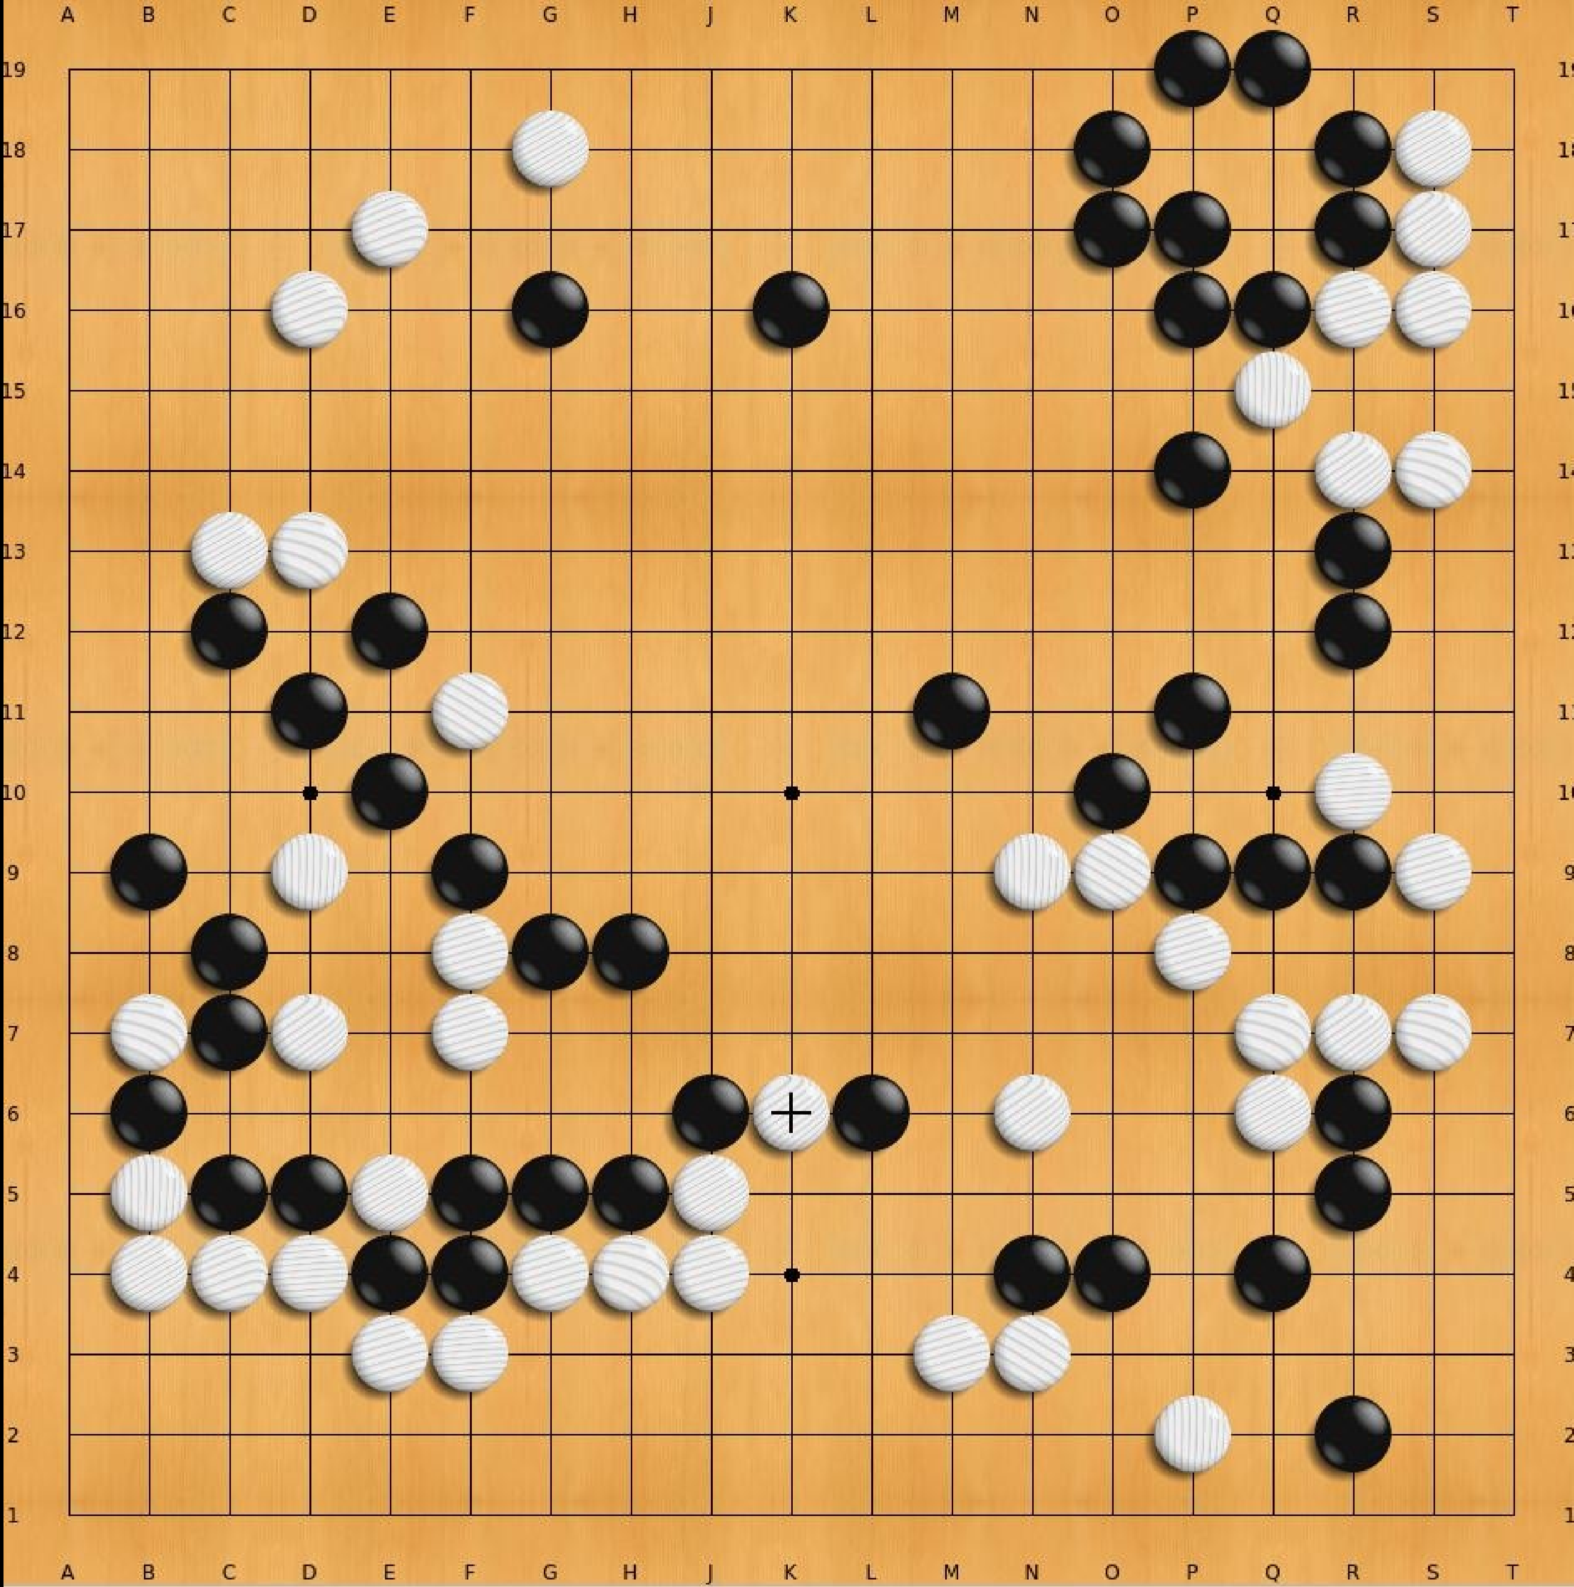
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1


Monte Carlo Method

What is Pi?

How do you determine Pi?

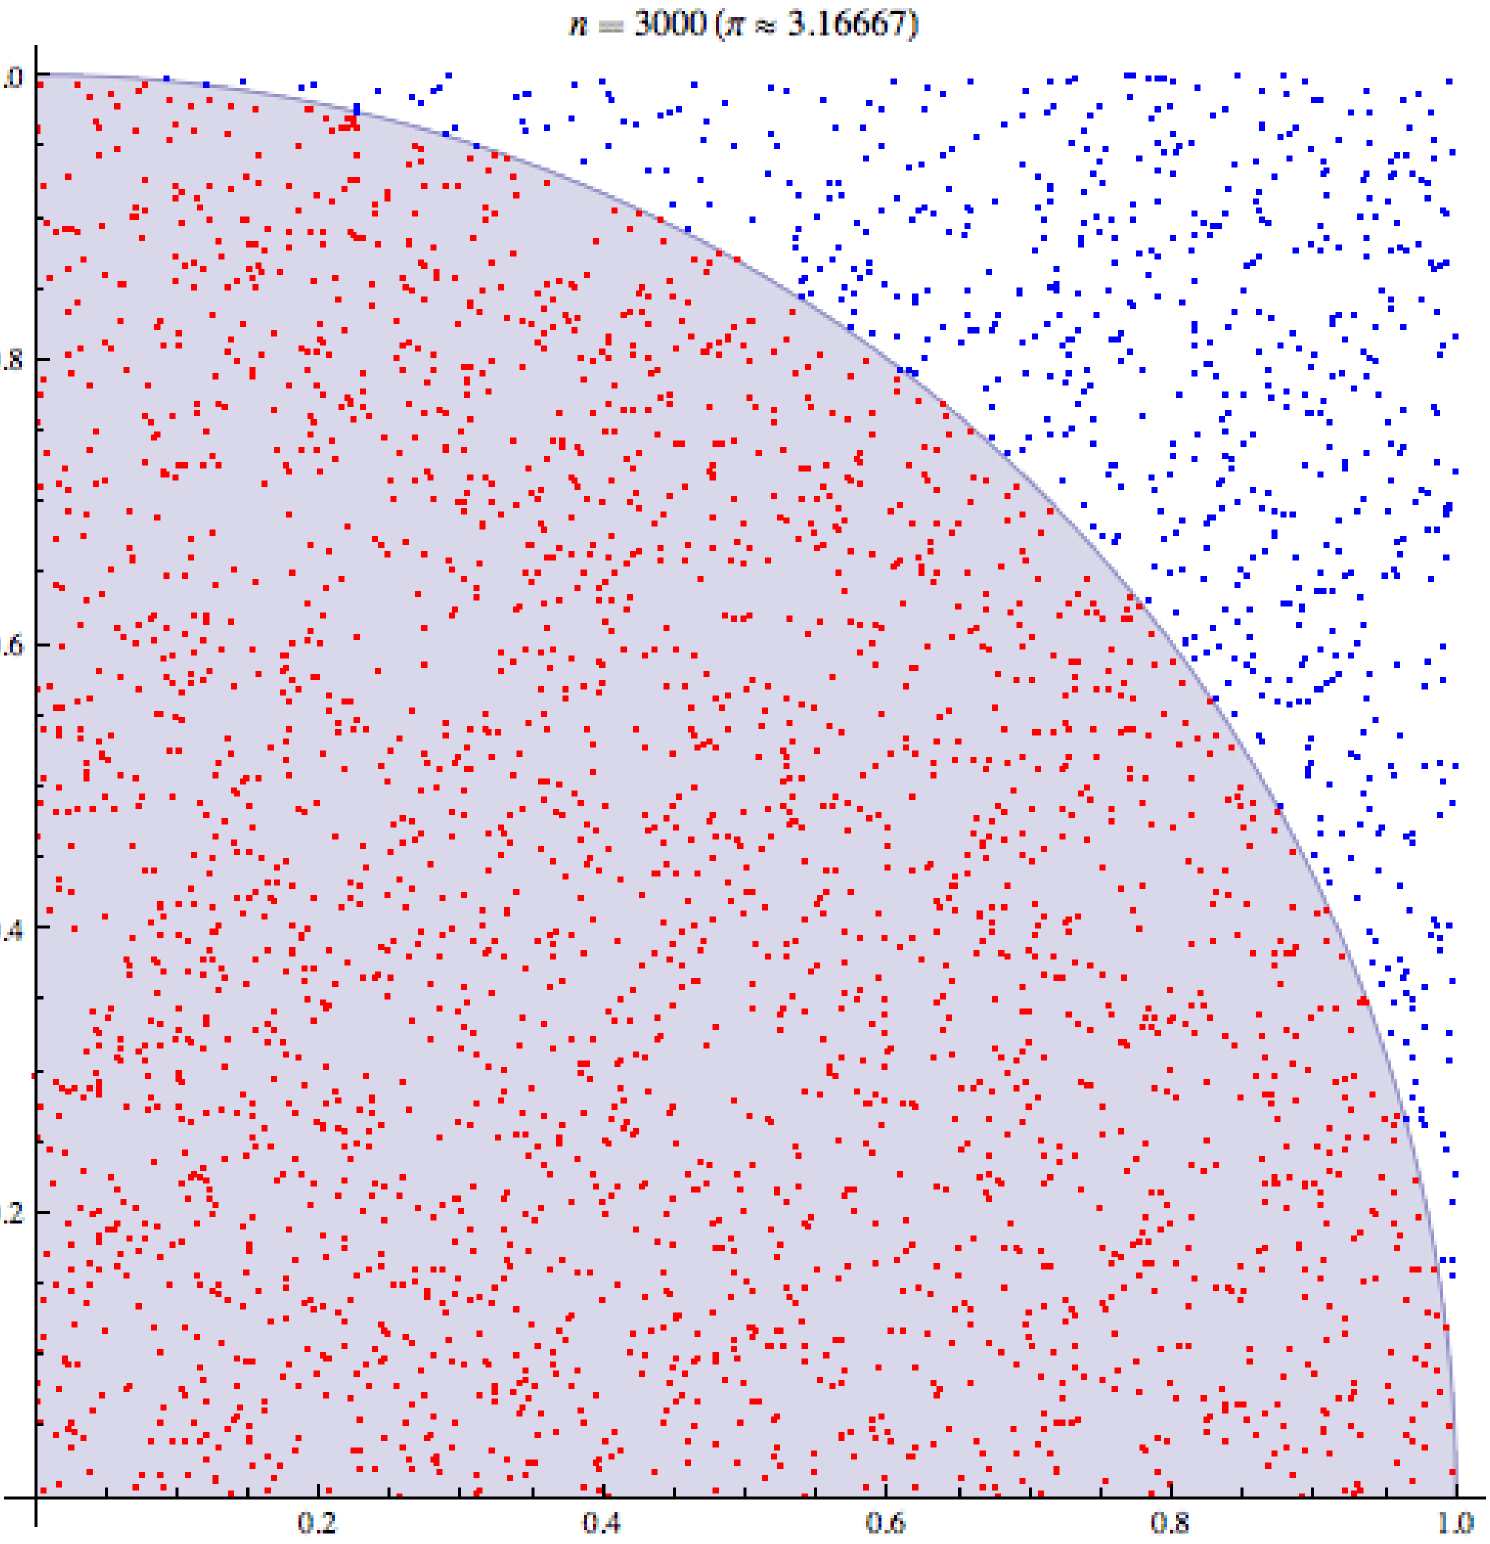
n = 3000 (π ≈ 3.16667)
1.0
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1.0

2006

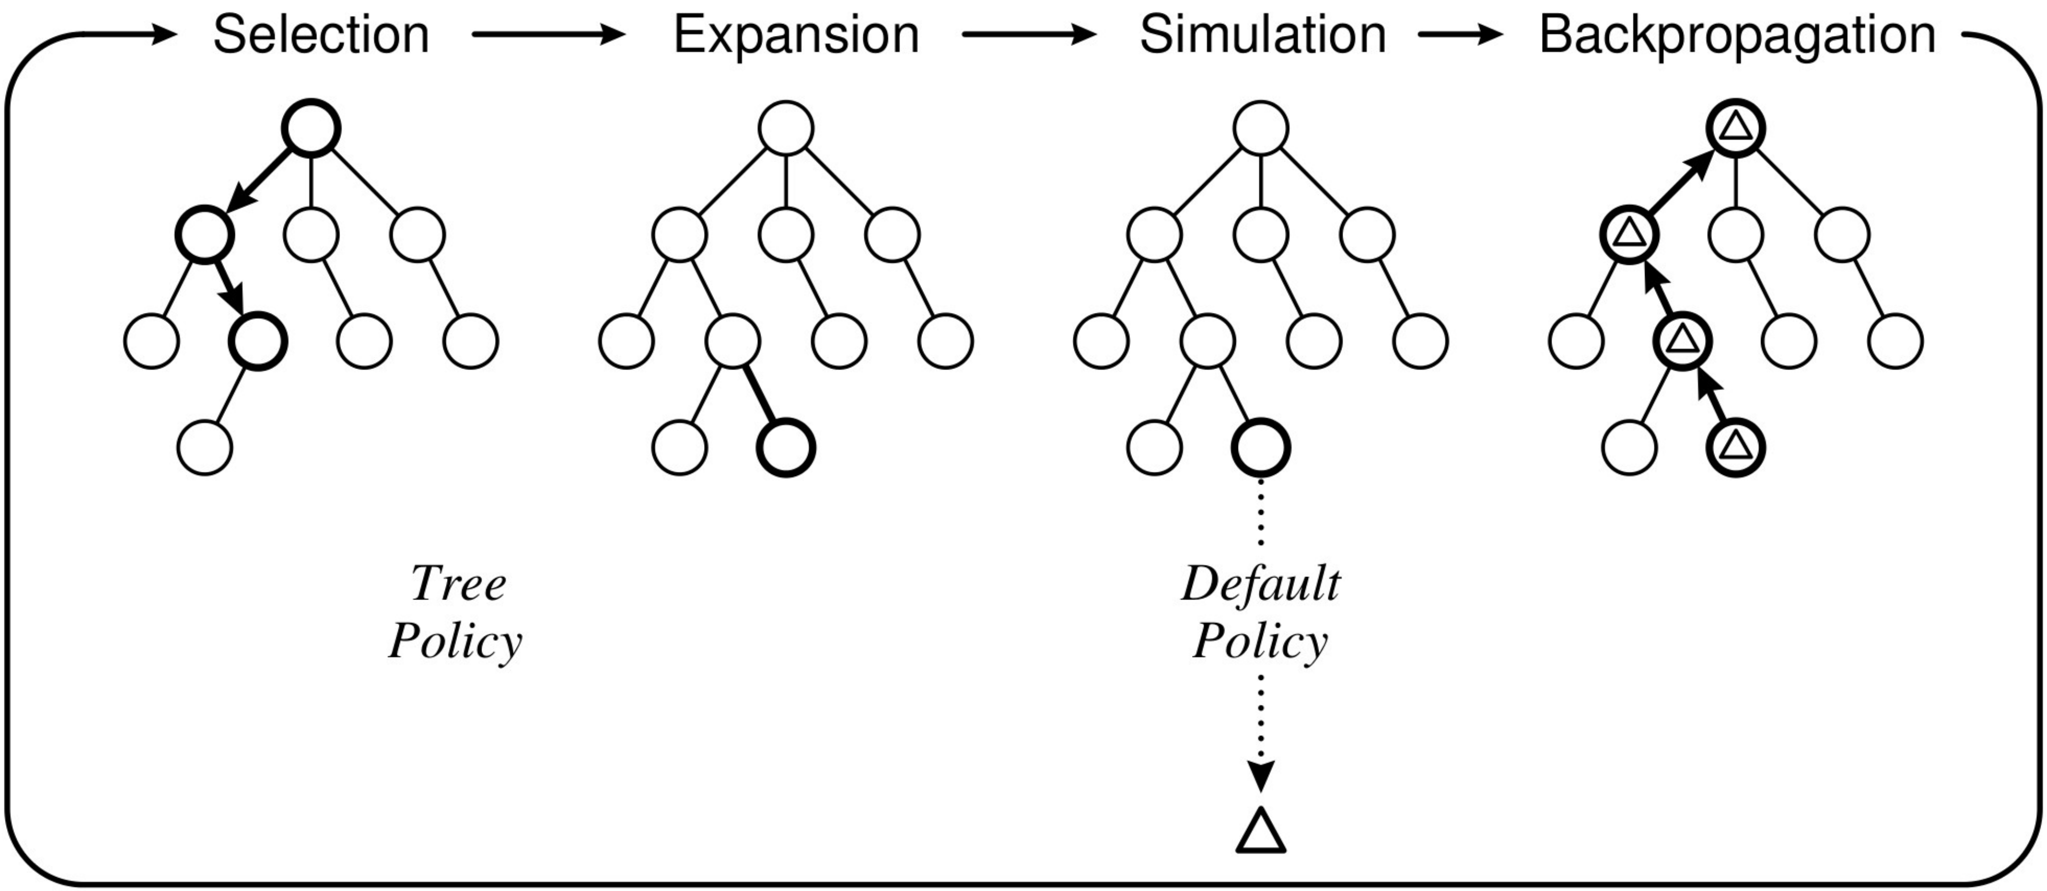

Browne, Cb, and Edward Powley. 2012. A survey of monte carlo tree search methods. Intelligence and AI 4, no. 1: 1-49

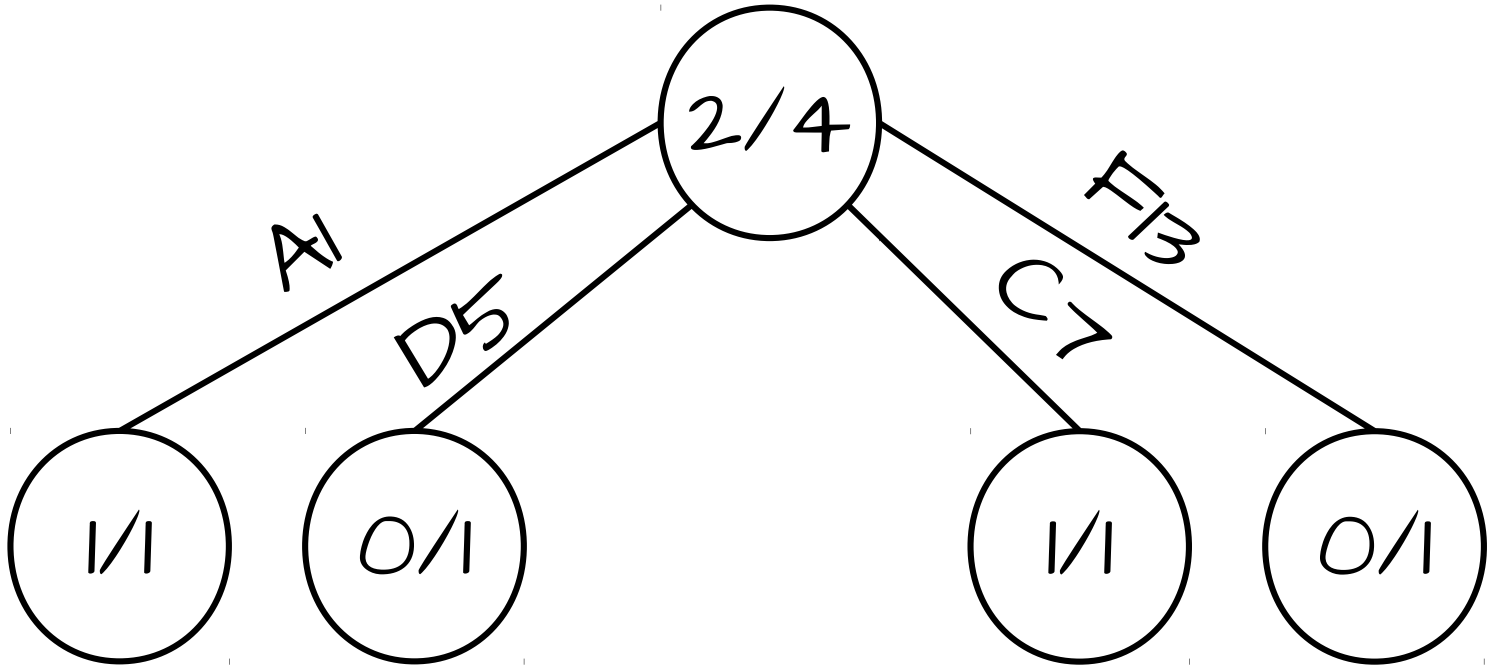
2/4
A1
D5
C7
F13
1/1
0/1
1/1
0/1

Selection

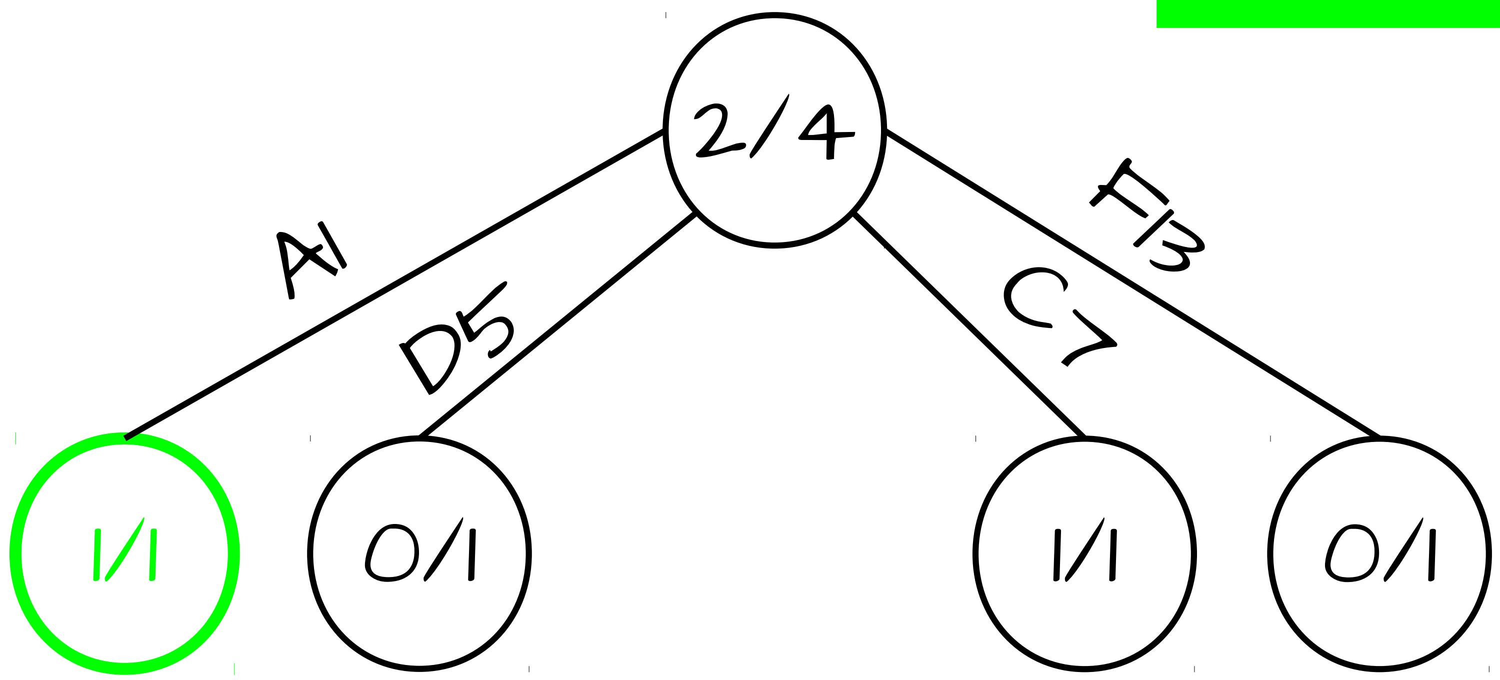

Expansion

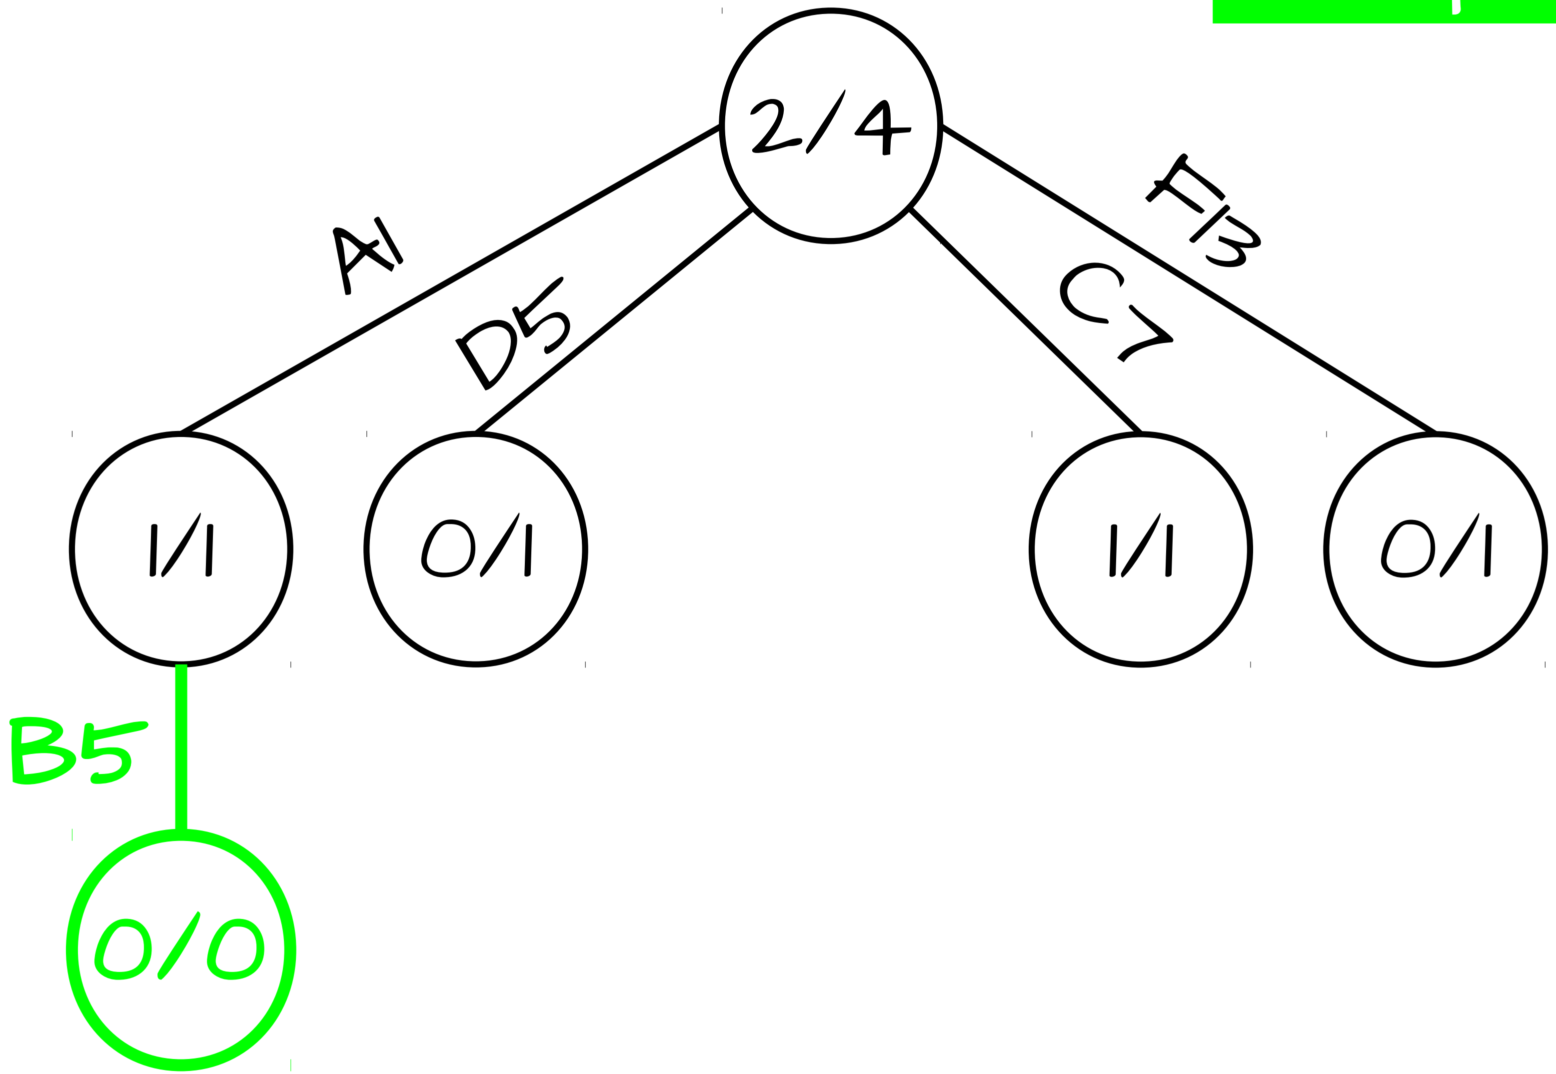

Simulation

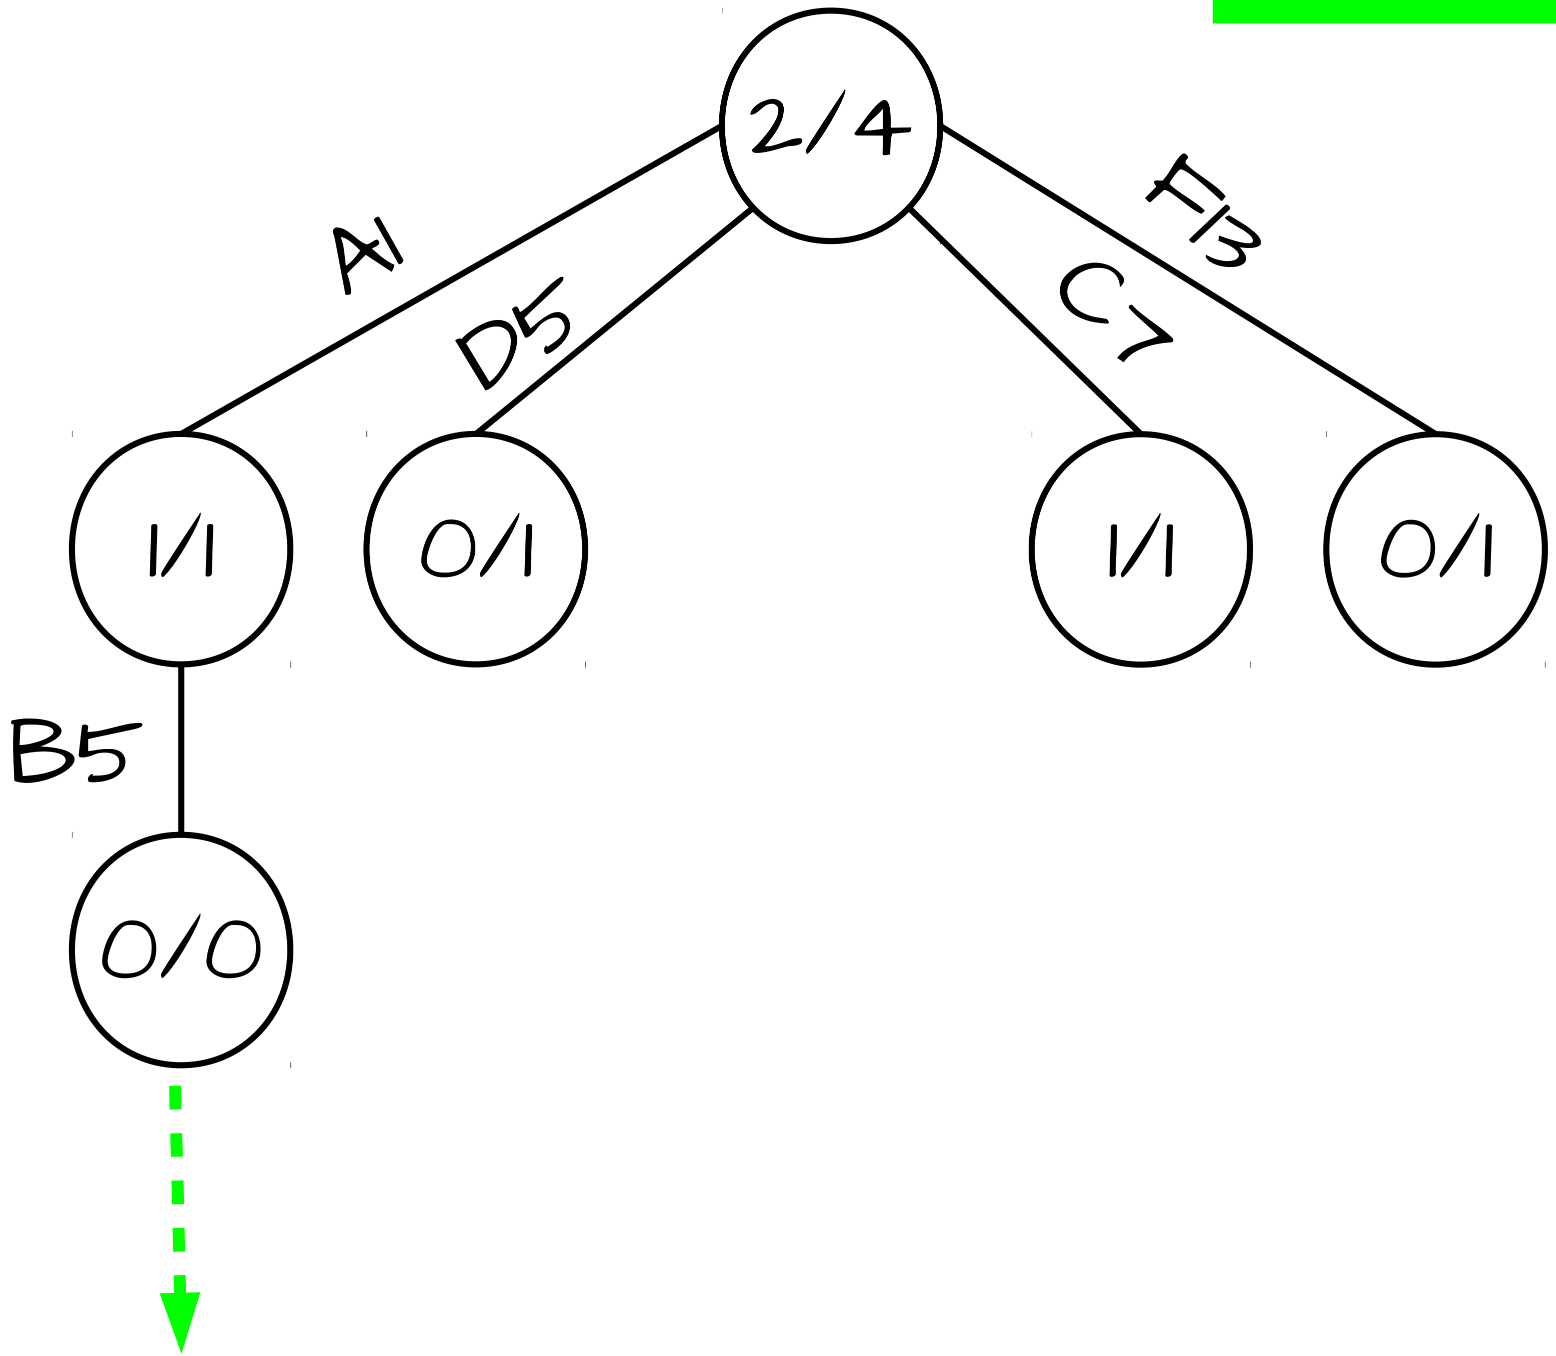

Random

Backpropagation

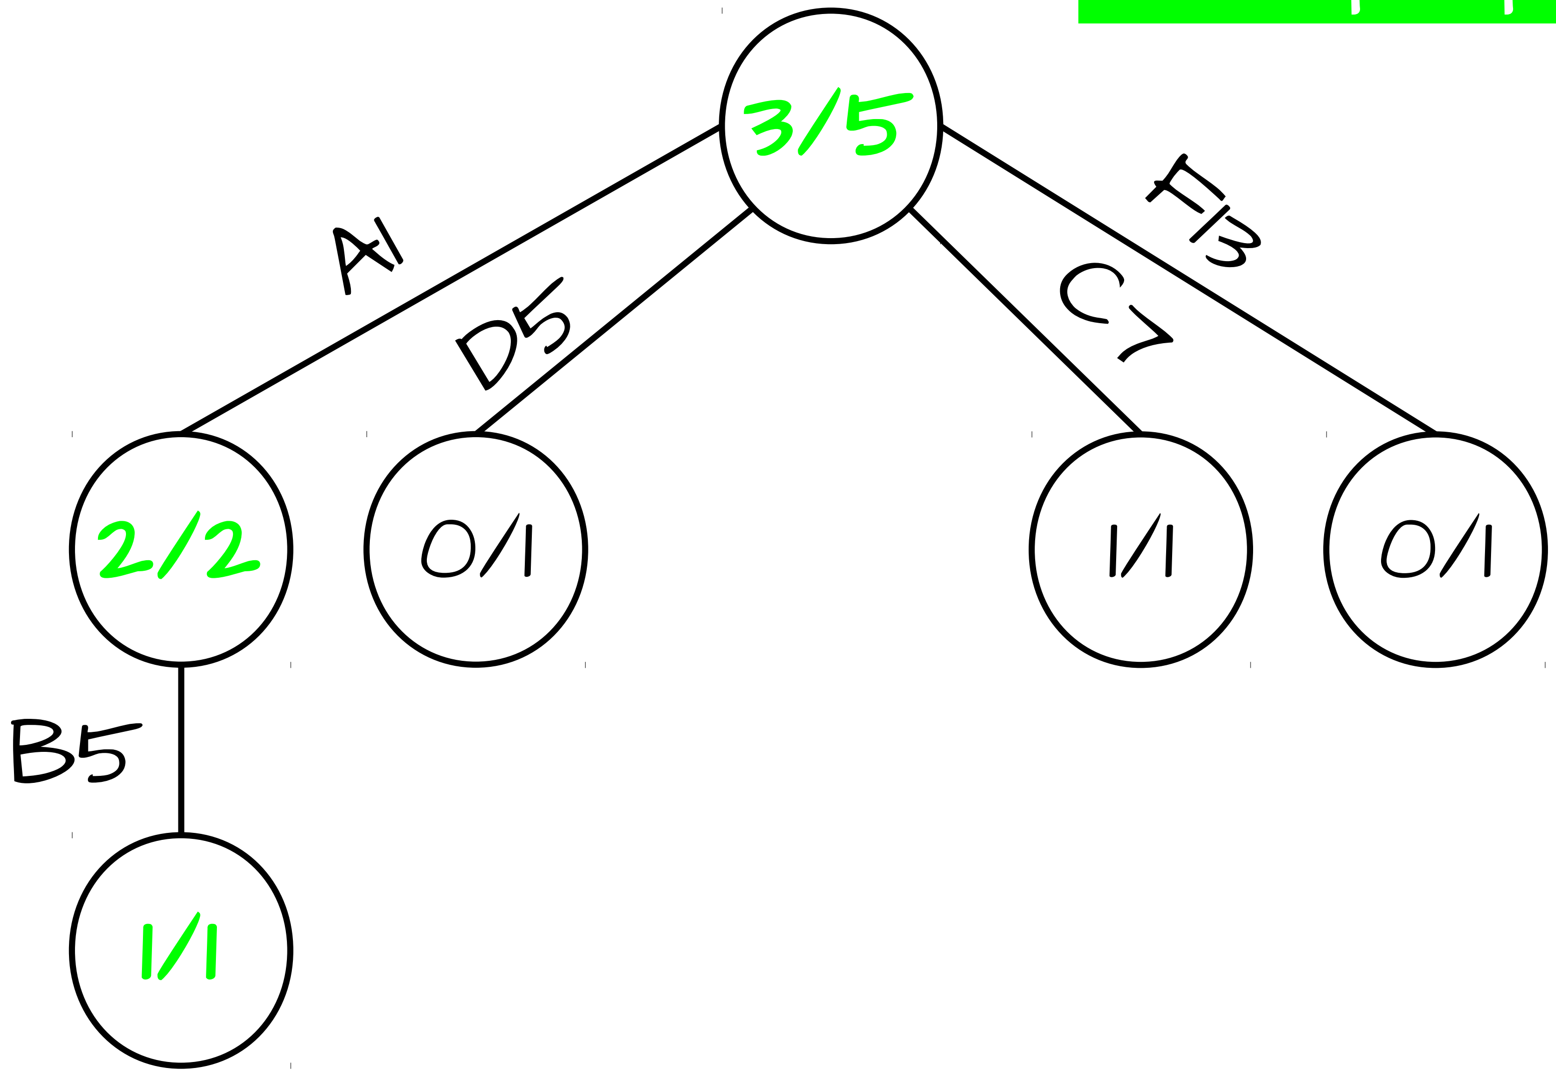

Perspective

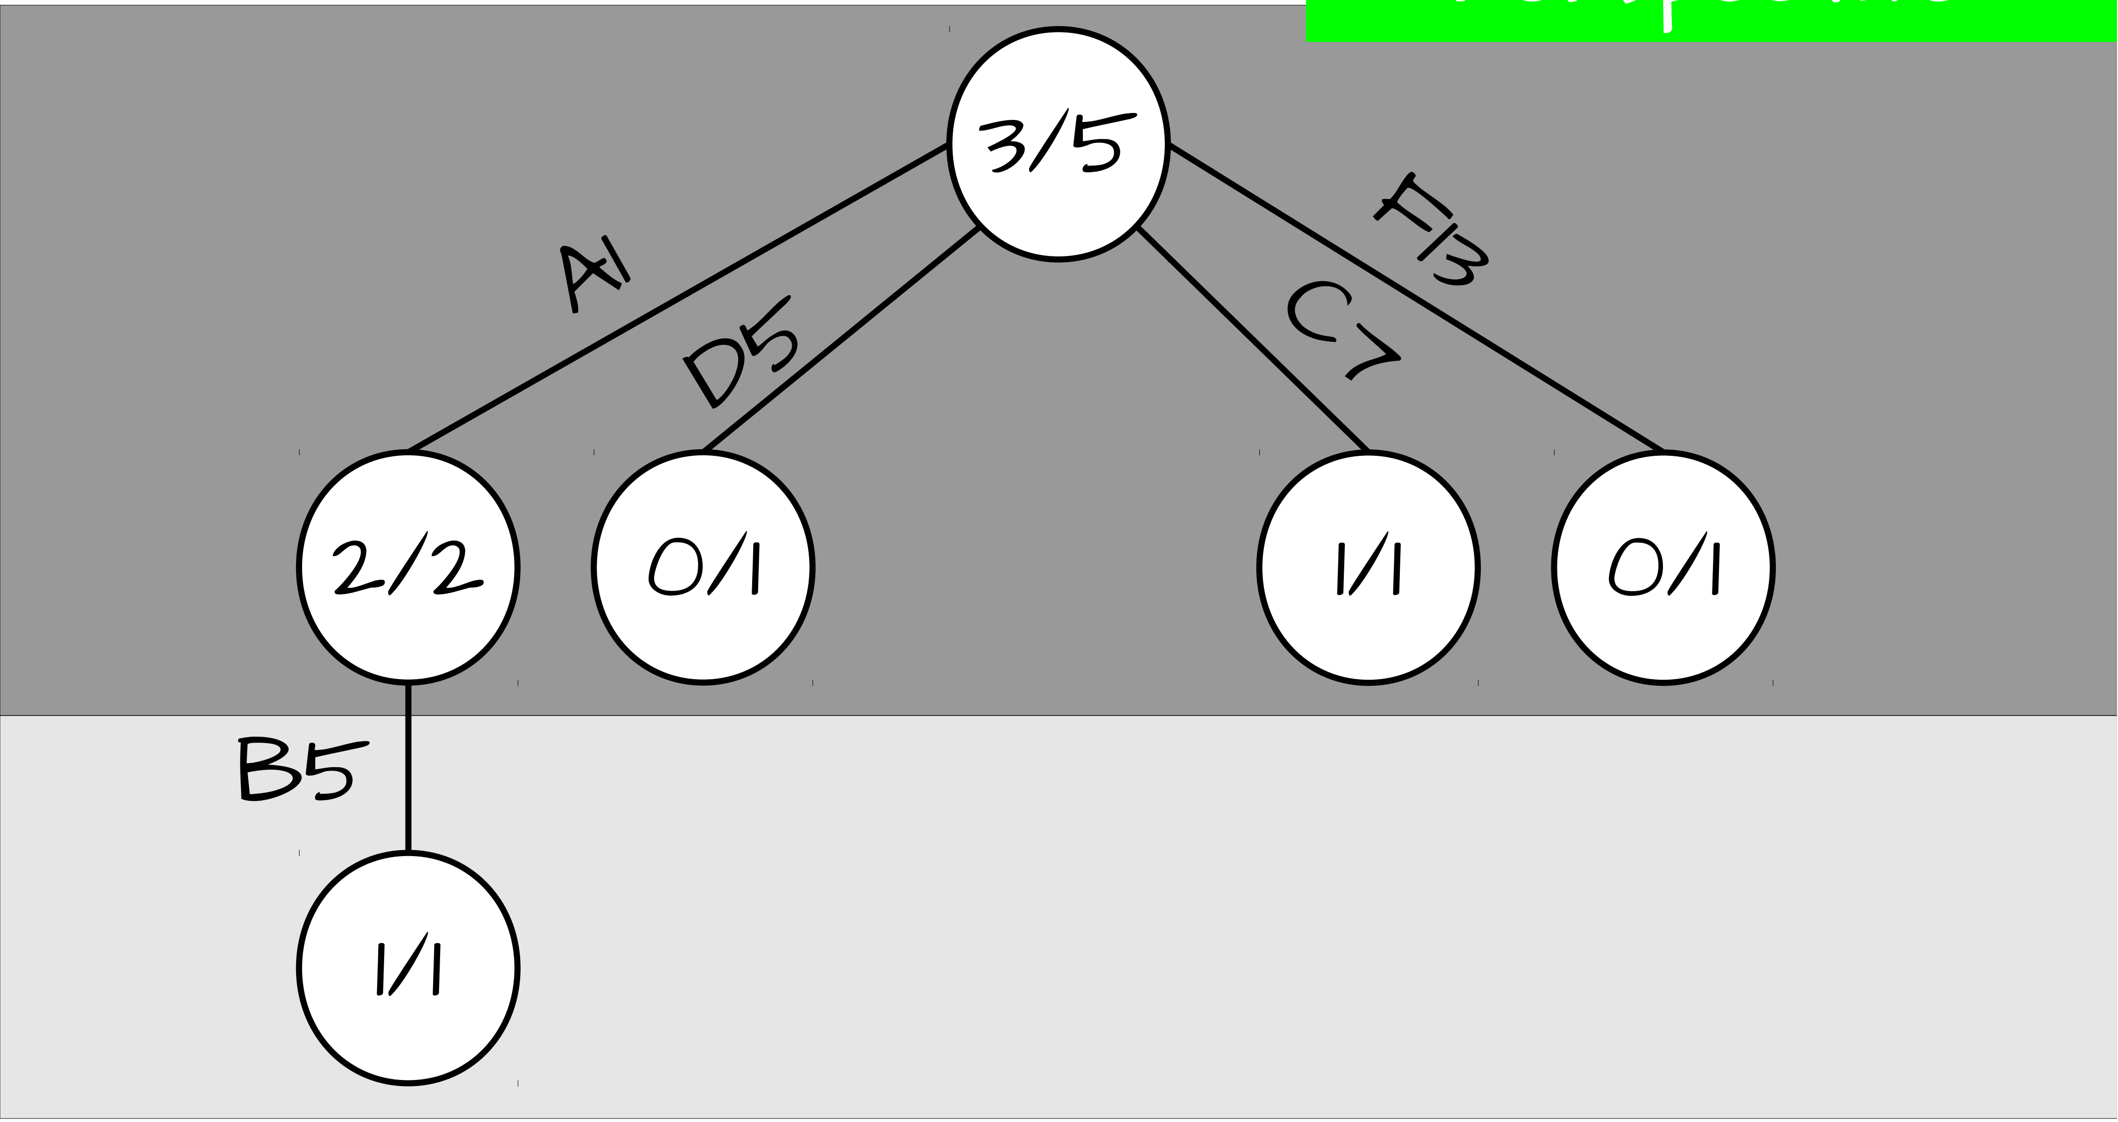

Perspective

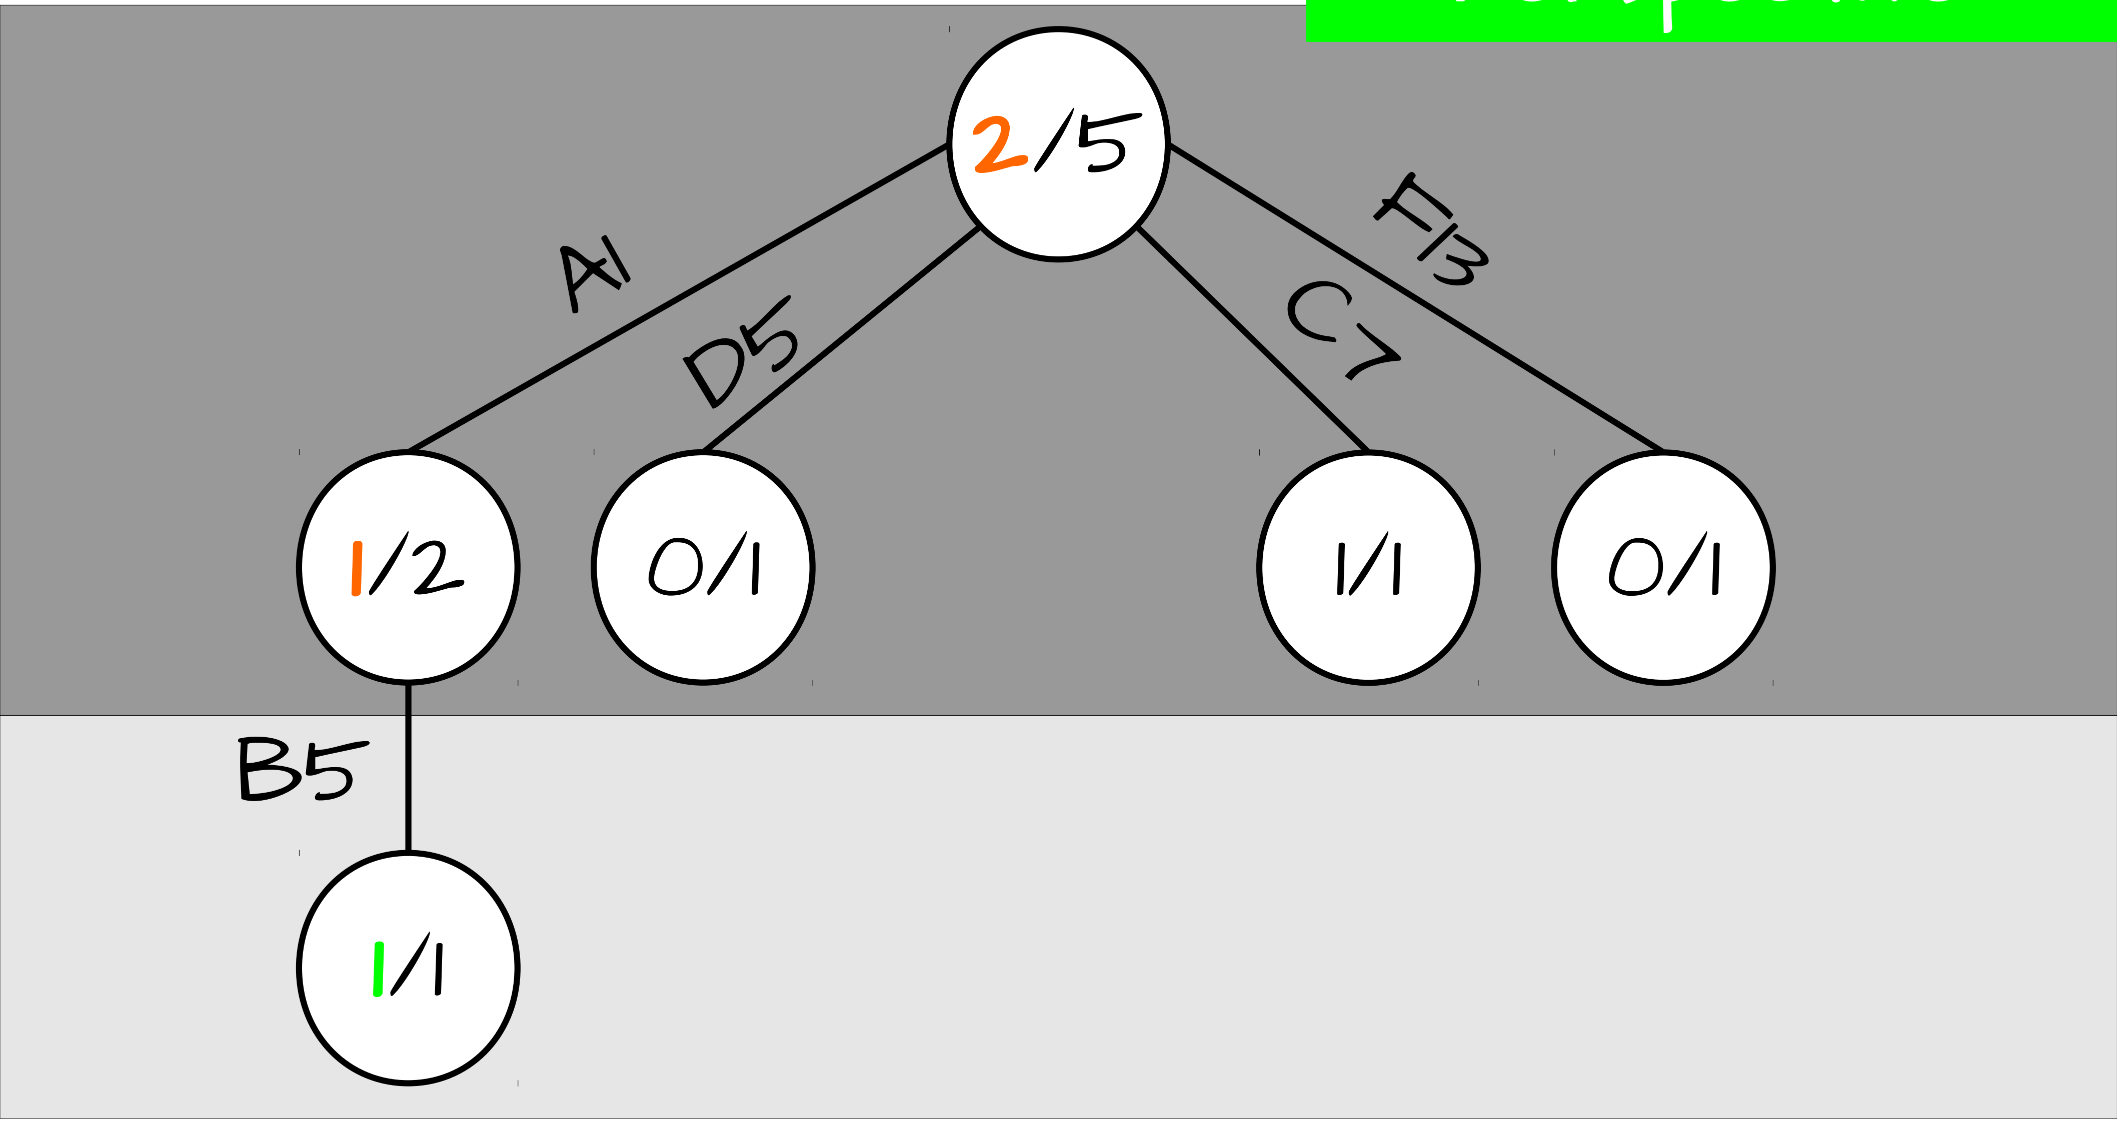


Multi Armed Bandit
Lil Red
SPARTACUS
GLADIATOR OF ROME
WIN FREE SPINS WITH NUDGING WILDS!
FORBIDDEN DRAGONS
WMS
CQ

Exploitation vs Exploration

$$\frac{wins}{visits} + explorationFactor\sqrt{\frac{\ln\left(totalVisits\right)}{visits}}$$

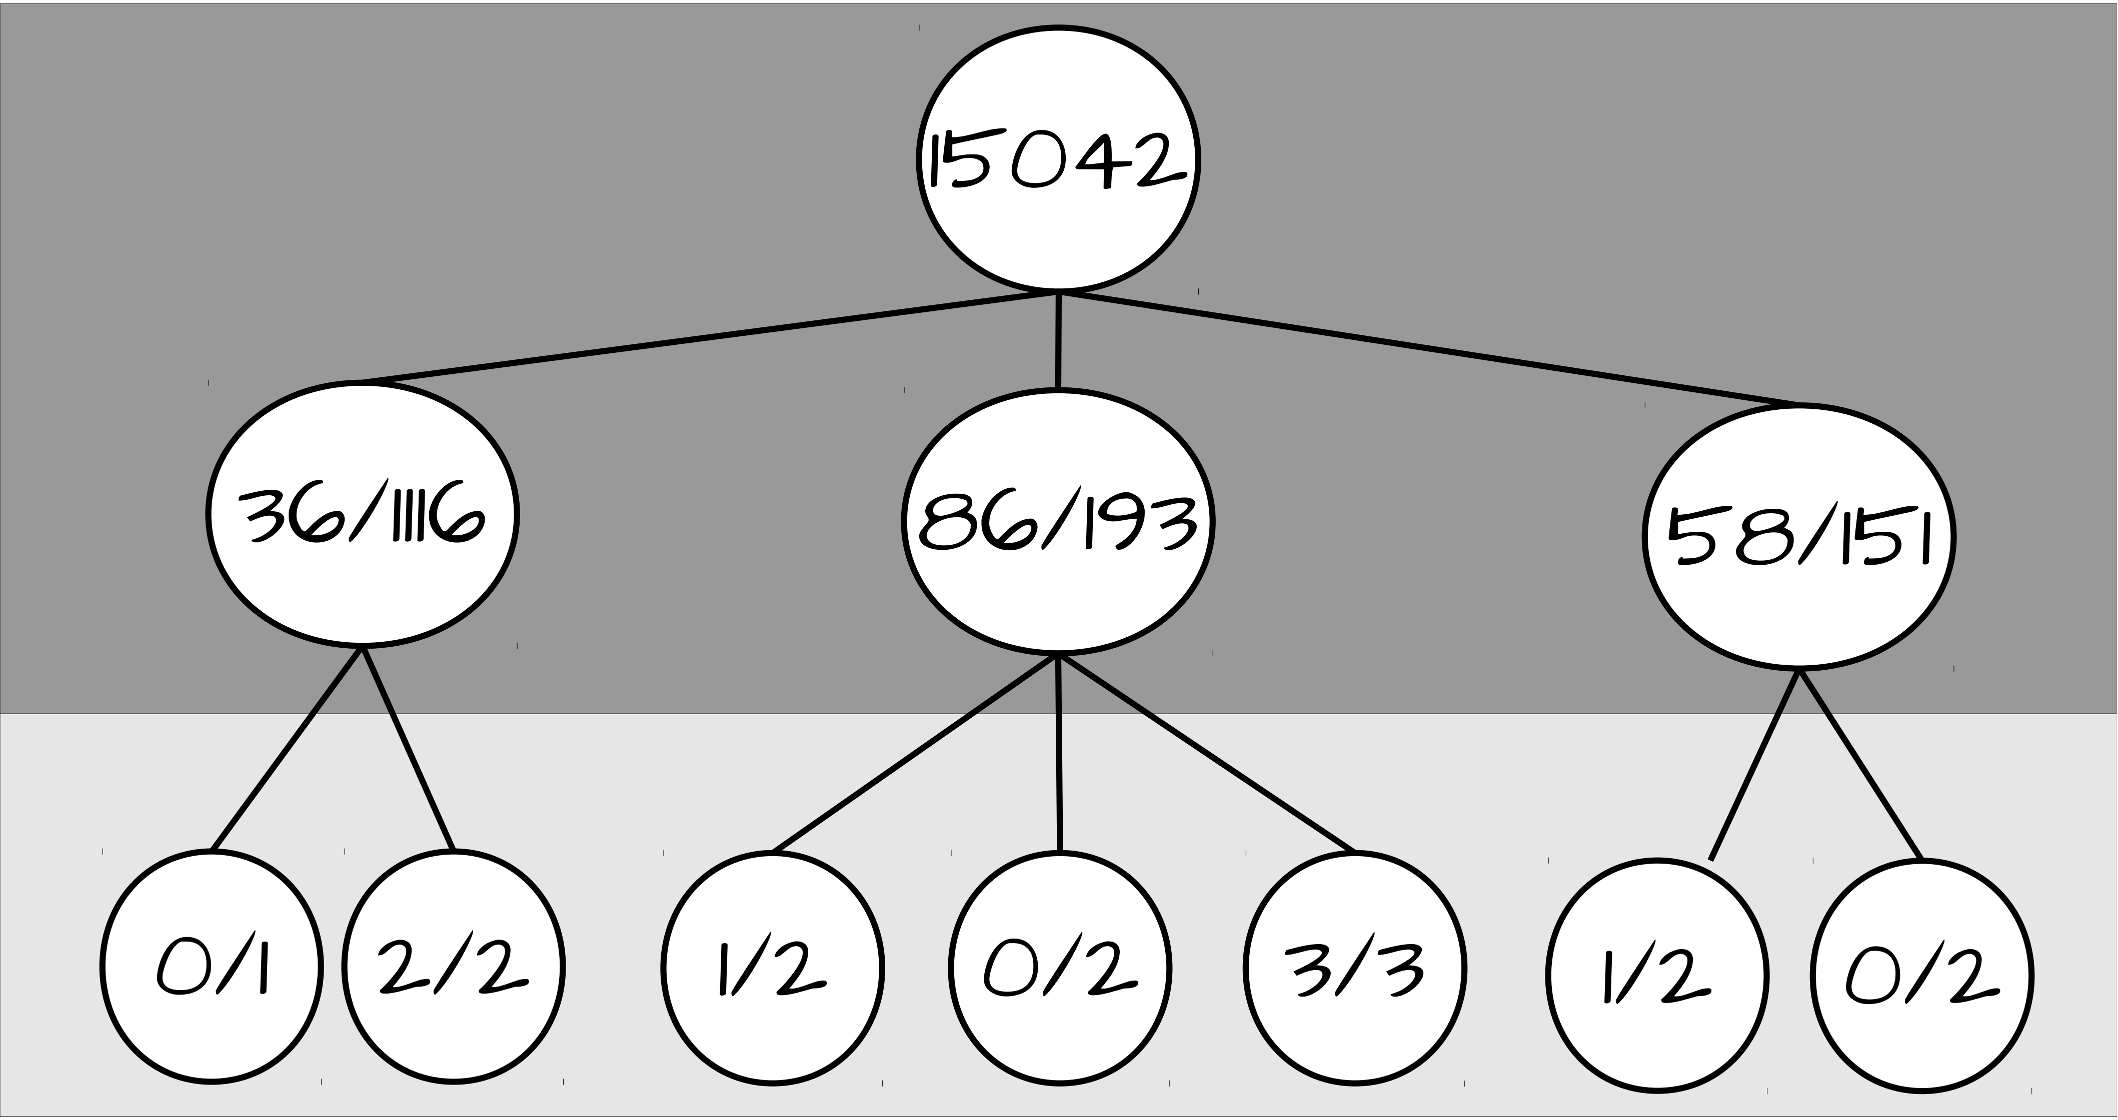
15042
36/1116
86/193
58/151
0/1
2/2
1/2
0/2
3/3
1/2
0/2

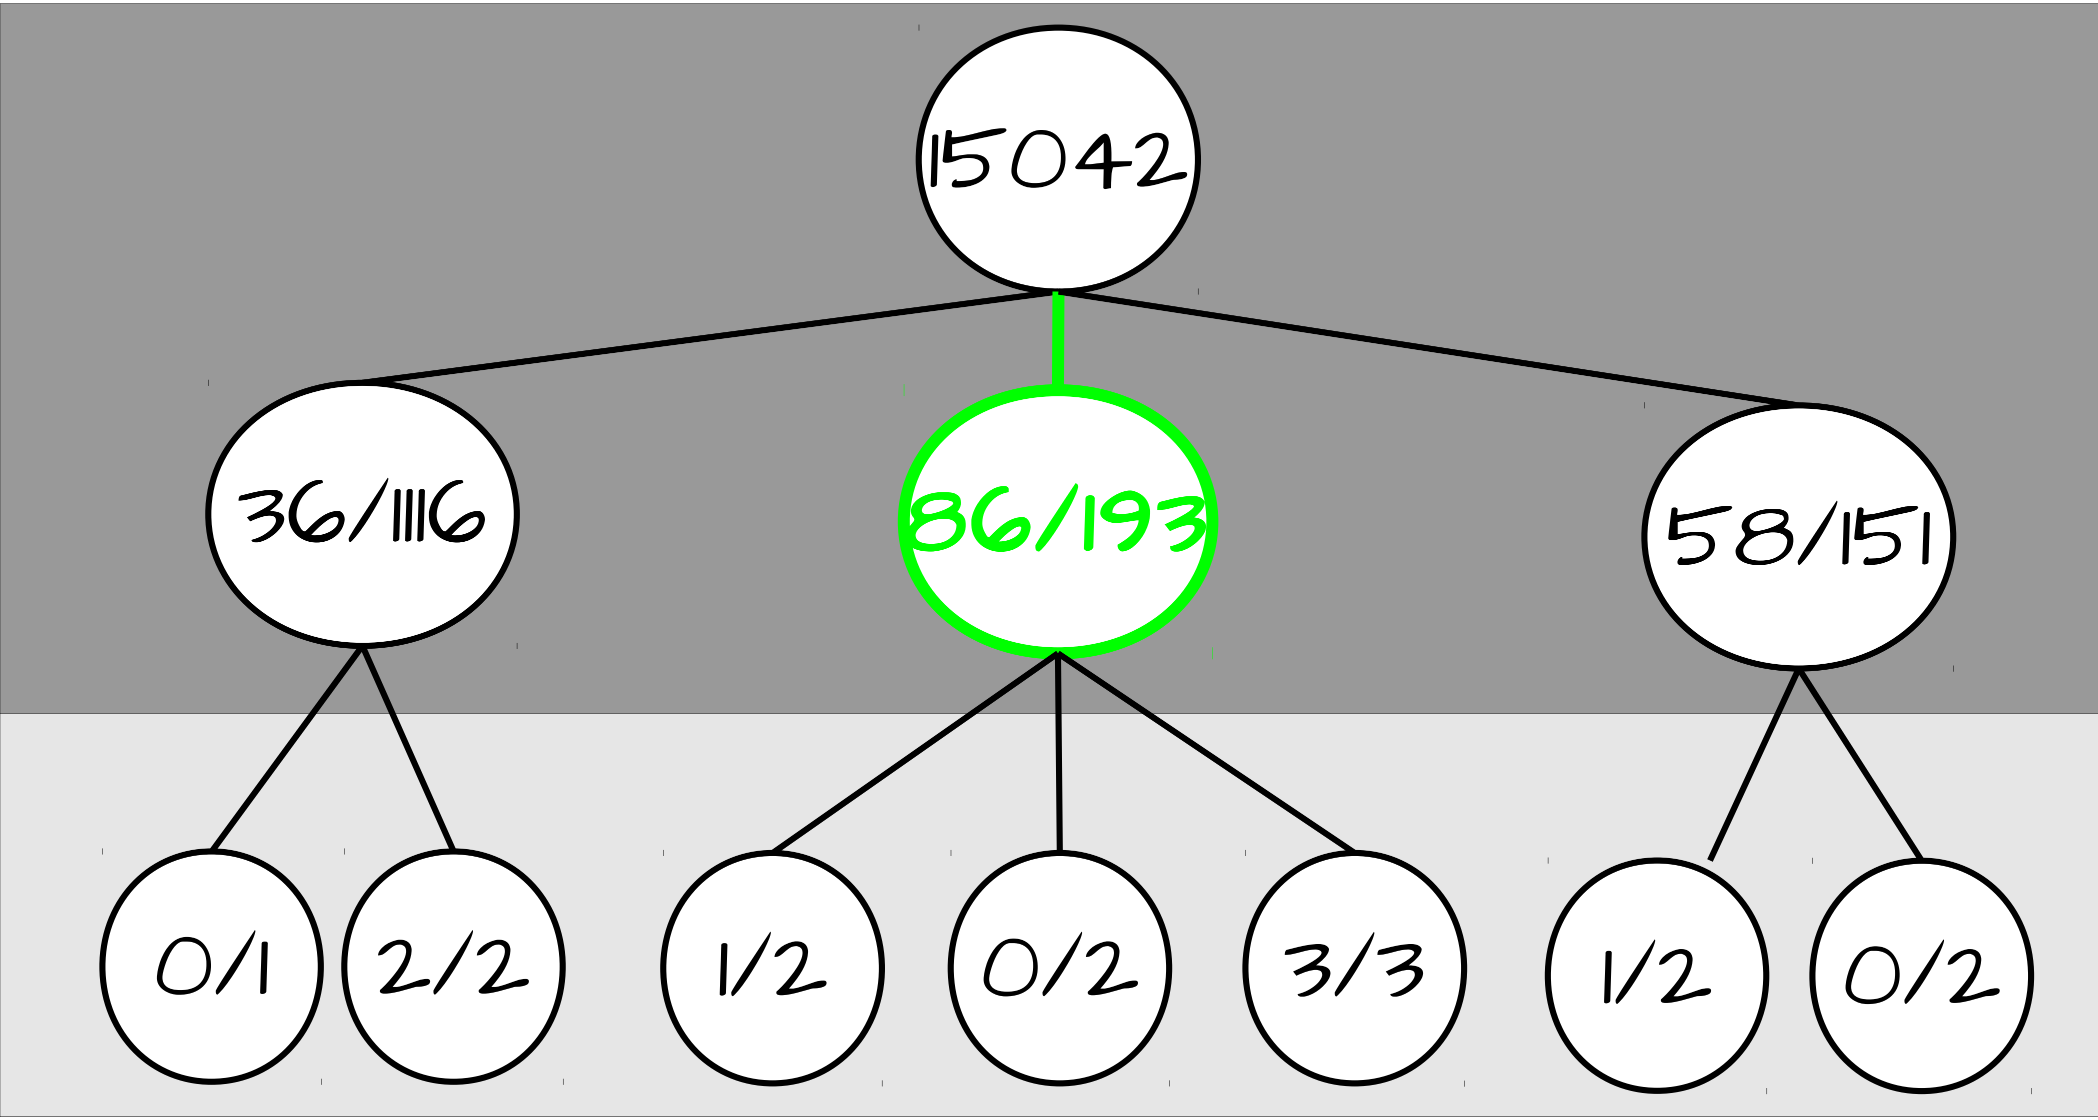
15042
36/1116
86/193
58/151
0/1
2/2
1/2
0/2
3/3
1/2
0/2

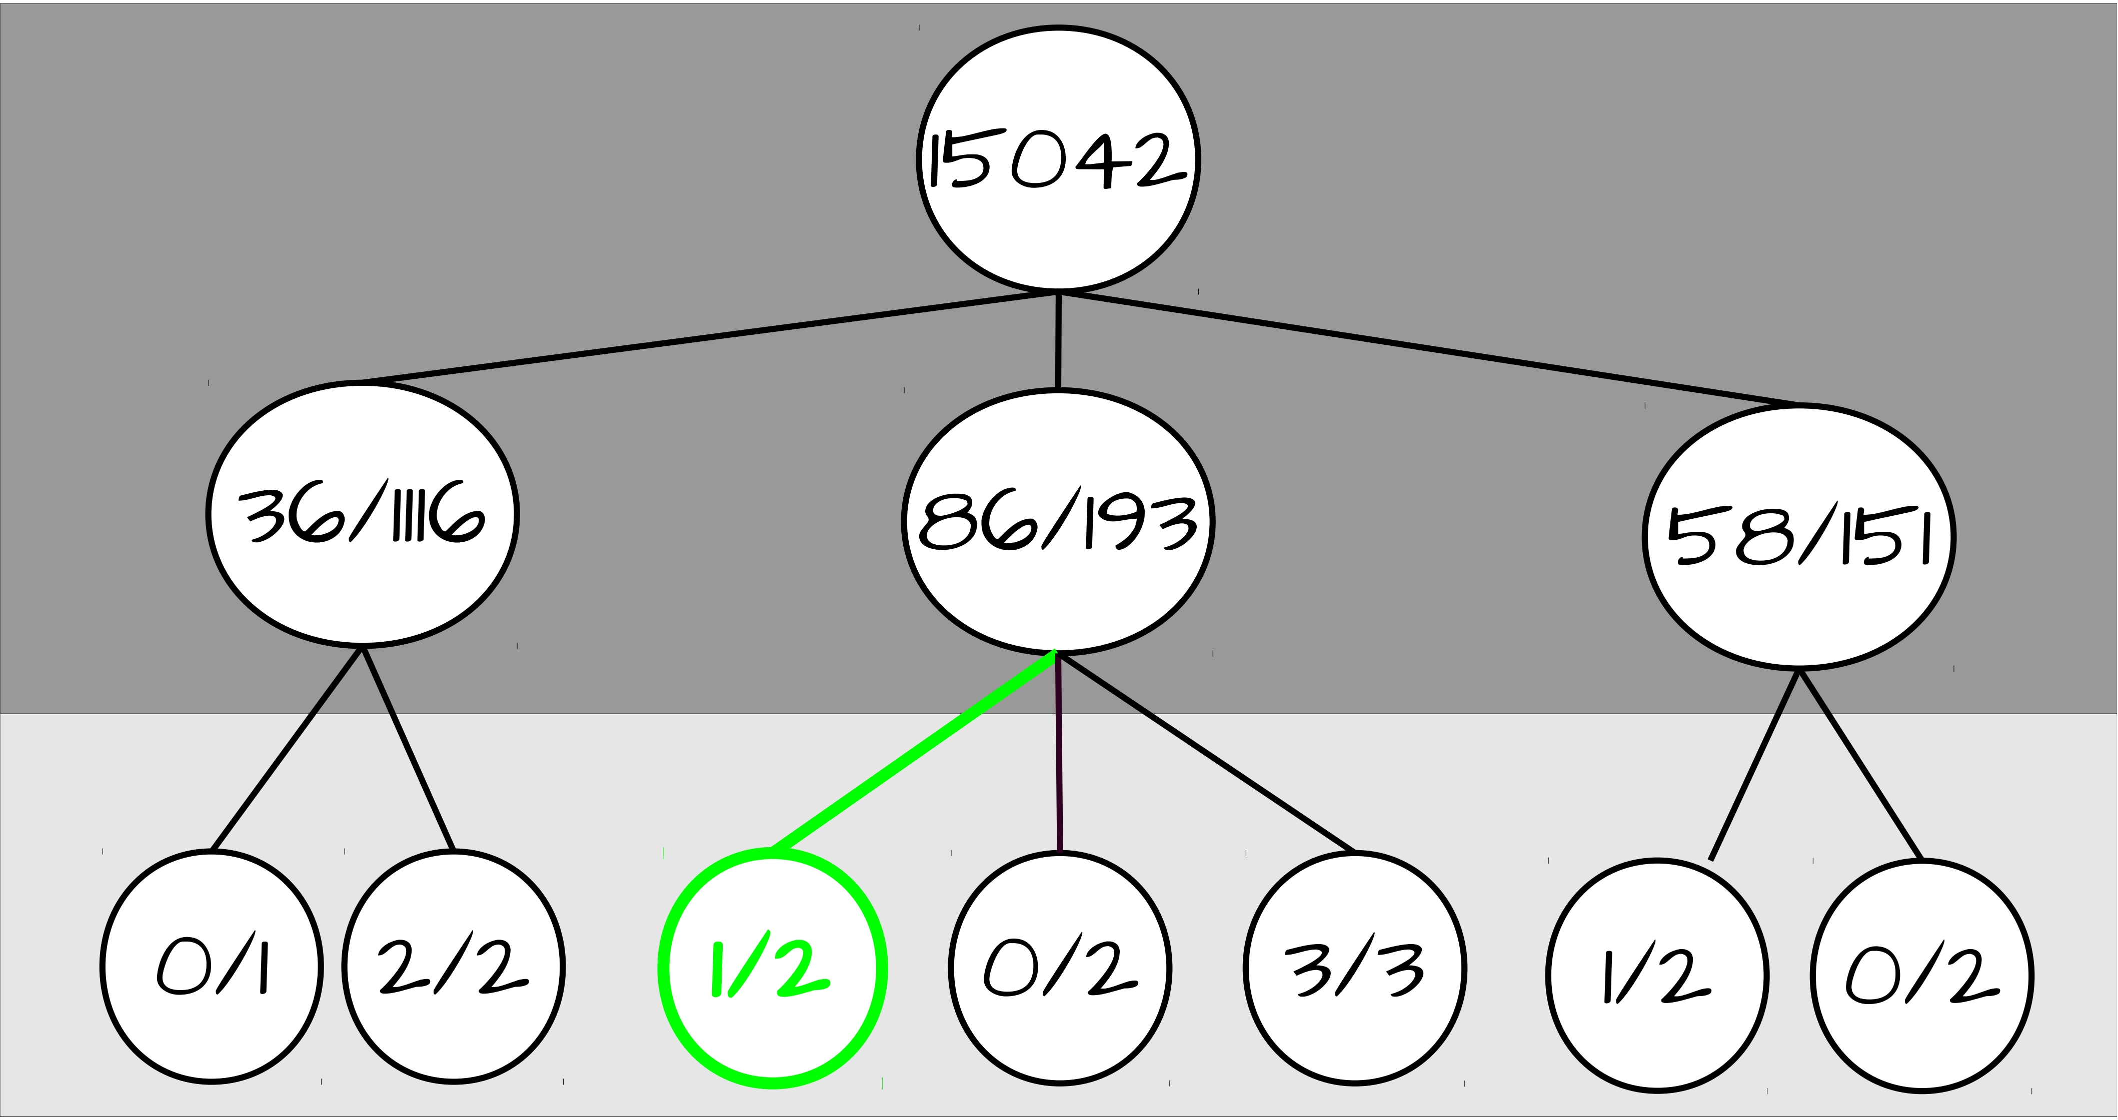
15042
36/1116
86/193
58/151
0/1
2/2
1/2
0/2
3/3
1/2
0/2


Not Human like?
© H. Michael Miley 2012 | mike@miley.us

Aheuristic

Generate a valid random move

Who has won?

General Game Playing

Anytime

Lazy

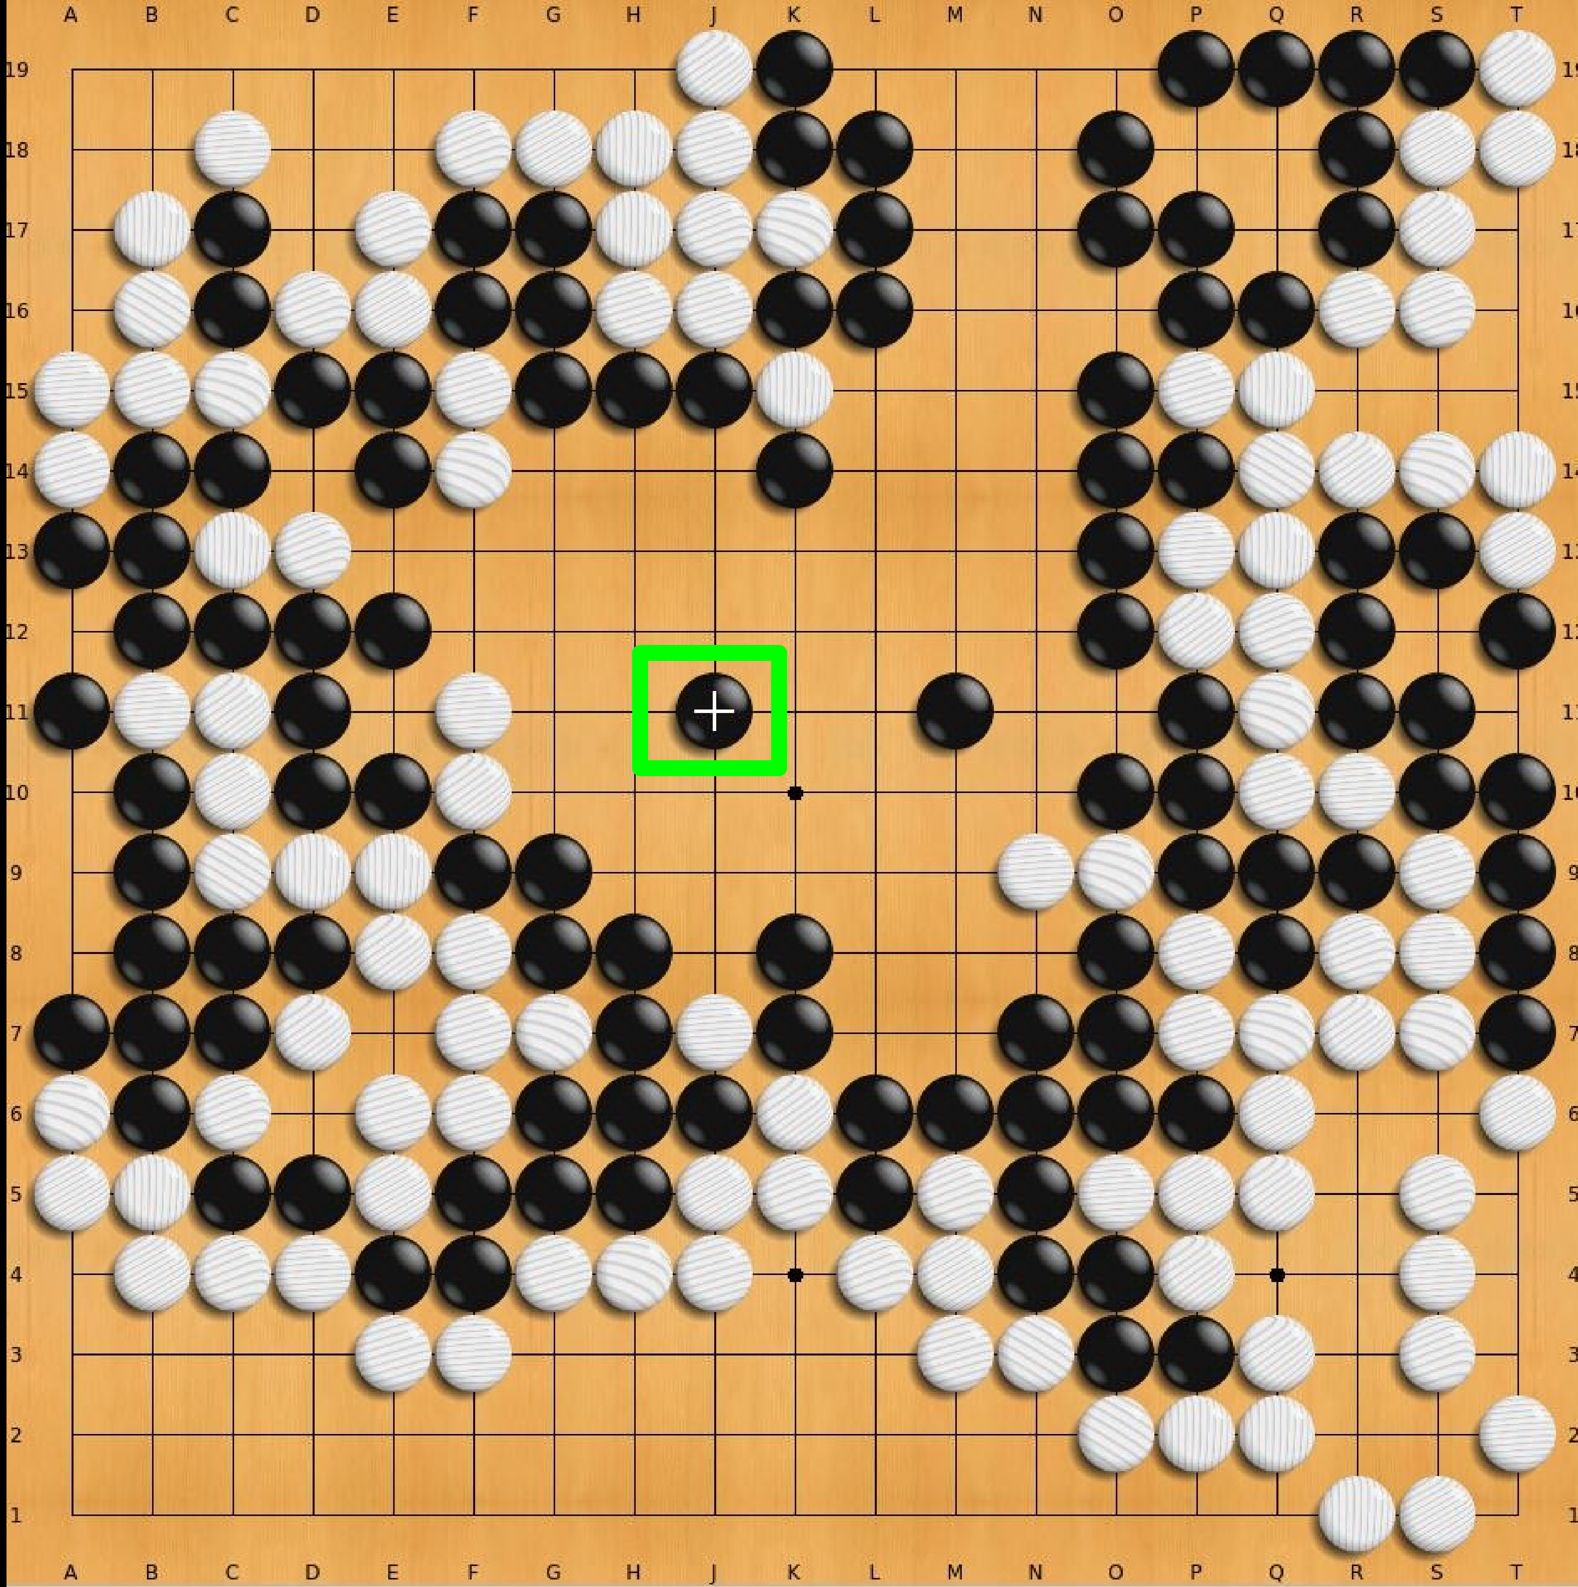

AMAF + RAVE

Expert Knowledge

Neural Networks

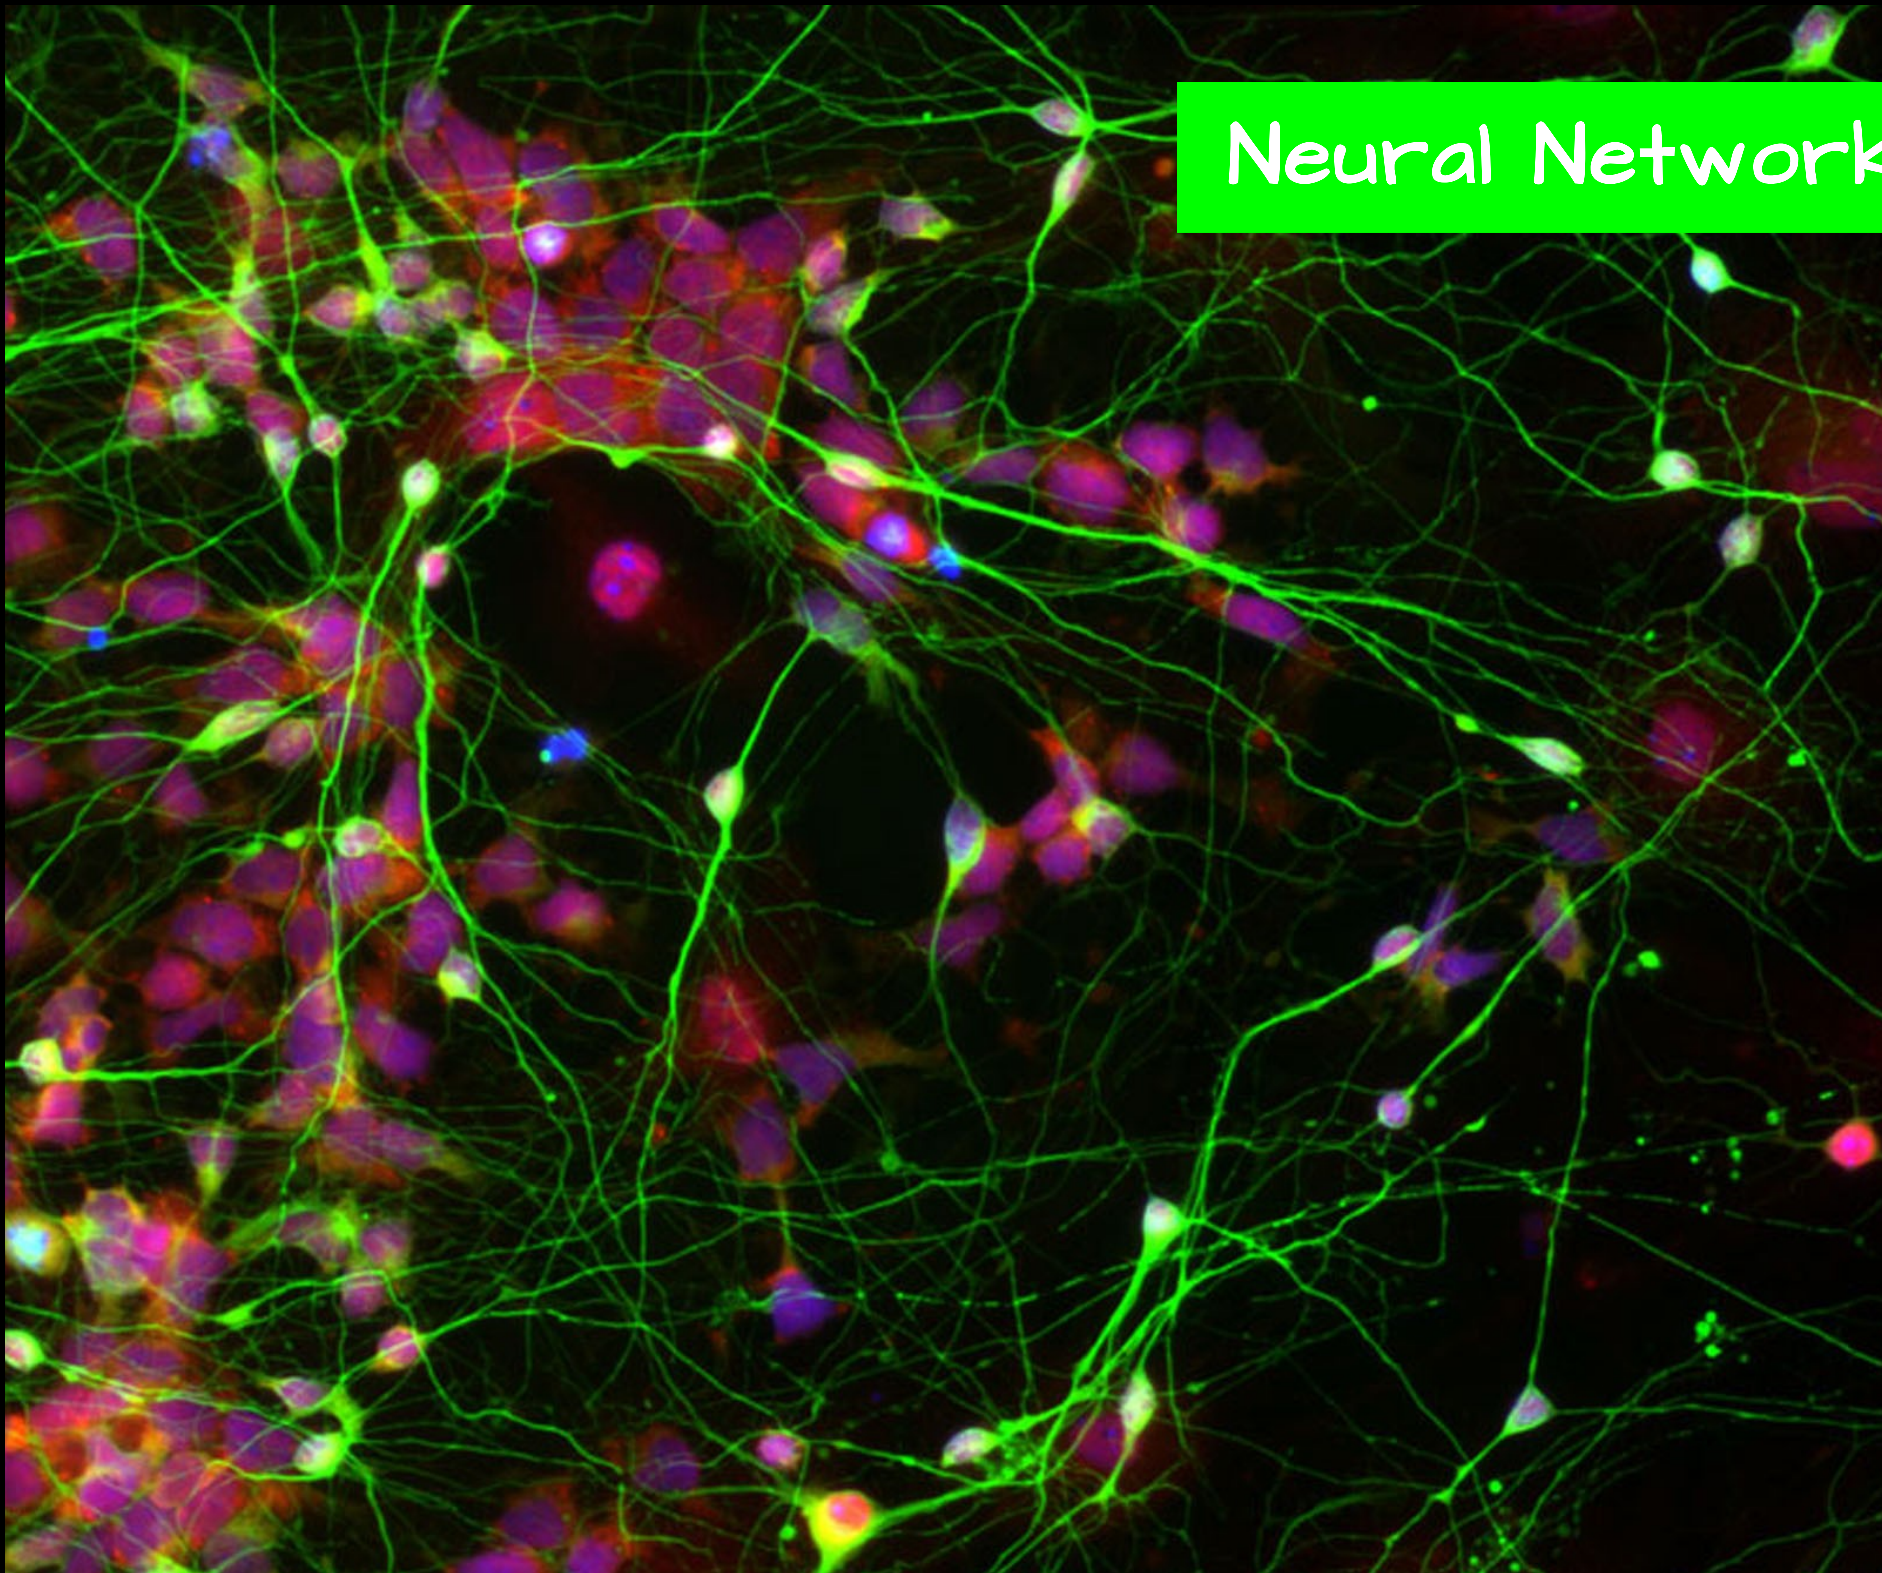

Move Evaluation in Go Using Deep Convolutional Neural Networks

Chris J. Maddison
University of Toronto
cmaddis@cs.toronto.edu

Aja Huang[1], Ilya Sutskever[2], David Silver[1]
Google DeepMind[1], Google Brain[2]
{ajahuang,ilyasu,davidsilver}@google.com

2014

Abstract

The game of Go is more challenging than other board games, due to the difficulty of constructing a position or move evaluation function. In this paper we investigate whether deep convolutional networks can be used to directly represent and learn this knowledge. We train a large 12-layer convolutional neural network by supervised learning from a database of human professional games. The network correctly predicts the expert move in 55% of positions, equalling the accuracy of a 6 dan human player. When the trained convolutional network was used di-

What does this even mean?

Neural Networks

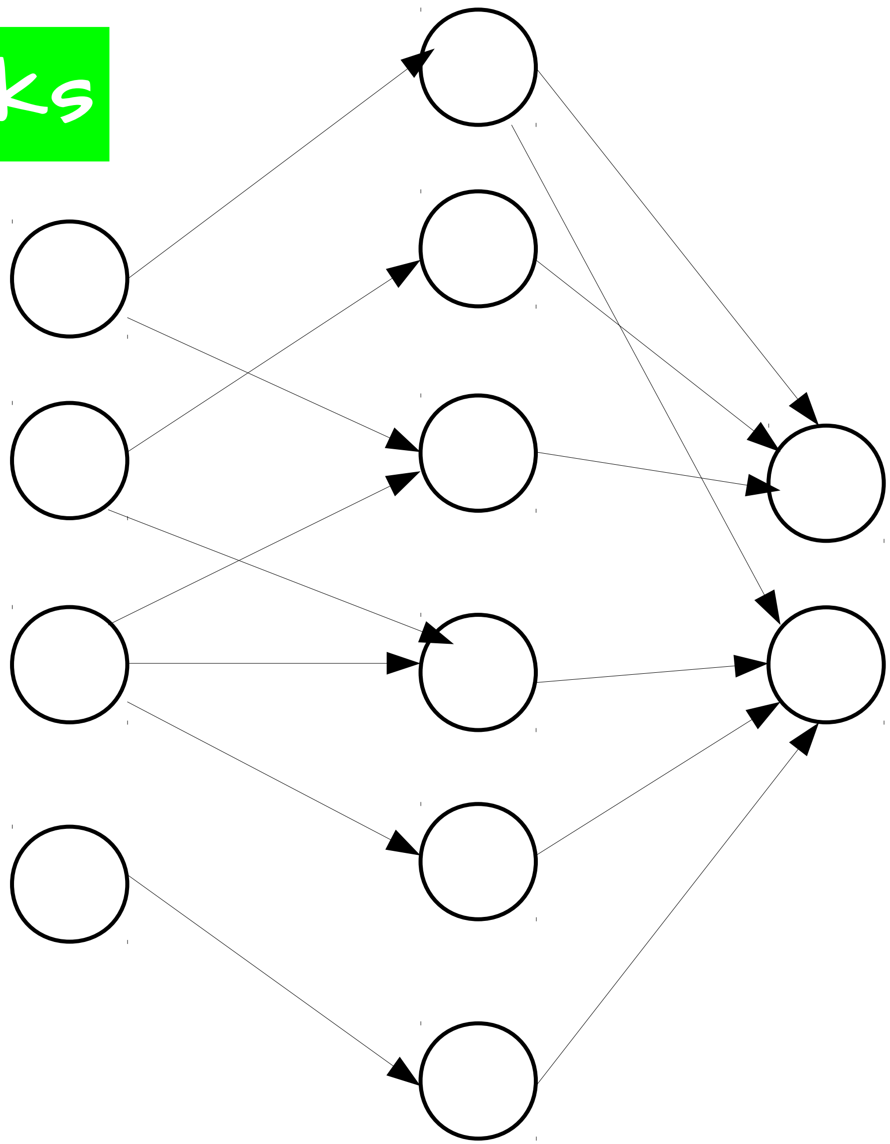

Neural Networks

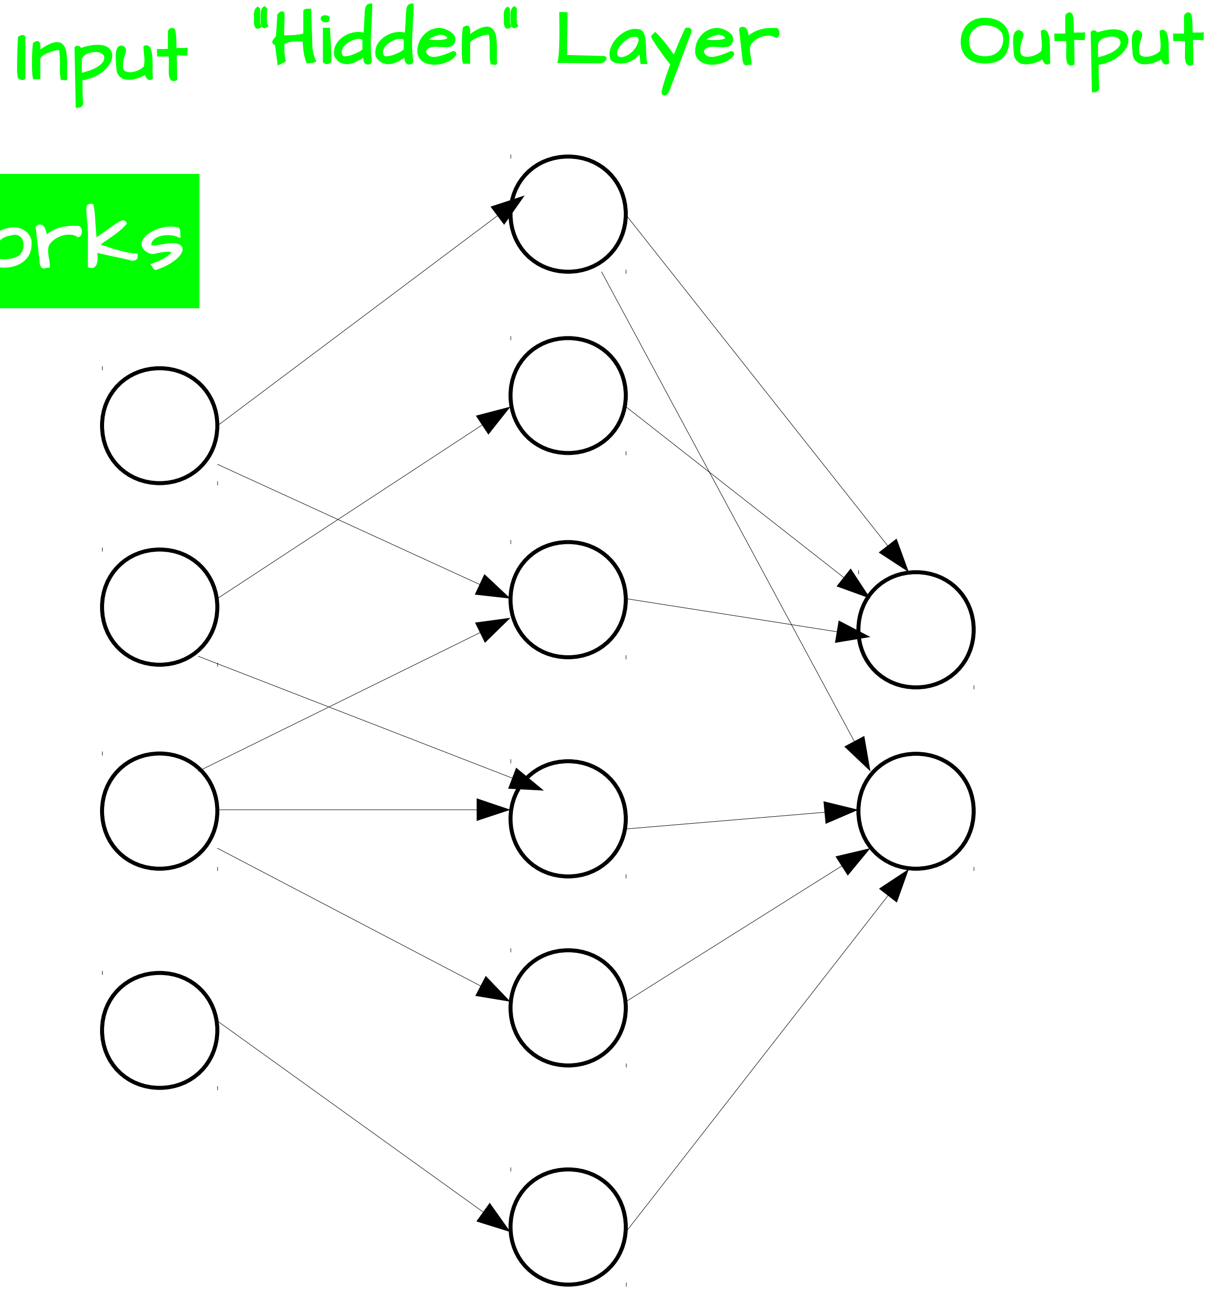

Weights

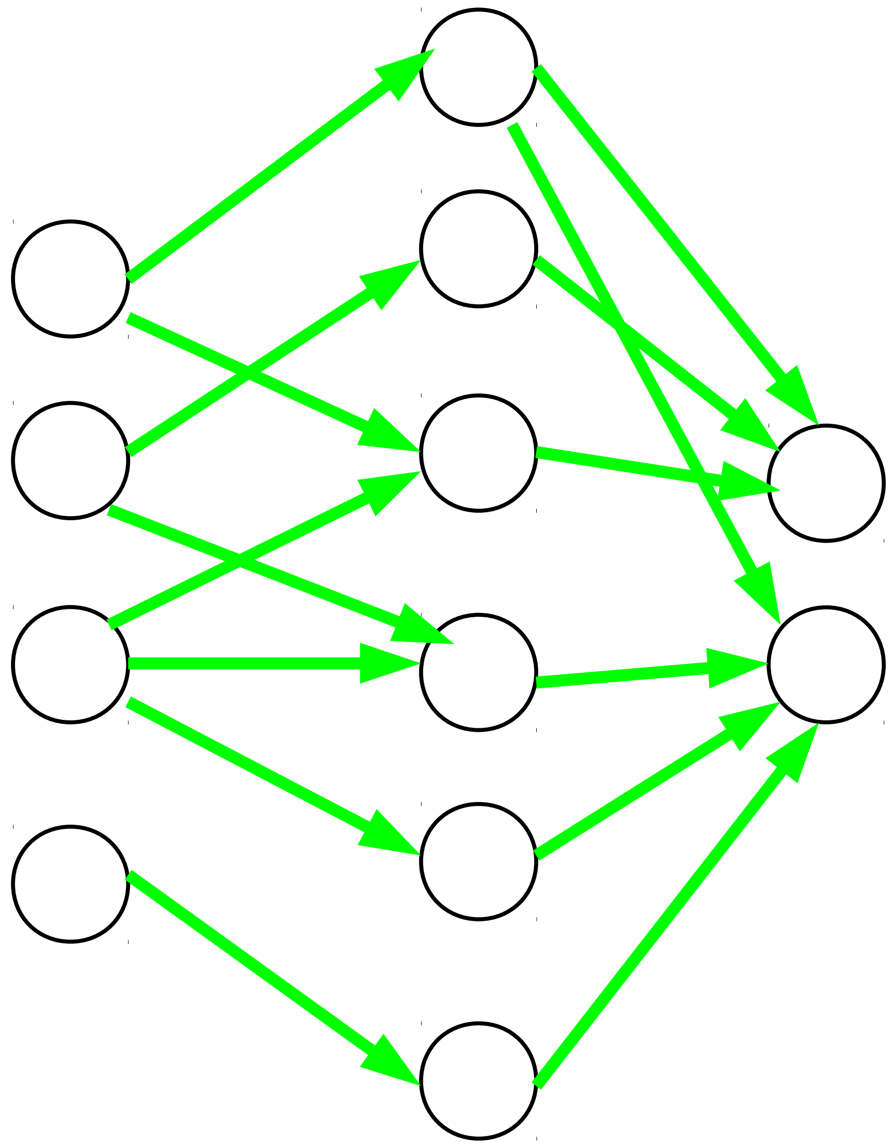

Bias/Threshold

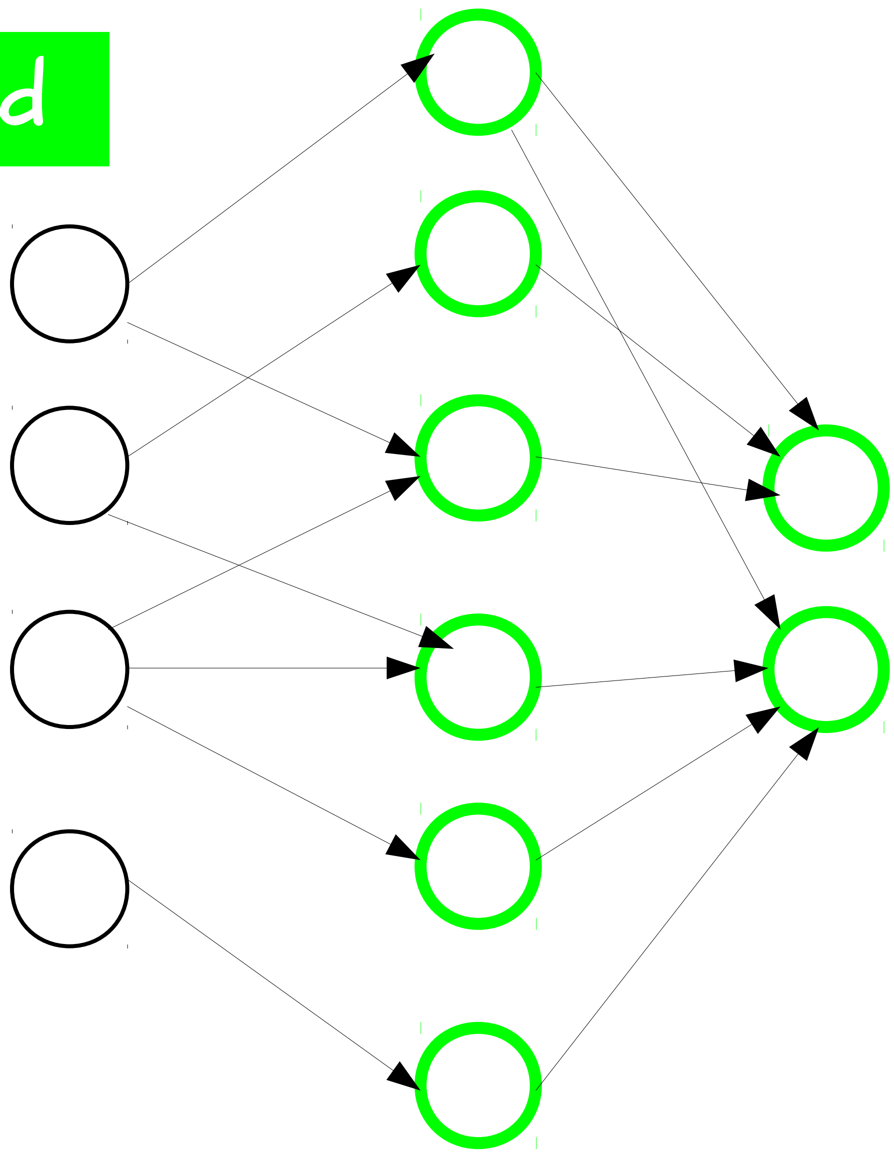

Sum of Weights >= Threshold

Activation

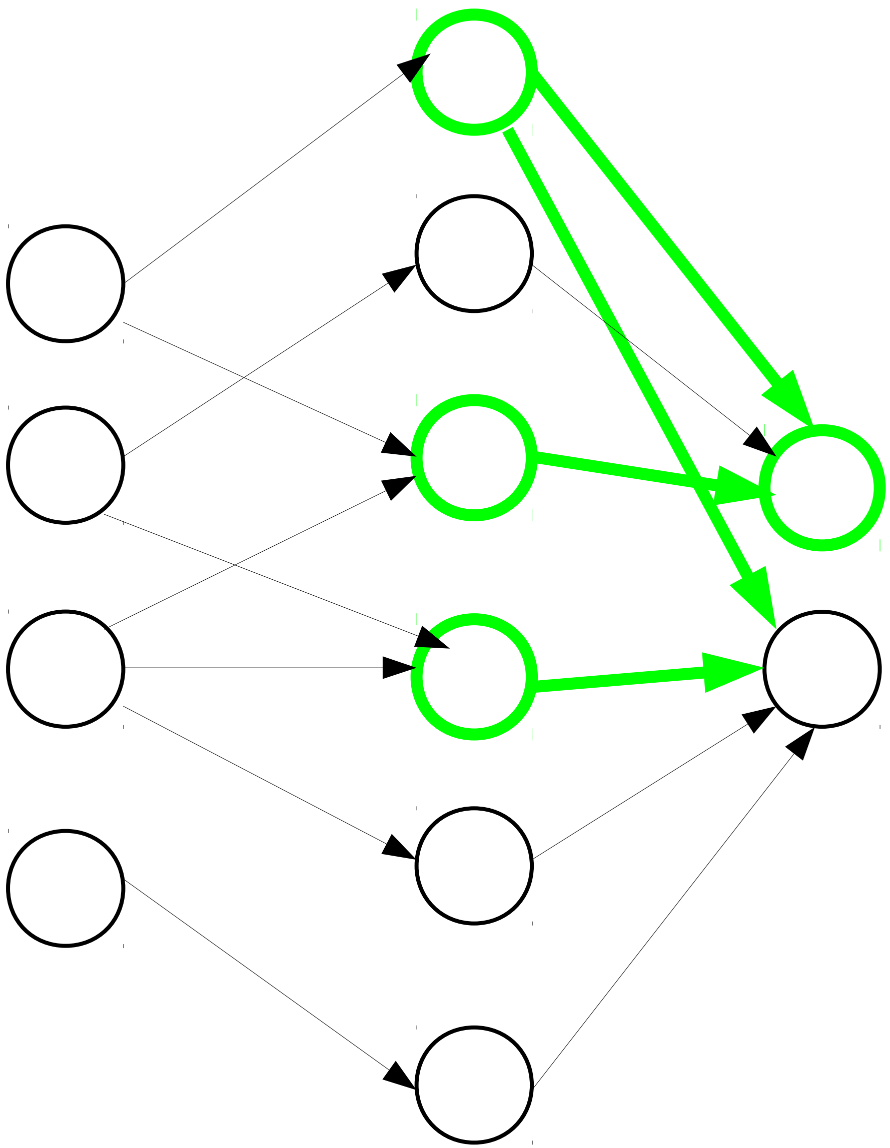

Training

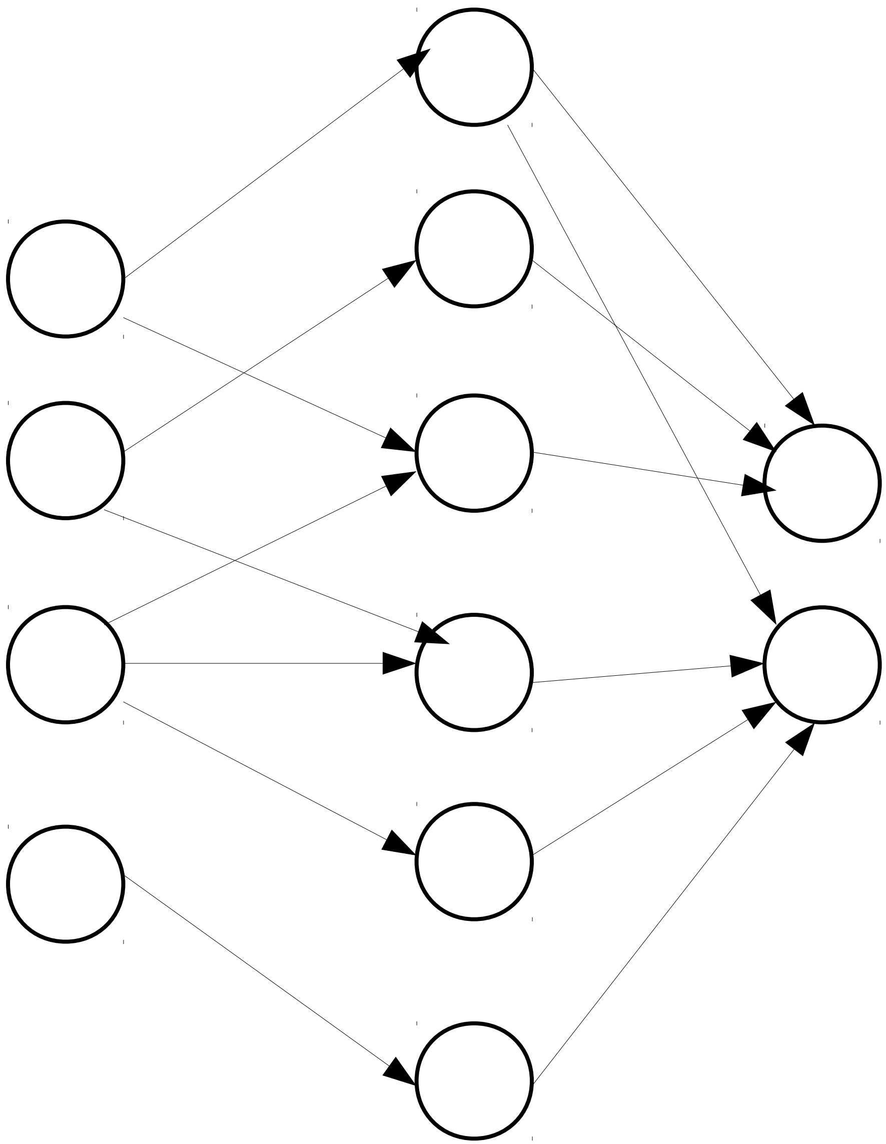

Adjust parameters

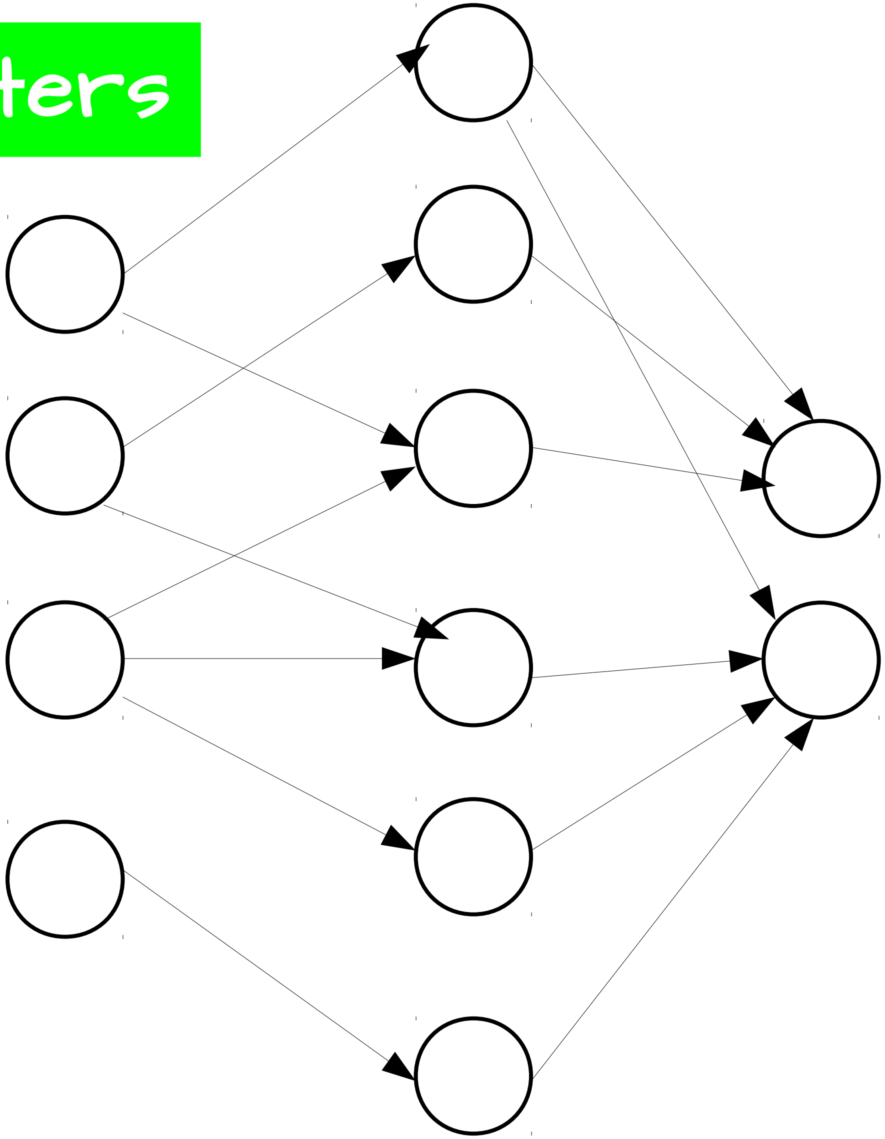

Supervised Learning

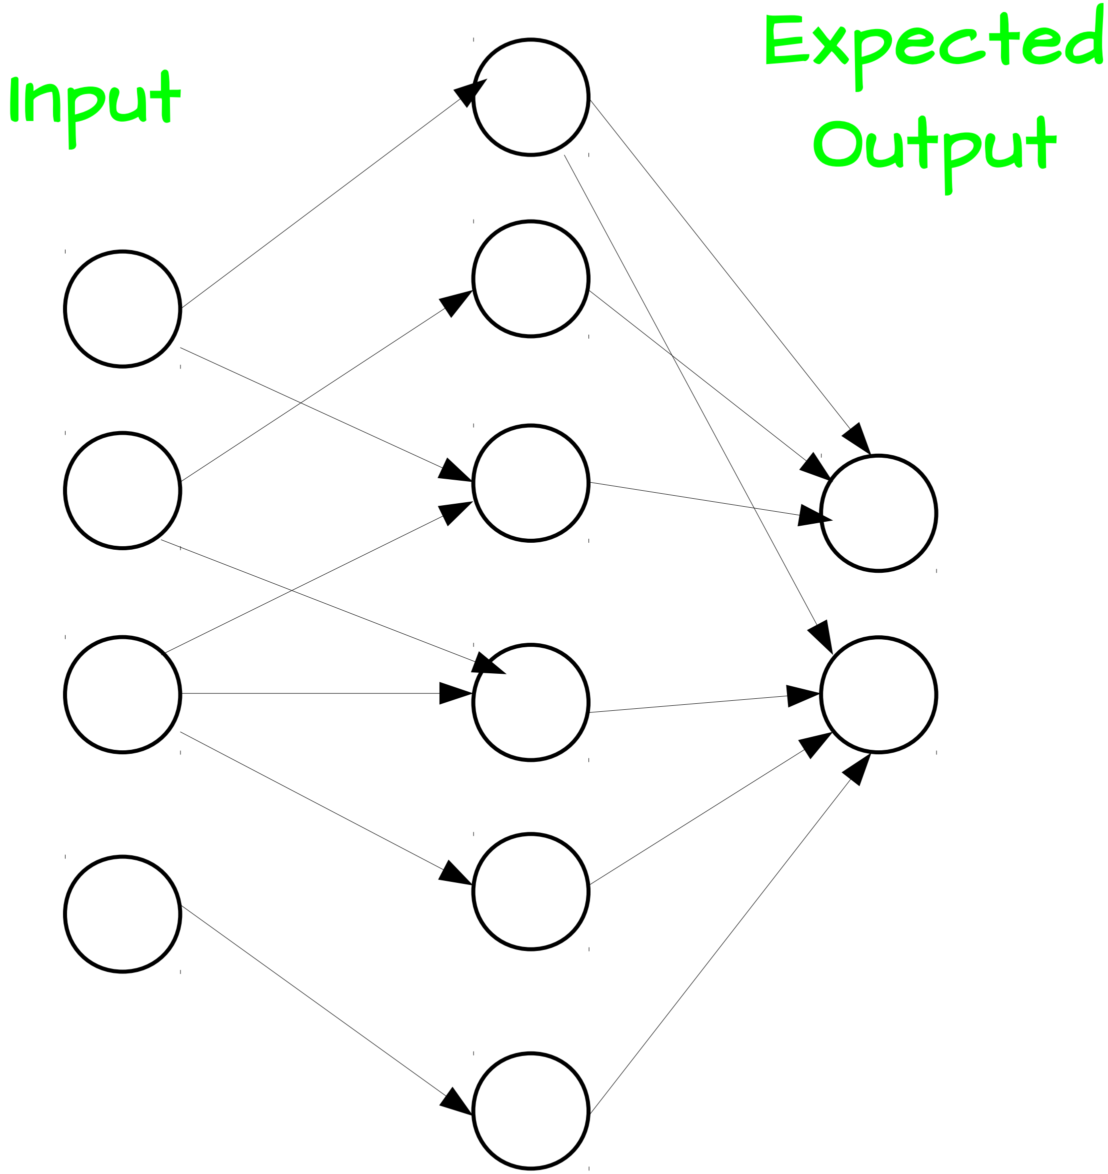

Backpropagation

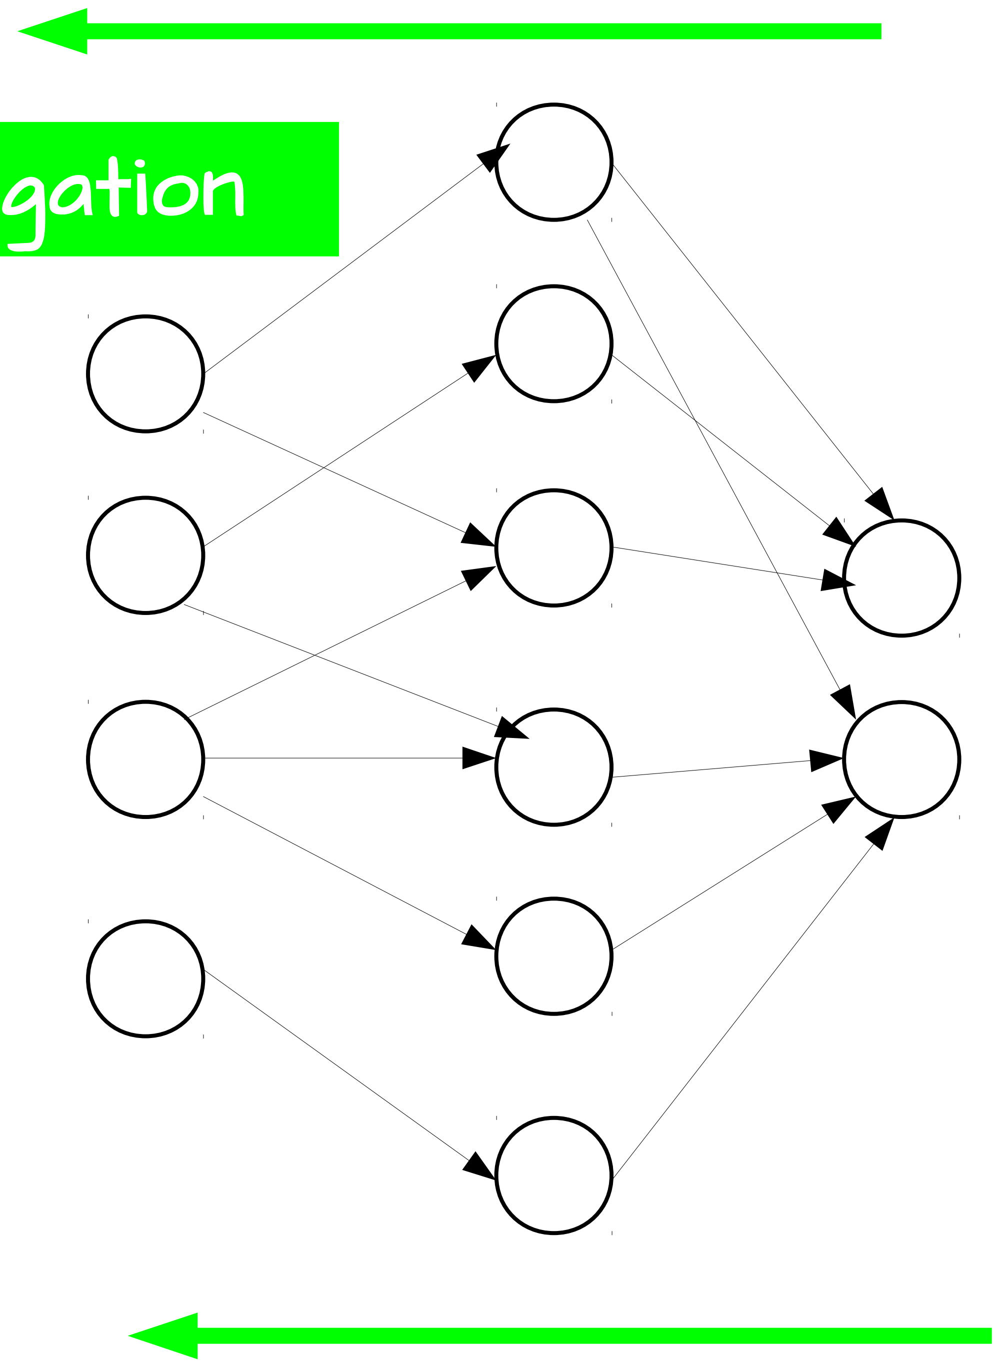

Data set

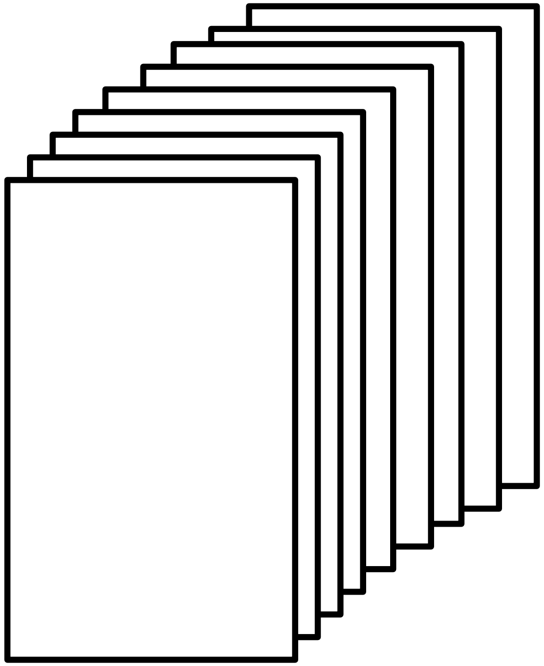

Training data + test data

Training

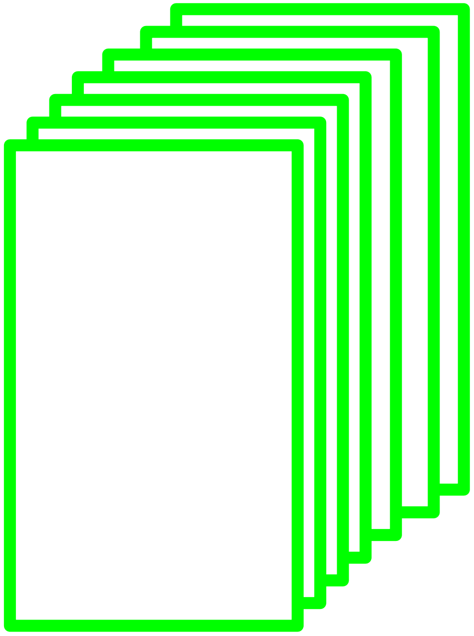

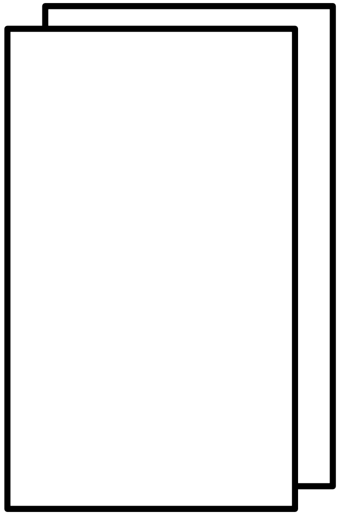

Verify

Overfitting

Deep Neural Networks

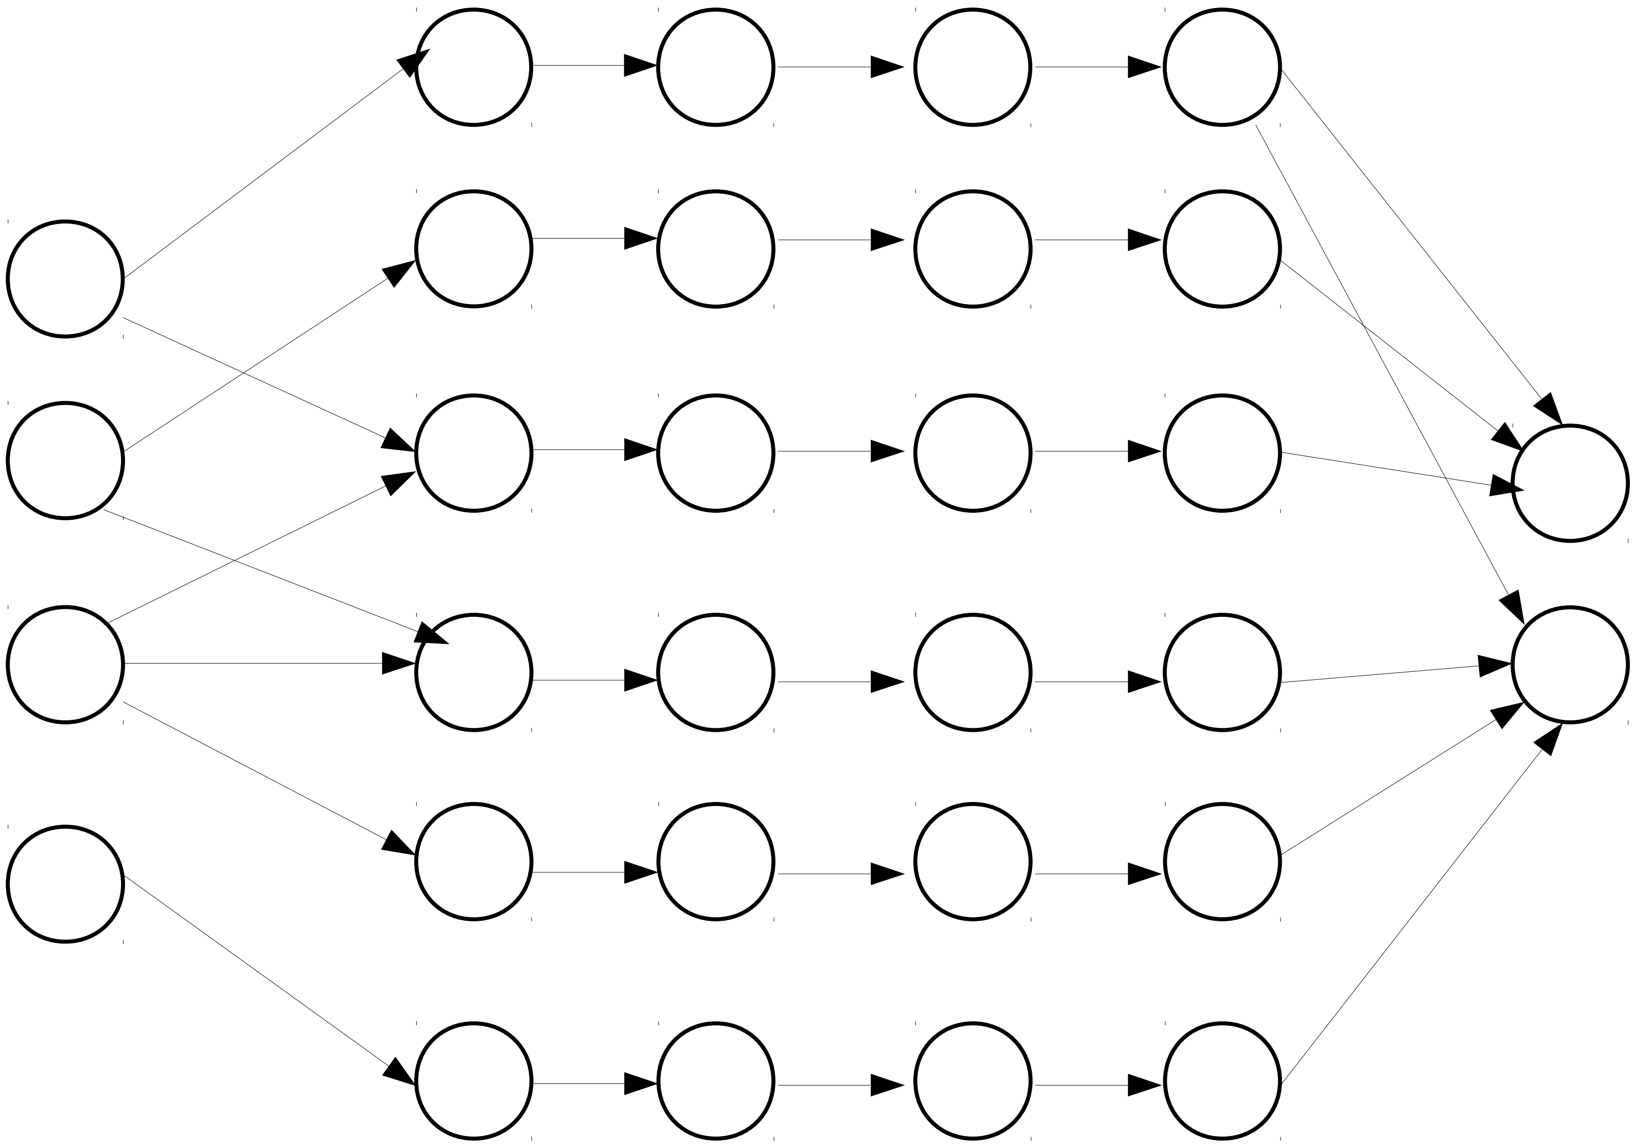

Convolutional Neural Networks

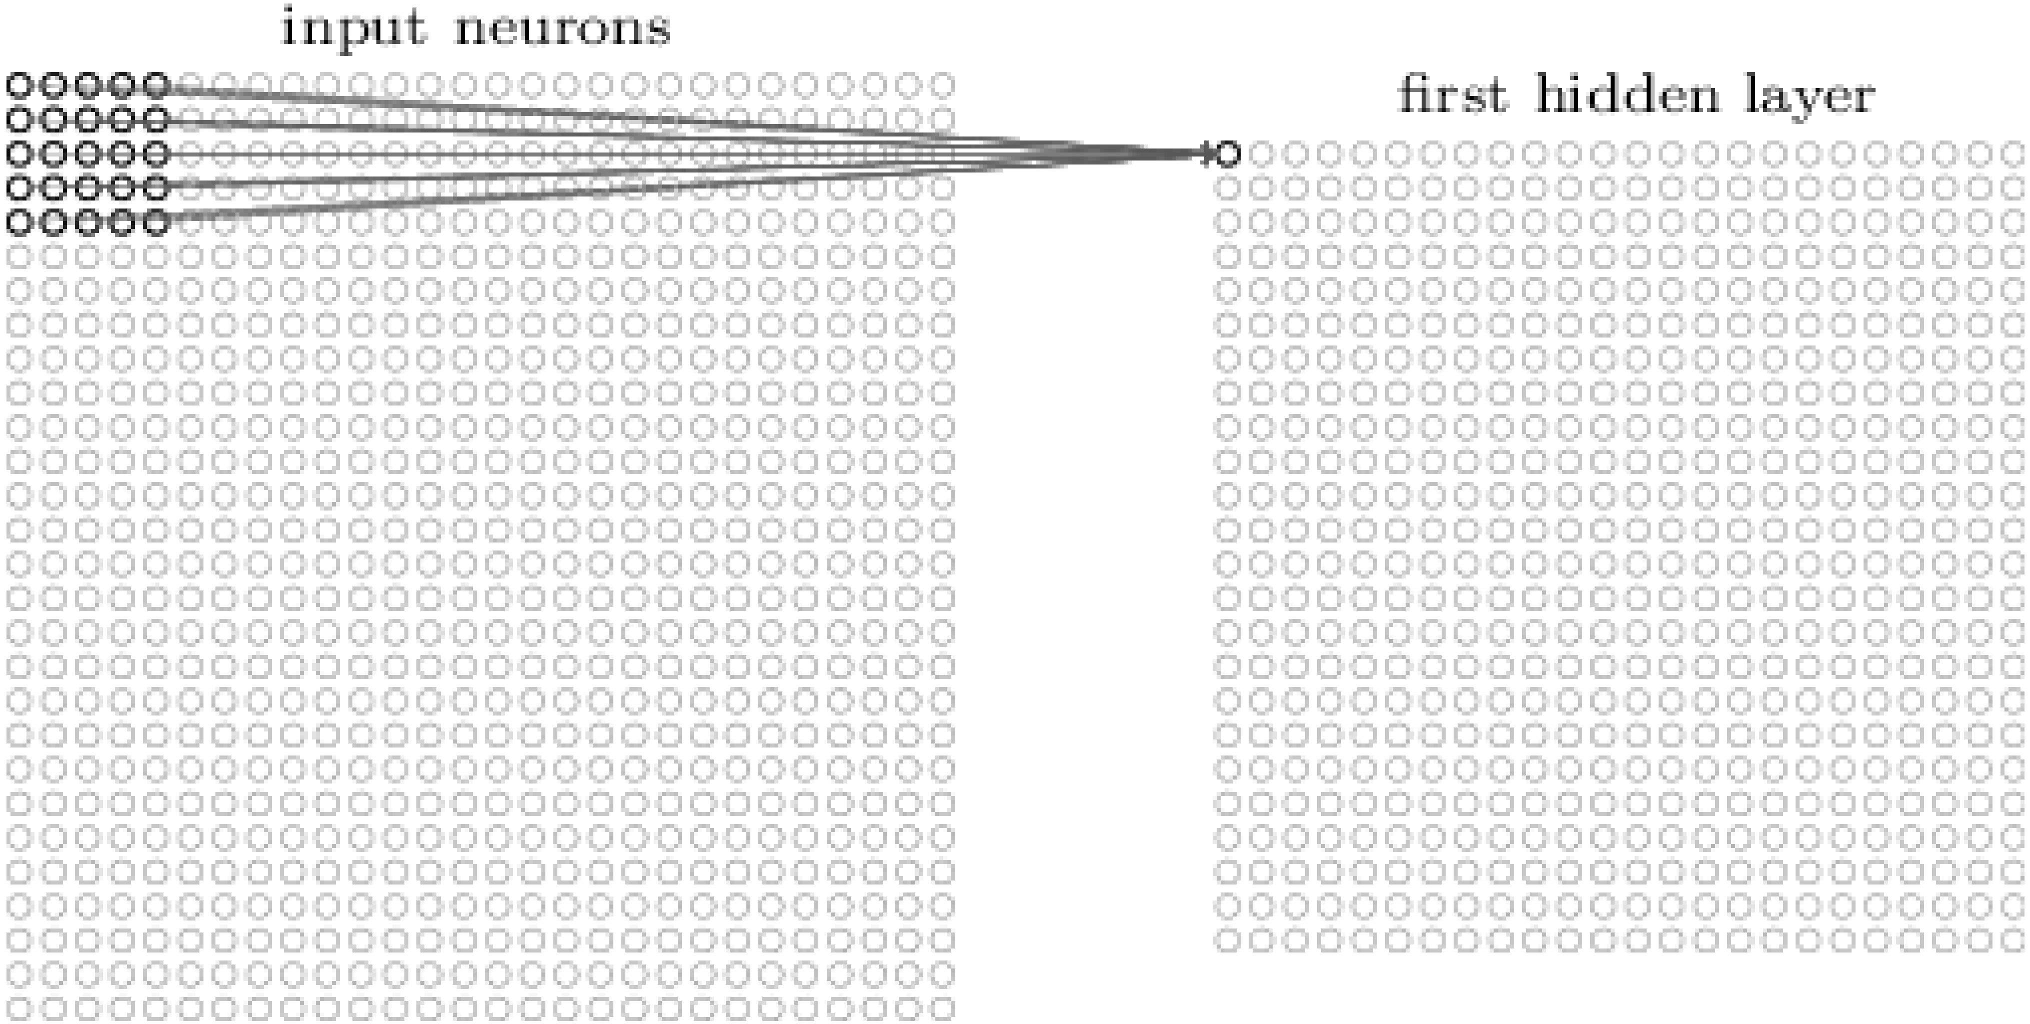

Michael A. Nielsen, "Neural Networks and Deep Learning", Determination Press, 2015
http://neuralnetworksanddeeplearning.com

Local Receptive Field

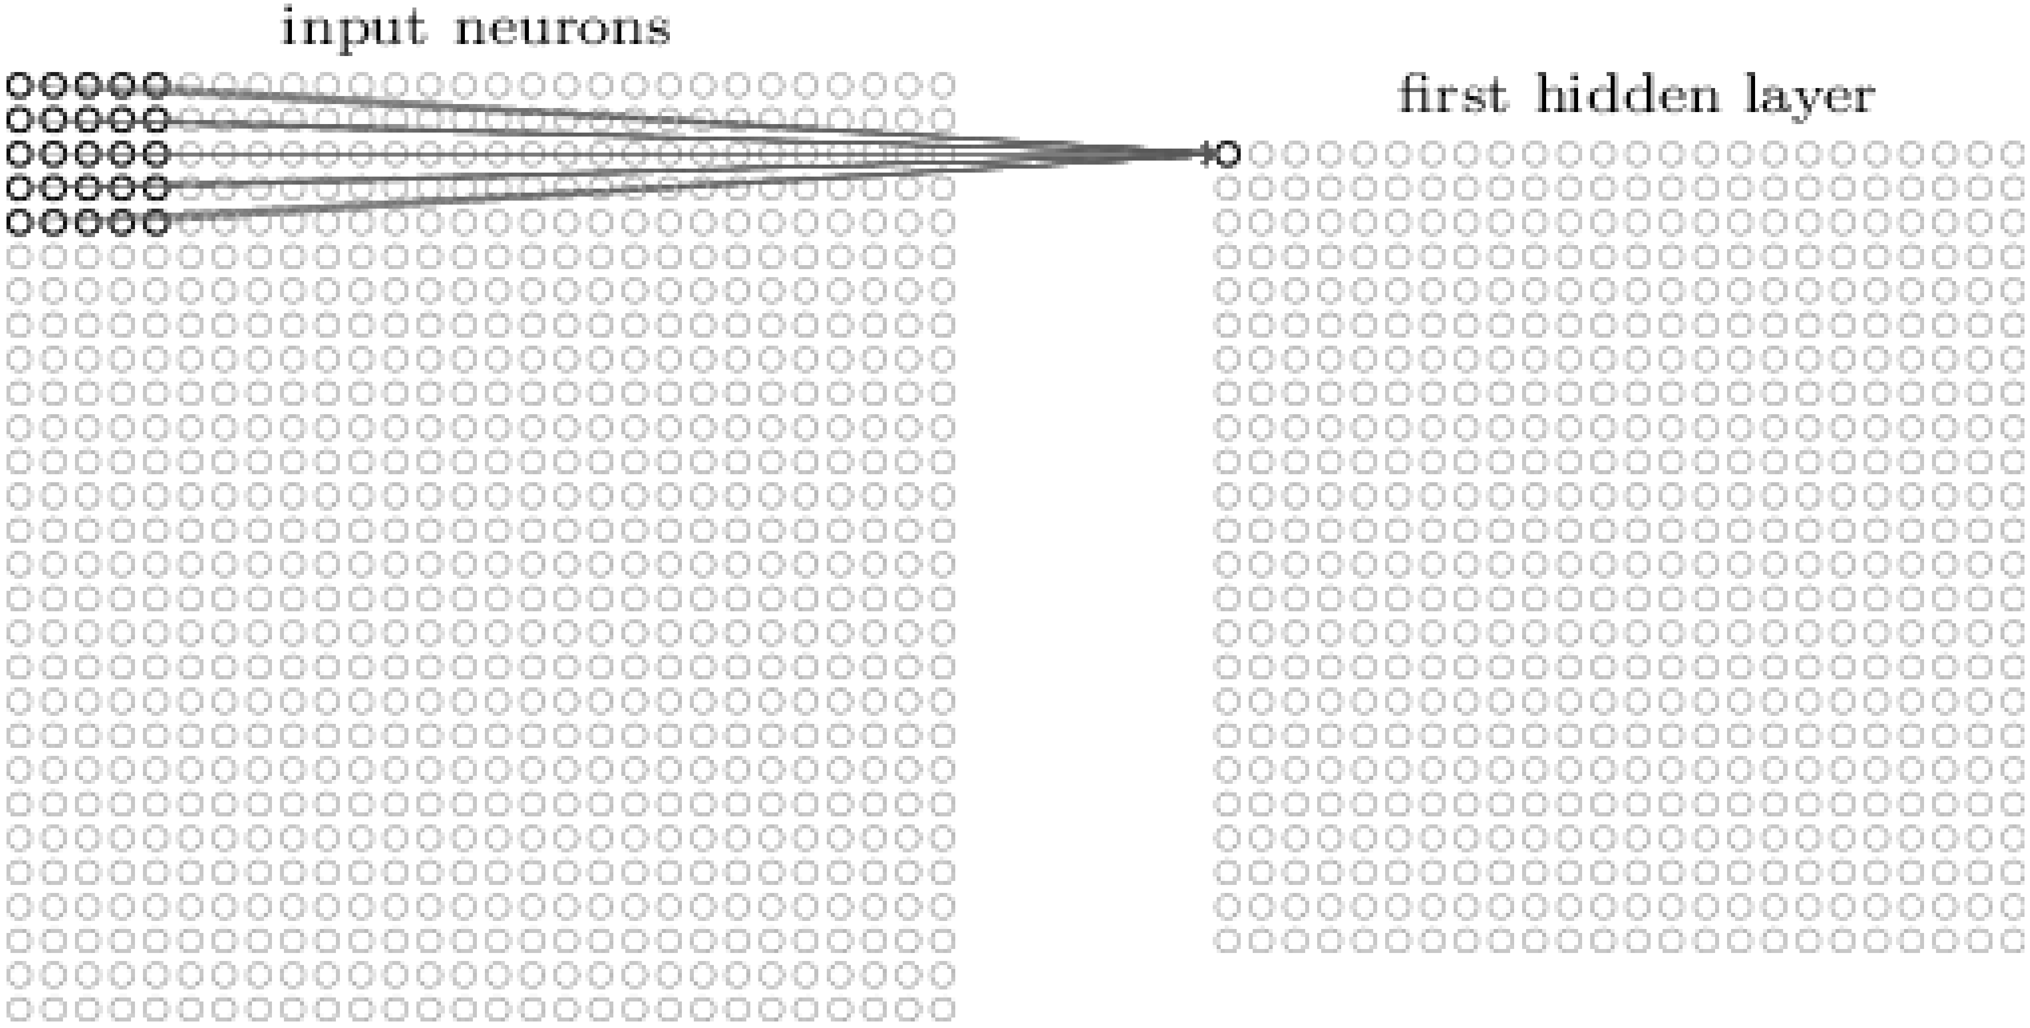

Michael A. Nielsen, "Neural Networks and Deep Learning", Determination Press, 2015
http://neuralnetworksanddeeplearning.com

Stride

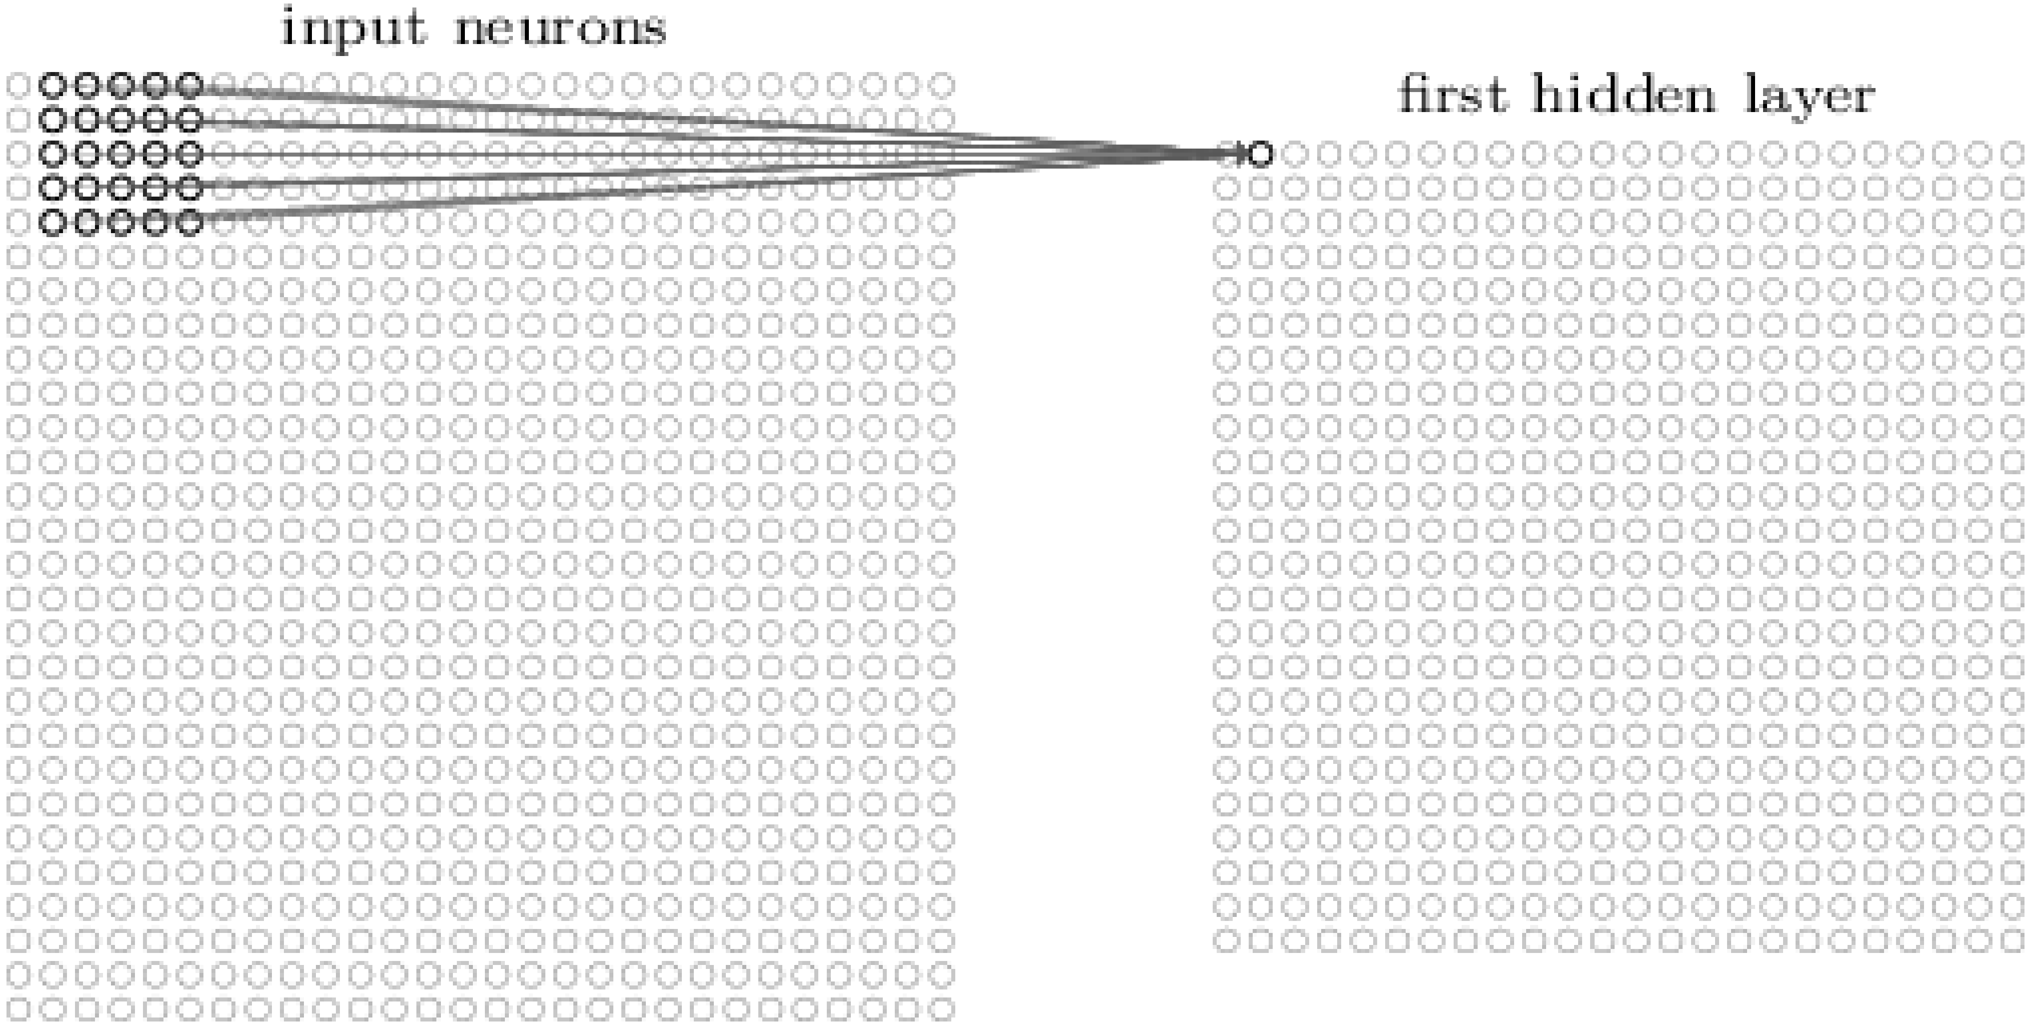

Michael A. Nielsen, "Neural Networks and Deep Learning", Determination Press, 2015
http://neuralnetworksanddeeplearning.com

Shared weights and biases

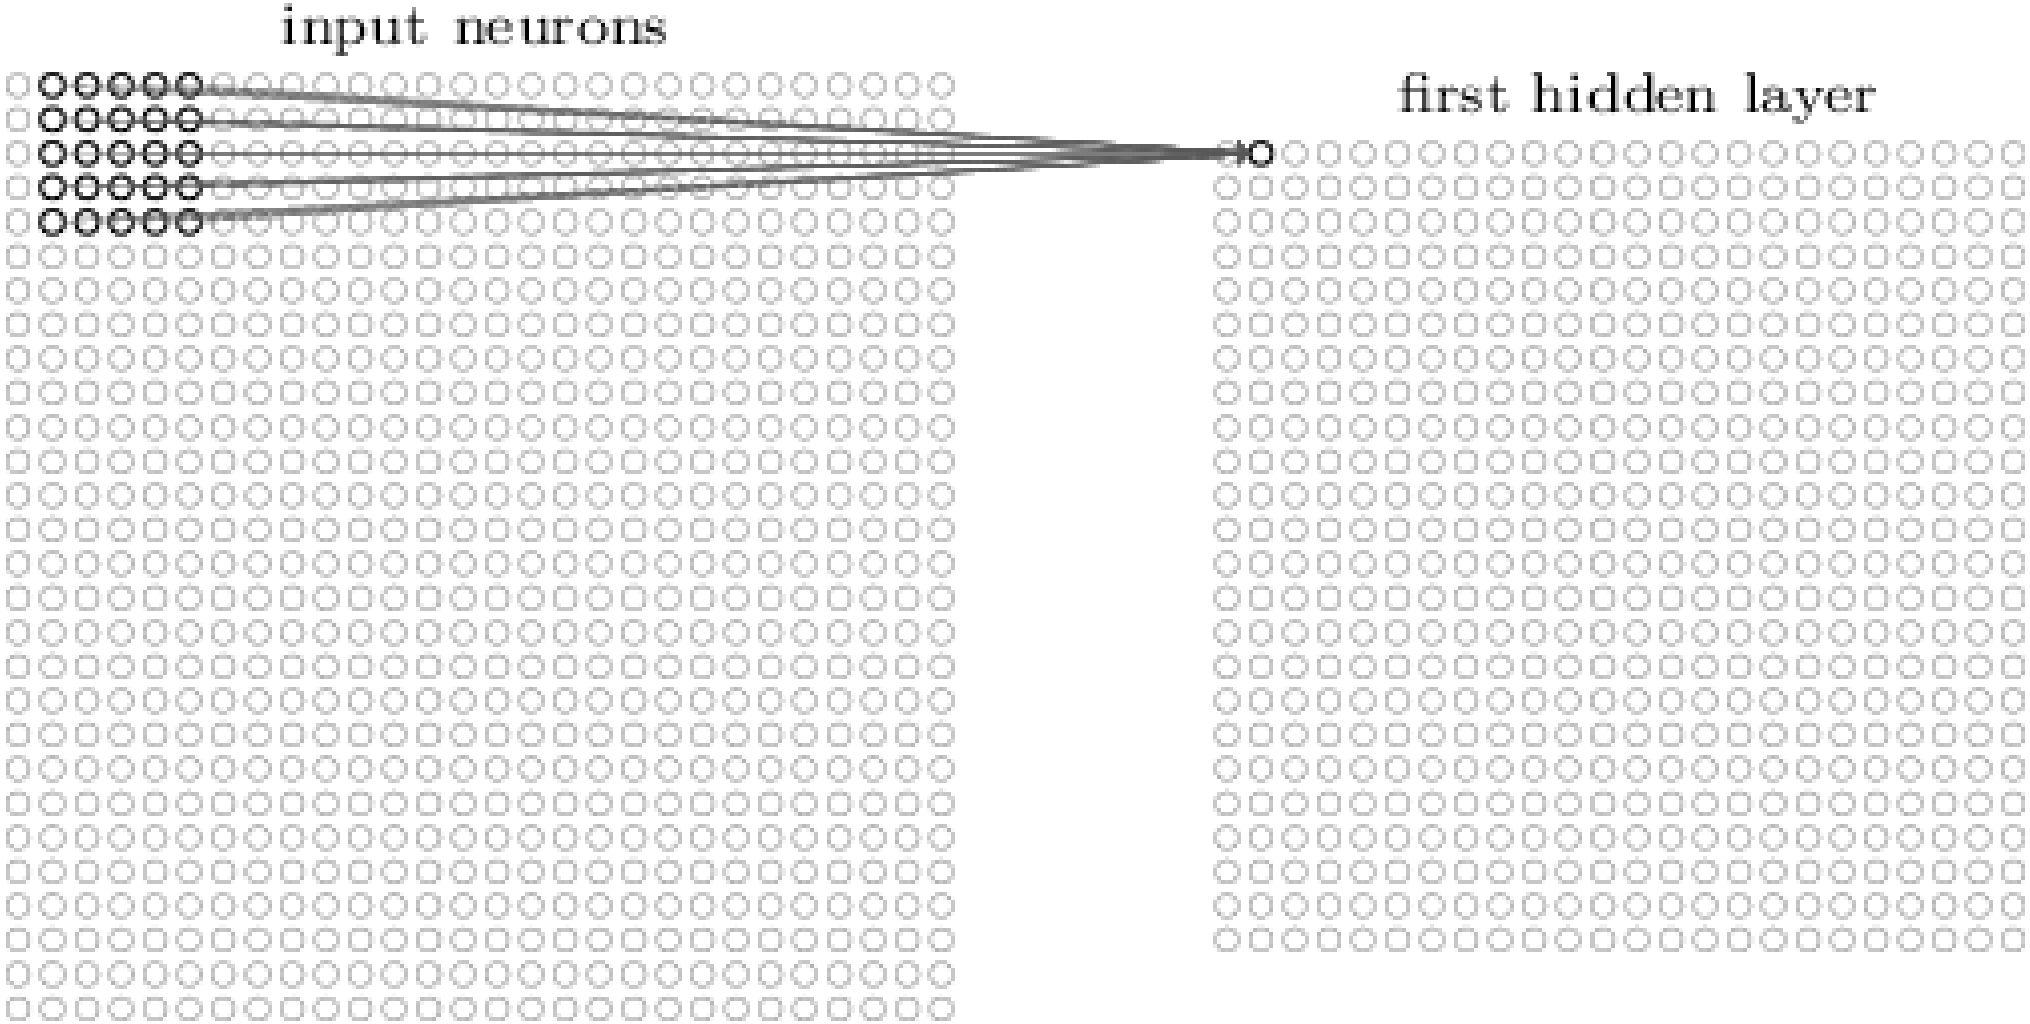

Michael A. Nielsen, "Neural Networks and Deep Learning", Determination Press, 2015
http://neuralnetworksanddeeplearning.com

Multiple Feature maps/filters

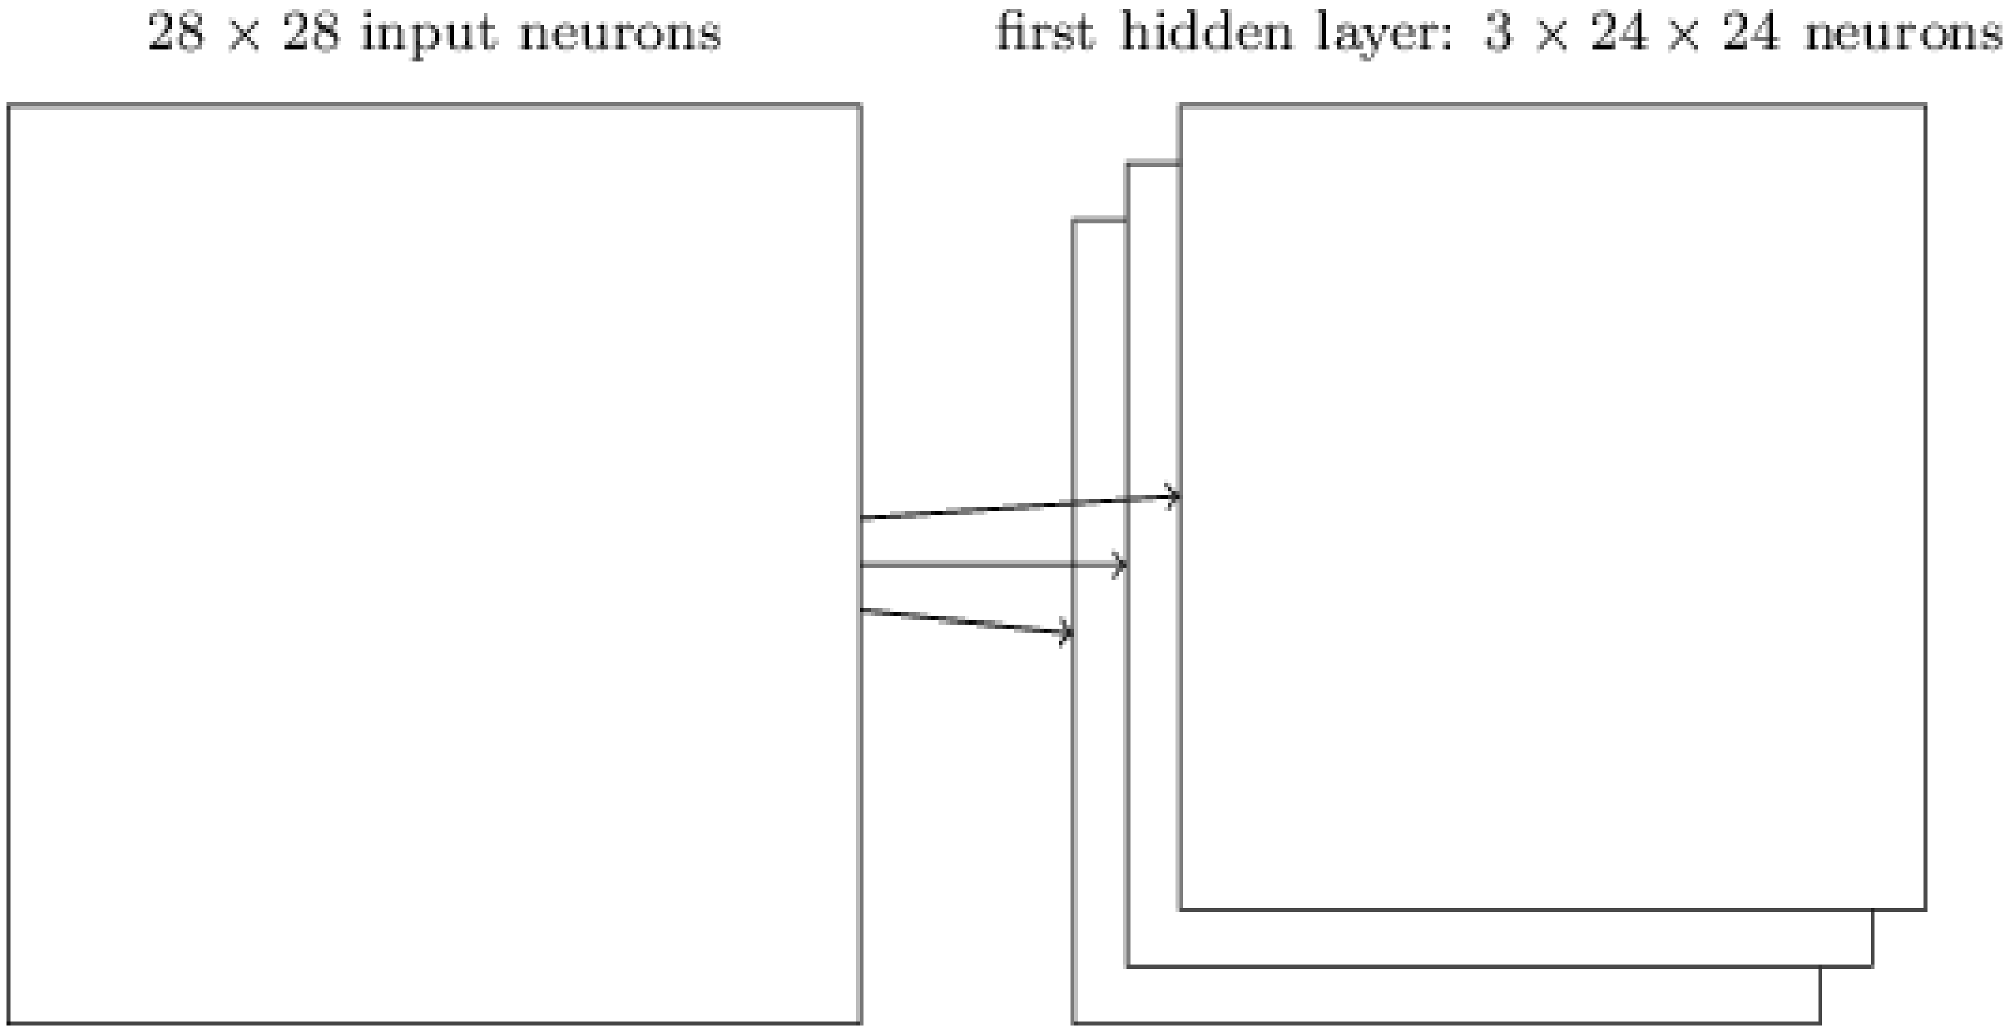

Michael A. Nielsen, "Neural Networks and Deep Learning", Determination Press, 2015
http://neuralnetworksanddeeplearning.com

Pooling

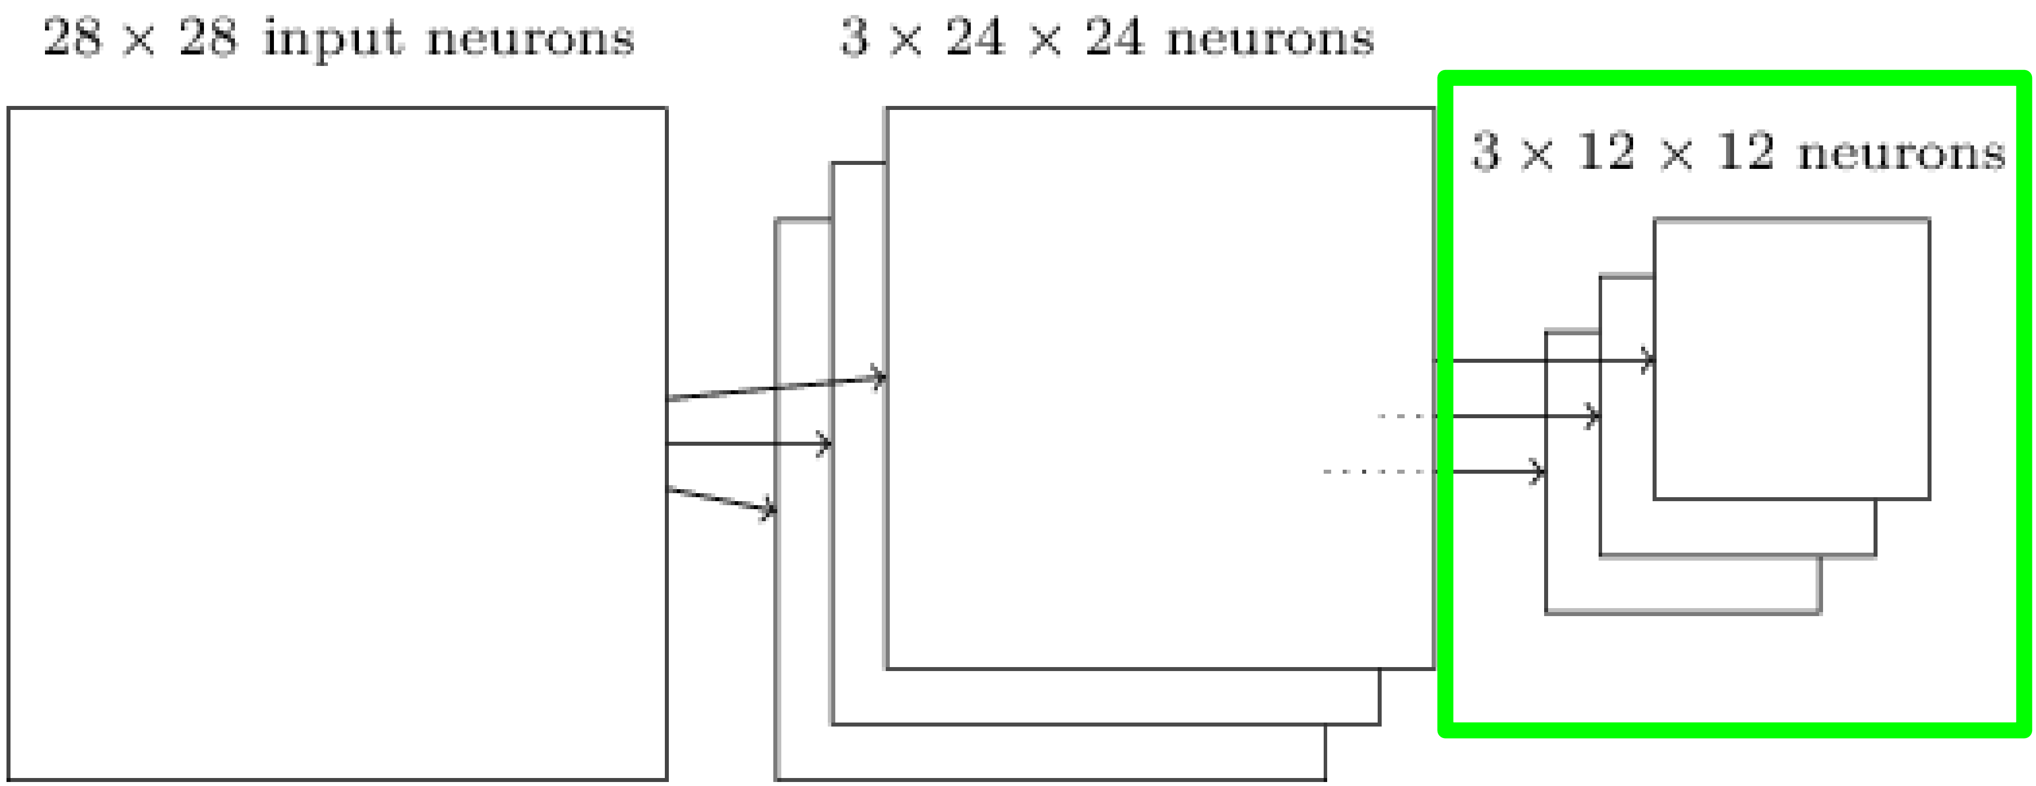

Michael A. Nielsen, "Neural Networks and Deep Learning", Determination Press, 2015
http://neuralnetworksanddeeplearning.com

Training on game data predicting the next move

12 layered DCNN

64 to 192 feature maps **per layer**

2.3 million parameters
630 million connections

Input Features

• Stone Colour	x 3
• Liberties	x 4
• Liberties after move played	x 6
• Legal Move	x 1
• Turns since	x 5
• Capture Size	x 7
• Ladder Move	x 1
• KGS Rank	x 9

55% Accuracy

Mostly beats GnuGo

Combined with MCTS

Selection

Asynchronous GPU Power

Revolution

AlphaGo

Networks in Training

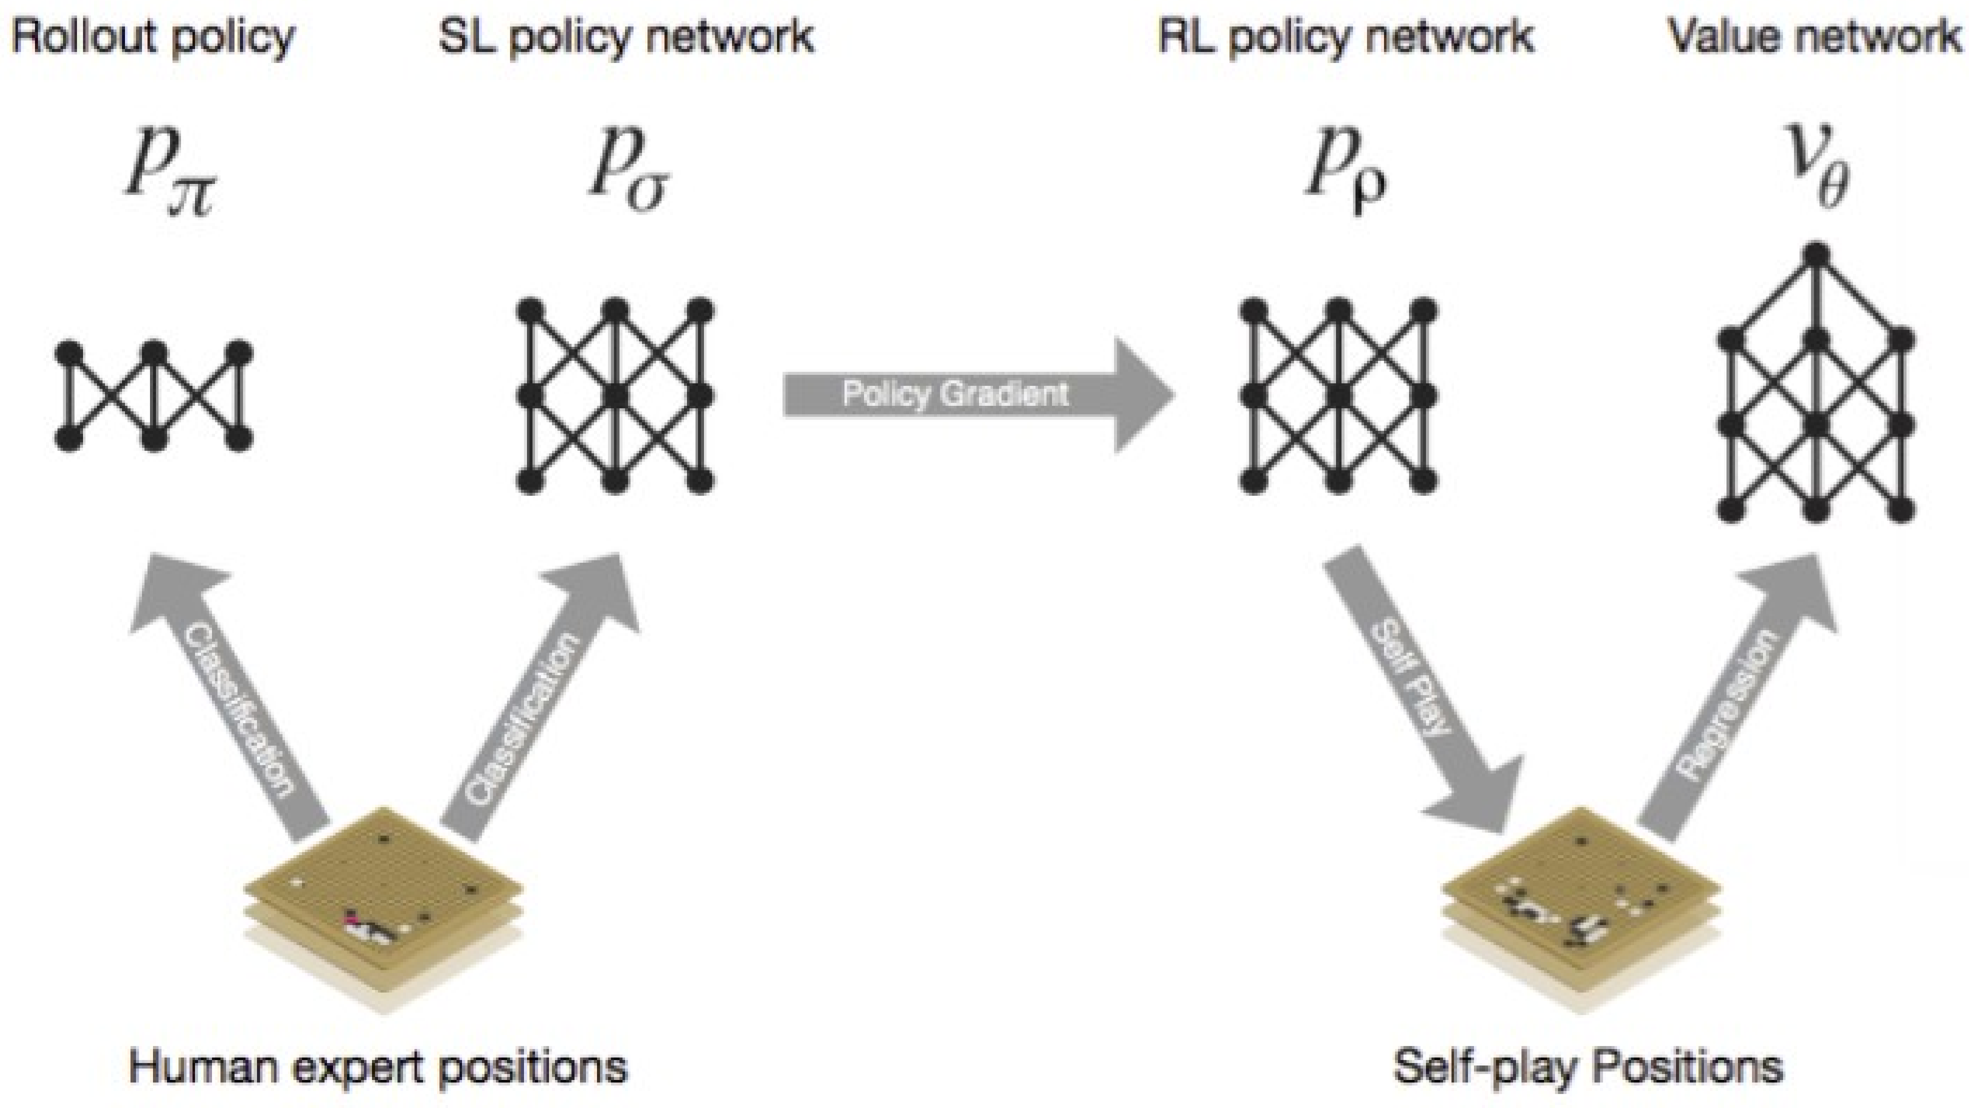

Silver, D. et al., 2016. Mastering the game of Go with deep neural networks and tree search. Nature, 529(7587), p.484-489.

AlphaGo Search

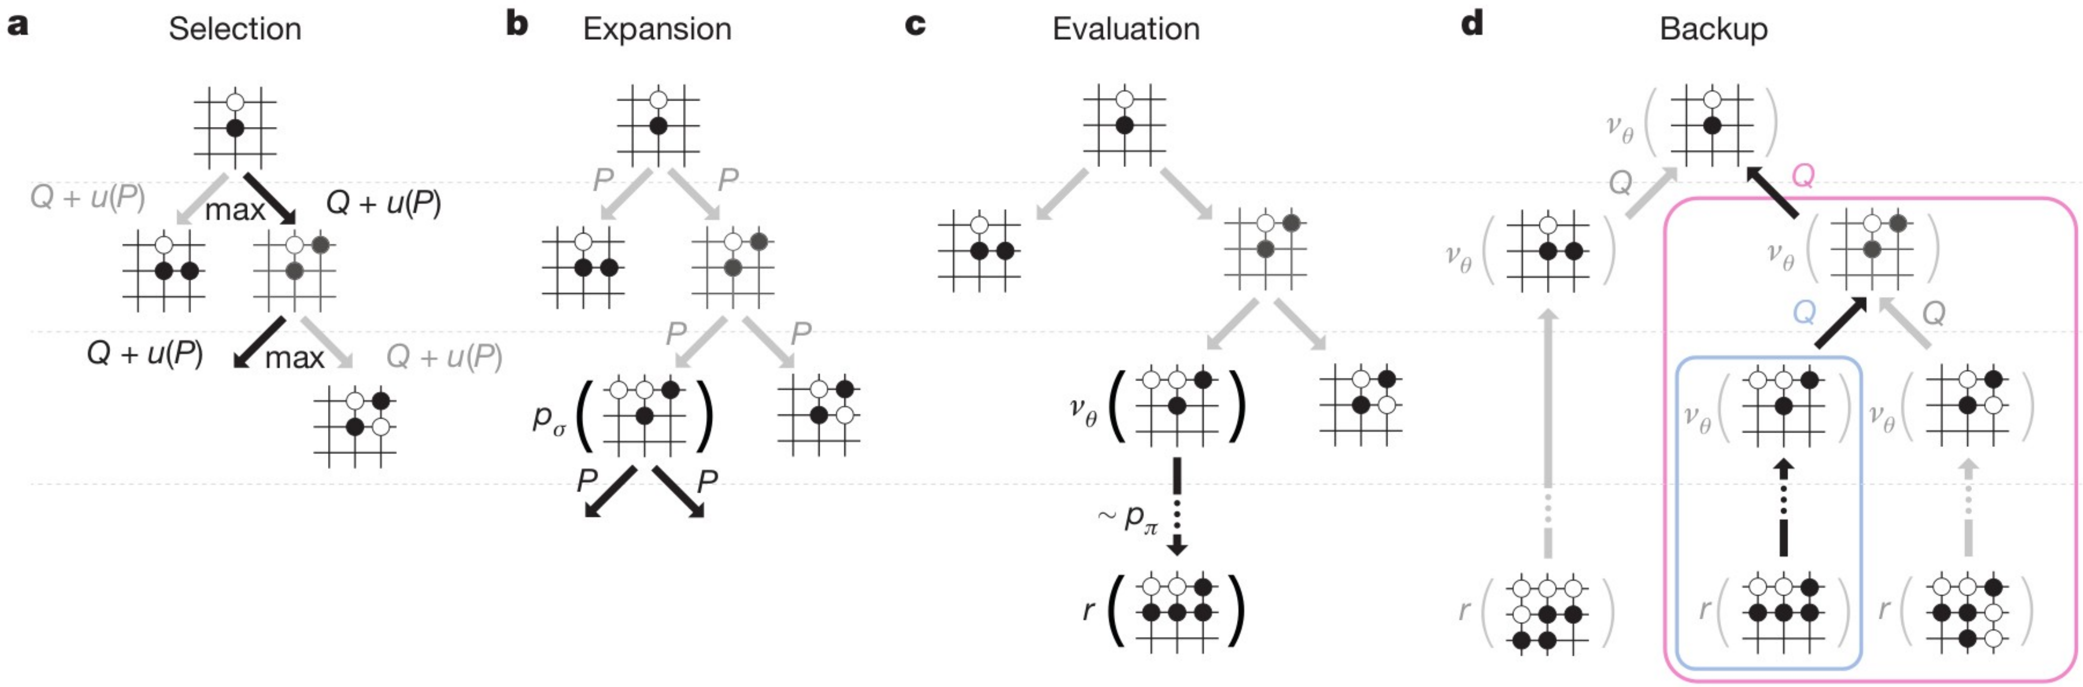

Silver, D. et al., 2016. Mastering the game of Go with deep neural networks and tree search. Nature, 529(7587), p.484-489.

1202 CPUs and 176 GPUs

Tensor PU

3 Strengths of AlphaGo

Human Instinct	Policy Network
Reading Capability	Search
Positional Judgement	Value Network

Most Important Strength

Human Instinct	Policy Network
Reading Capability	Search
Positional Judgement	Value Network


More Natural

Style

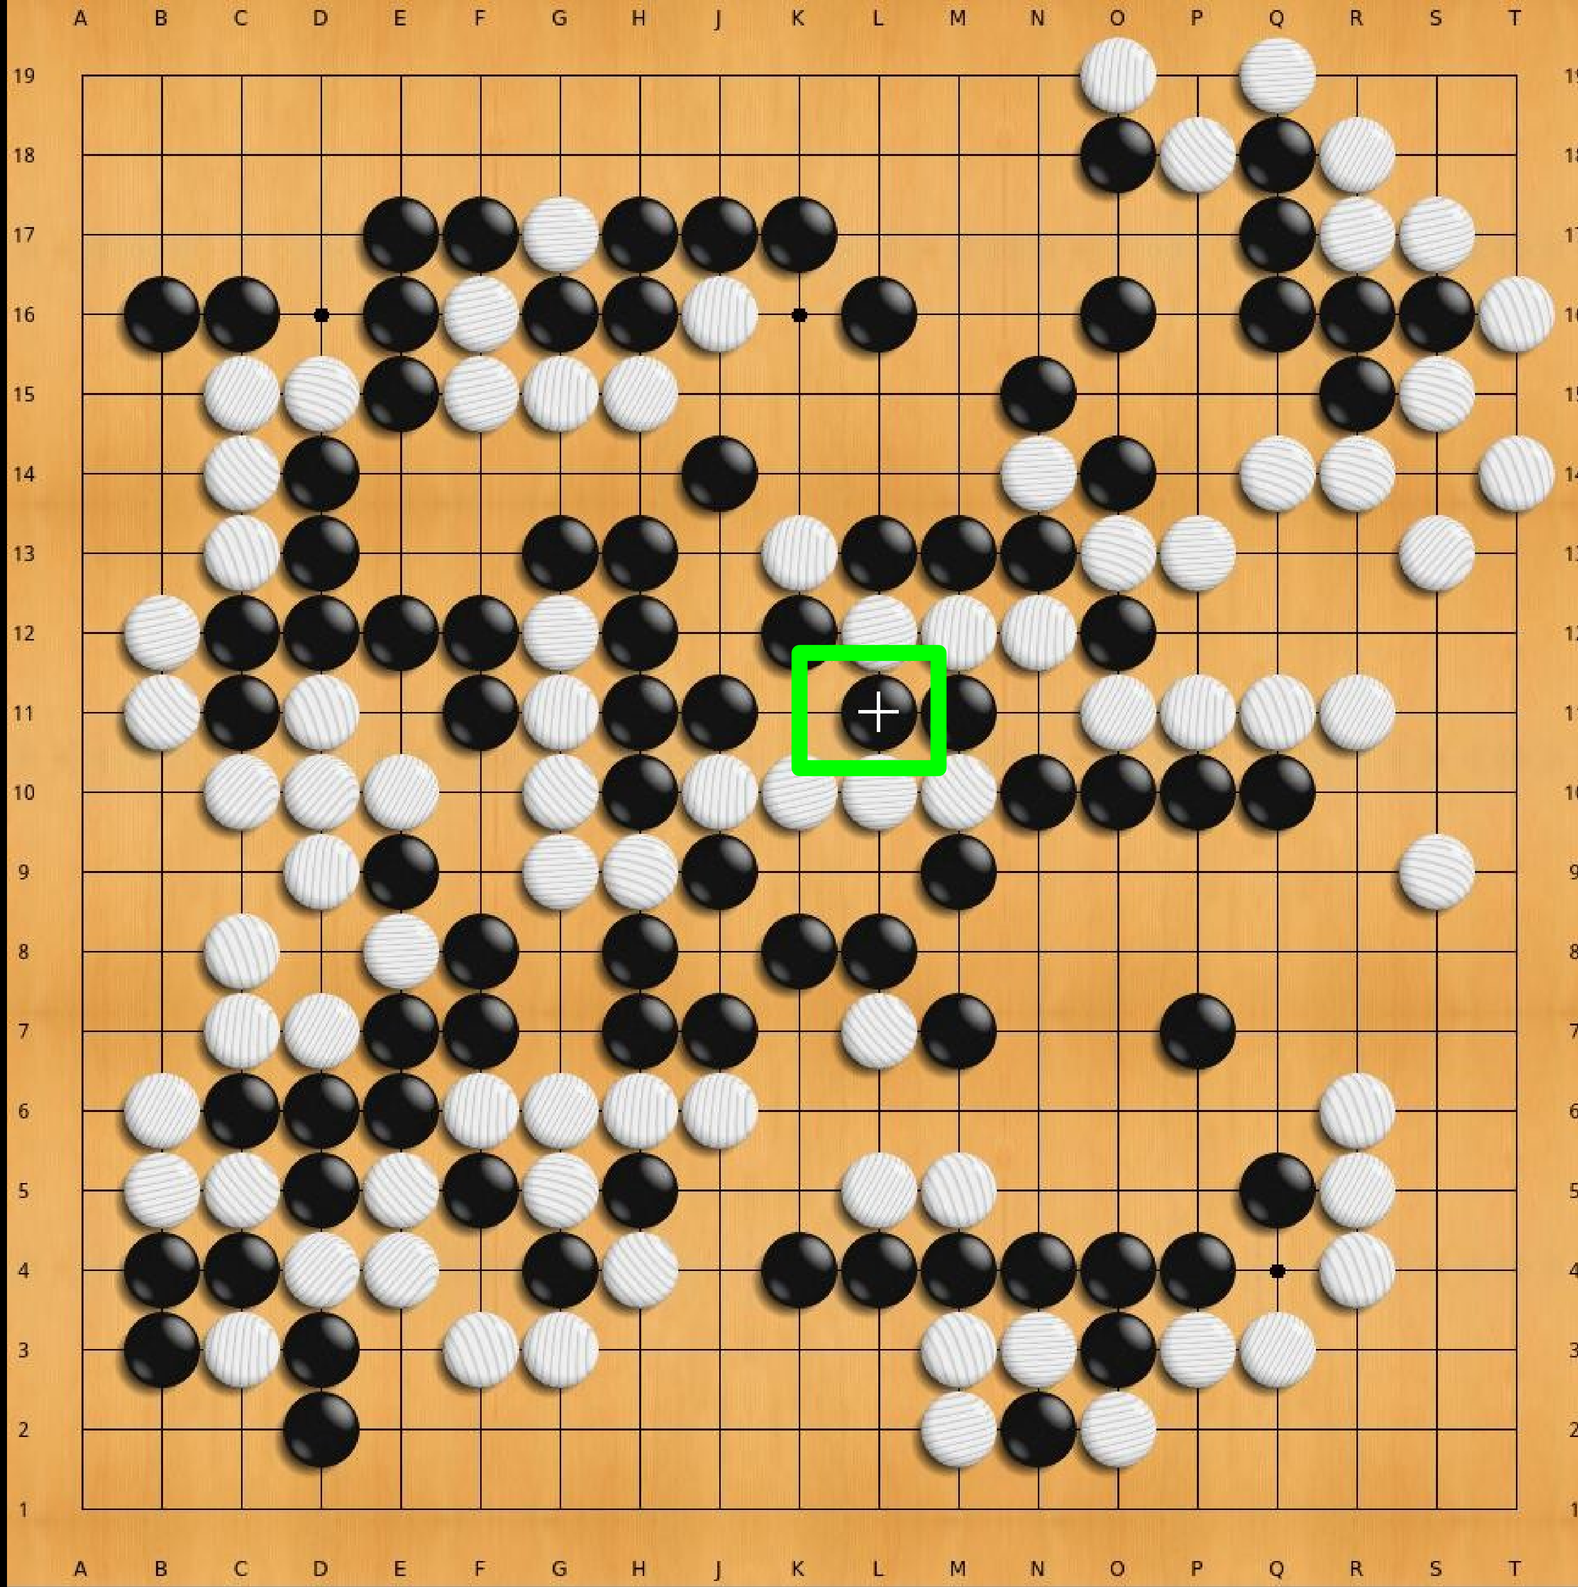

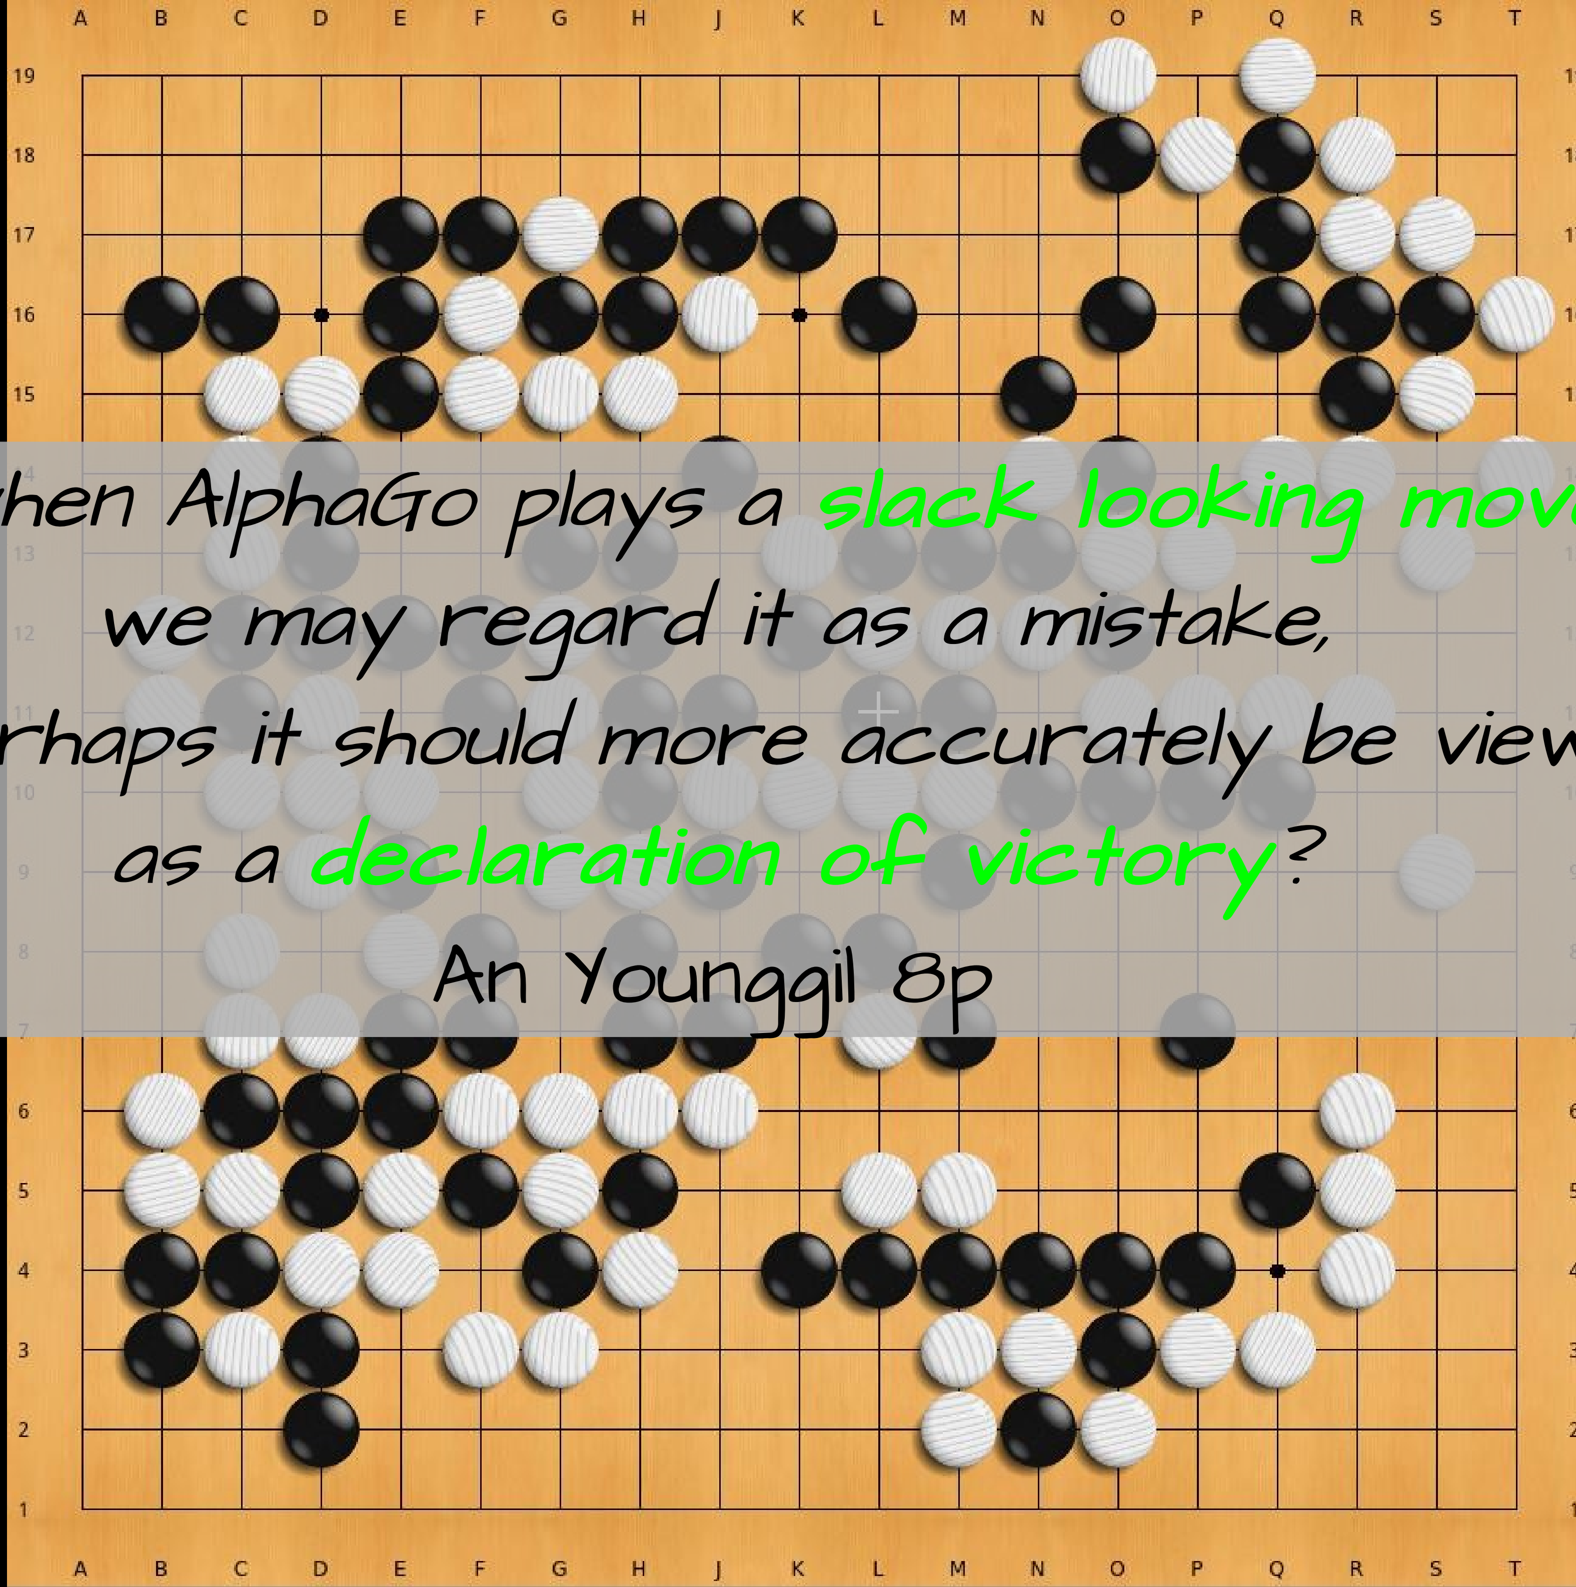

So when AlphaGo plays a slack looking move,
we may regard it as a mistake,
but perhaps it should more accurately be viewed
as a declaration of victory?

An Younggil 8p

Game 2

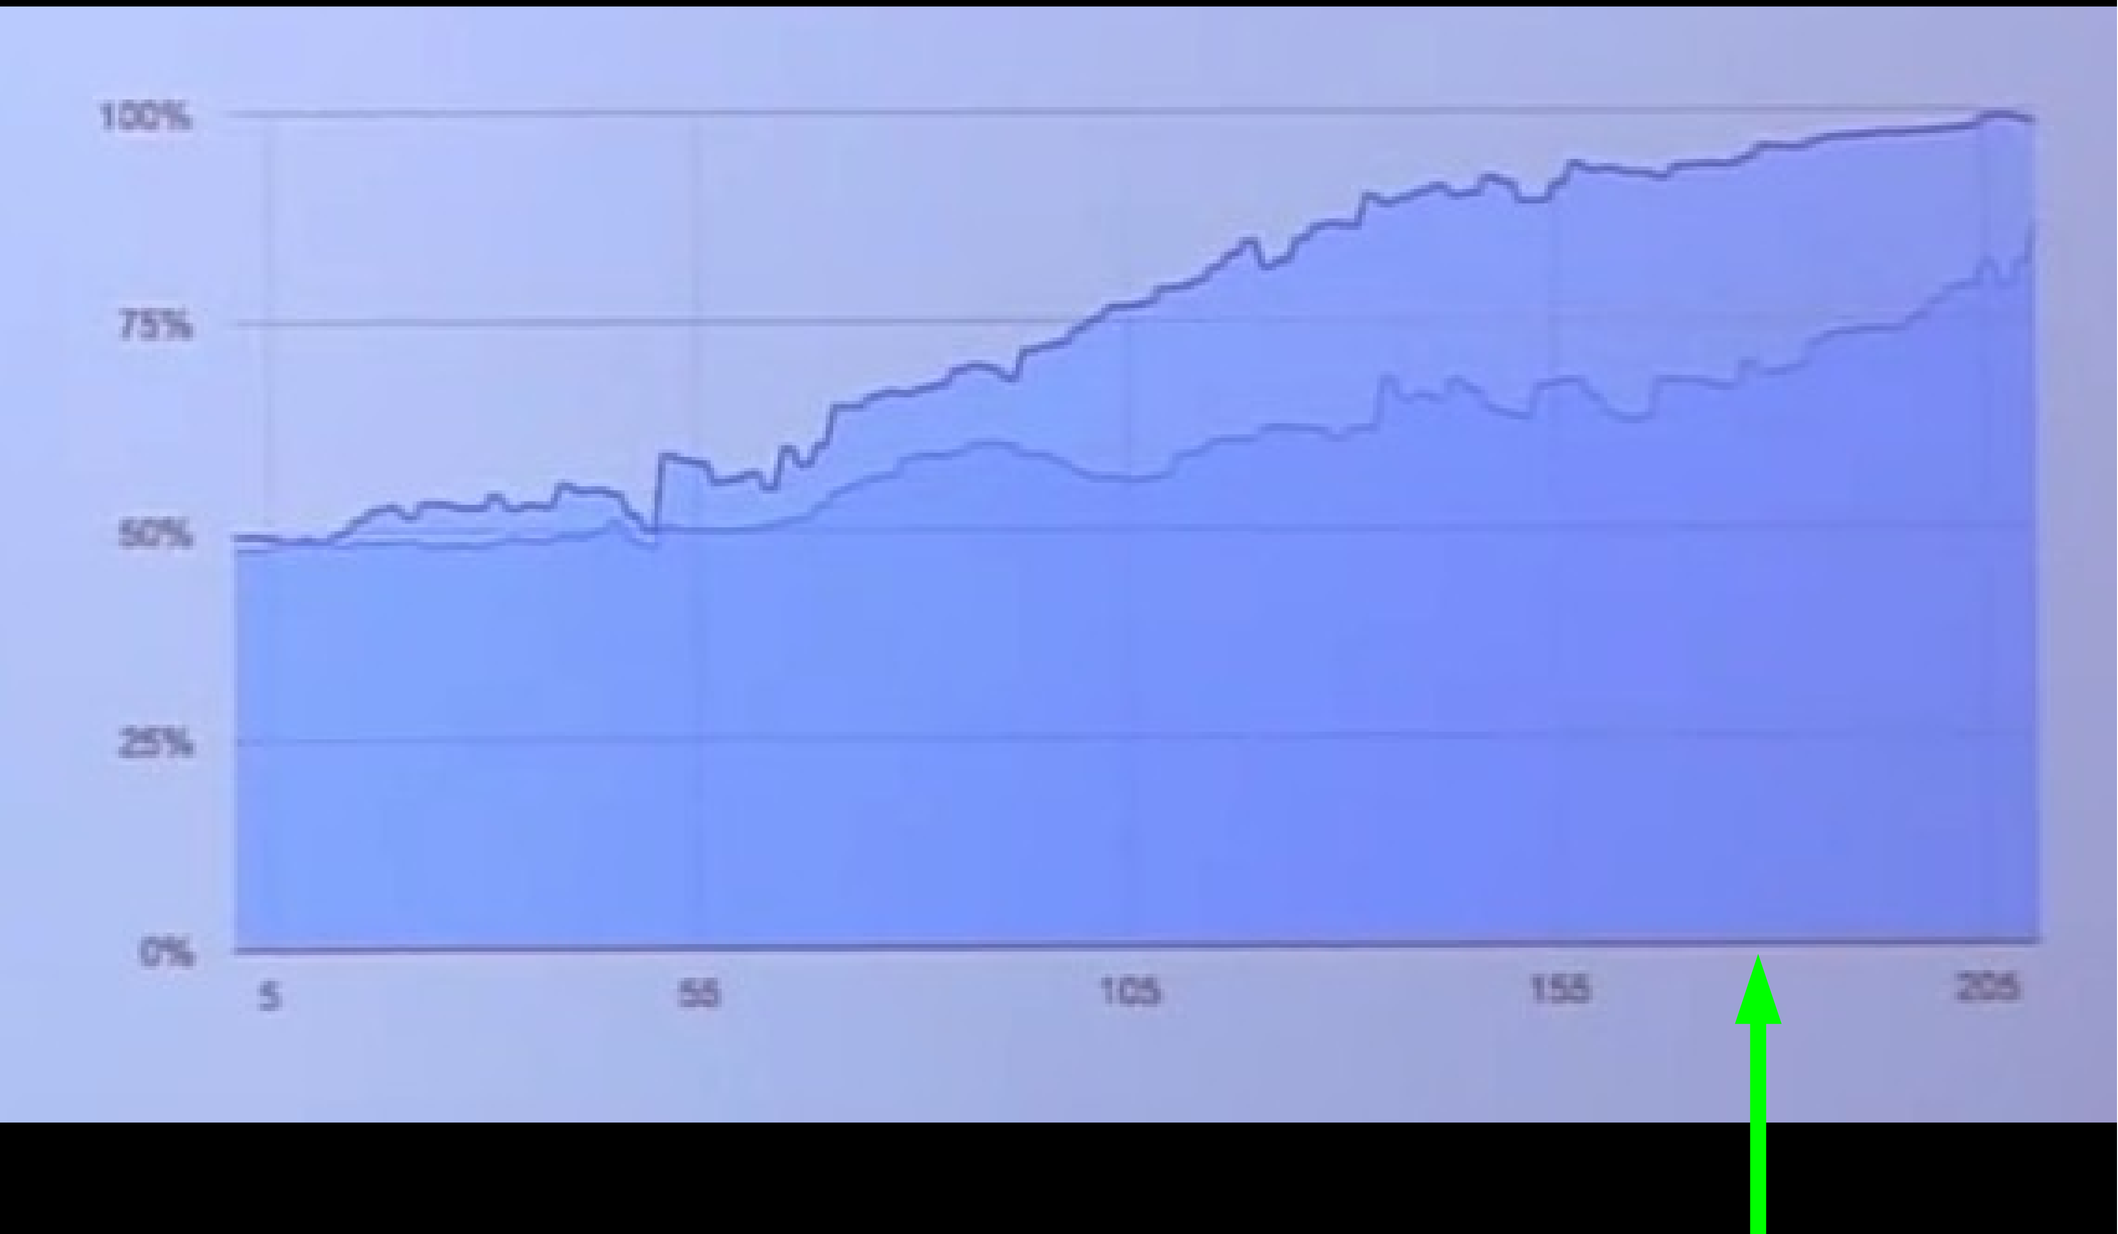

Game 4

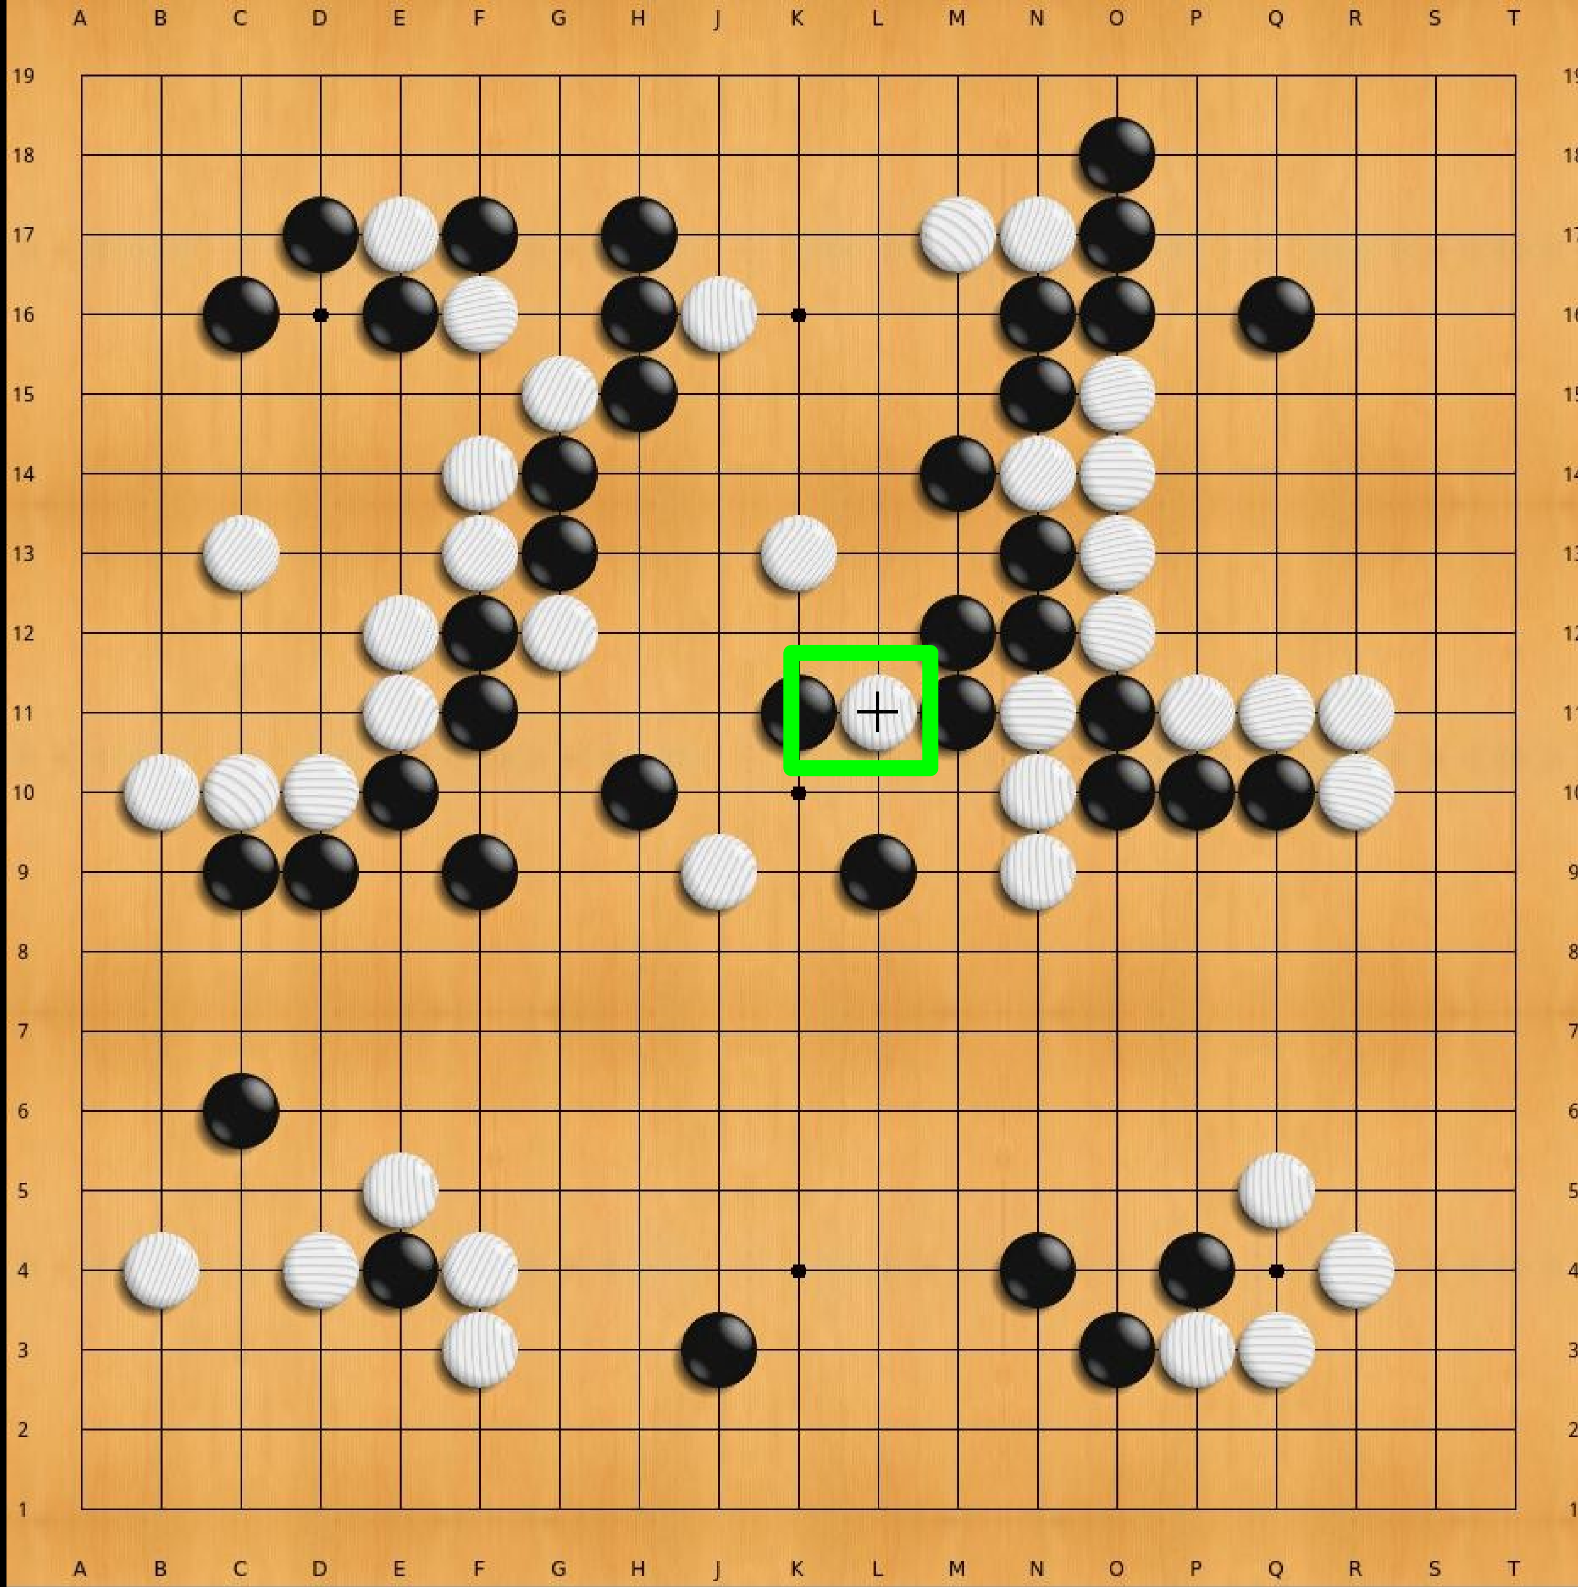
A B C D E F G H J K L M N O P Q R S T
19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

Game 4

Game 4

What can we learn?

Making X faster

vs

Doing less of X

Modularizing small components

Benchmark everything

Solving problems the human way

vs

Solving problems the computer way

Don't **blindly dismiss** approaches as infeasible

One Approach

vs

Combination of Approaches

Joy of Creation

```
white's turn to move!'
Make a move in the form XY, e.g. A19, D7 as the labels indicate!
> F8
    A B C D E F G H J
  9 . . . . . . . . .  9
  8 . X . . . O X . .  8
  7 . O O . . . O . .  7
  6 . . . . O . . . .  6
  5 . O O X . . . O .  5
  4 . X X . . . . . .  4
  3 . . X . . X . . .  3
  2 . . . . . . . X .  2
  1 . . . . . . . . .  1
    A B C D E F G H J
white played at F8
black's turn to move!'
Rubykon is thinking...
D2 => 0.4881516587677725
C8 => 0.4827586206896552
F4 => 0.4827586206896552
F5 => 0.47761194029850745
G4 => 0.47474747474747475
E8 => 0.4744897959183674
G6 => 0.47150259067357514
D6 => 0.46875
D4 => 0.4627659574468085
F7 => 0.46236559139784944
    A B C D E F G H J
  9 . . . . . . . . .  9
  8 . X . . . O X . .  8
  7 . O O . . . O . .  7
  6 . . . . O . . . .  6
  5 . O O X . . . O .  5
  4 . X X . . . . . .  4
  3 . . X . . X . . .  3
  2 . . . X . . . X .  2
  1 . . . . . . . . .  1
    A B C D E F G H J
black played at D2
white's turn to move!'
Make a move in the form XY, e.g. A19, D7 as the labels indicate!
> F5
    A B C D E F G H J
  9 . . . . . . . . .  9
  8 . X . . . O X . .  8
  7 . O O . . . O . .  7
  6 . . . . O . . . .  6
  5 . O O X . O . O .  5
  4 . X X . . . . . .  4
  3 . . X . . X . . .  3
  2 . . . X . . . X .  2
  1 . . . . . . . . .  1
    A B C D E F G H J
white played at F5
black's turn to move!'
Rubykon is thinking...
```

PragTob/Rubykon

Welcome to web-go

This is work in progress, please enjoy your game.
Counting rules are chinese /area counting.

PragTob/web-go

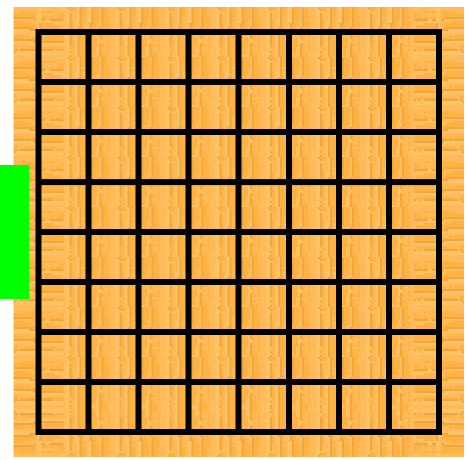

pasky / michi

Watch 18

Minimalistic Go MCTS Engine

117 commits | 2 branches | 0 releases | 5 contributors

Branch: master | michi / + | Fork this project and create a new file

pasky Merge pull request #5 from traveller42/master ... Latest commit 2f22d84 18 days

.gitignore	Keep only repository specific entries in repository .gitignore	a month
README.md	Update links to michi-c and michi-c2	a month
michi.py	Test for Pass in the GTP code was too strict	18 days

pasky/michi

Michi --- Minimalistic Go MCTS Engine

Michi aims to be a minimalistic but full-fledged Computer Go program based on state-of-art methods (Monte Carlo Tree Search) and written in Python. Our goal is to make it easier for new people to enter the domain of Computer Go, peek under the hood of a "real" playing engine and be able to learn by

What did AlphaGo do to beat the strongest human Go player?

Tobias Pfeiffer

@PragTob

pragtob.info

bitcrowd

Sources

- Maddison, C.J. et al., 2014. Move Evaluation in Go Using Deep Convolutional Neural Networks.
- Silver, D. et al., 2016. Mastering the game of Go with deep neural networks and tree search. Nature, 529(7587), p.484-489.
- Michael A. Nielsen, "Neural Networks and Deep Learning", Determination Press, 2015 http://neuralnetworksanddeeplearning.com
- Gelly, S. & Silver, D., 2011. Monte-Carlo tree search and rapid action value estimation in computer Go. Artificial Intelligence, 175(11), p.1856-1876.
- I. Althöfer, “On the Laziness of Monte-Carlo Game Tree Search In Non-tight Situations,” Friedrich-Schiller Univ., Jena, Tech. Rep., 2008.
- Browne, C. & Powley, E., 2012. A survey of monte carlo tree search methods. IEEE Transactions on Intelligence and AI in Games, 4(1), p.1-49.
- Gelly, S. & Silver, D., 2007. Combining online and offline knowledge in UCT. Machine Learning, p.273-280.
- https://www.youtube.com/watch?v=LX8Knl0g0LE&index=9&list=WL

Photo Credit

- http://www.computer-go.info/events/ing/2000/images/bigcup.jpg
- https://en.wikipedia.org/wiki/File:Kasparov-29.jpg
- http://www.geforce.com/hardware/desktop-gpus/geforce-gtx-titan-black/product-images
- http://giphy.com/gifs/dark-thread-after-lCP95tGSbMmWI
- https://cloudplatform.googleblog.com/2016/05/Google-supercharges-machine-learning-tasks-with-custom-chip.html
- https://gogameguru.com/i/2016/01/Fan-Hui-vs-AlphaGo-550x364.jpg
- CC BY 2.0
  - https://en.wikipedia.org/wiki/File:Deep_Blue.jpg
  - https://www.flickr.com/photos/luisbg/2094497611/
- CC BY-SA 3.0
  - https://en.wikipedia.org/wiki/Alpha%E2%80%93beta_pruning#/media/File:AB_pruning.svg
- CC BY-SA 2.0
  - https://flic.kr/p/cPUtny
  - https://flic.kr/p/dLSKTQ
  - https://www.flickr.com/photos/83633410@N07/7658272558/

Photo Credit

- CC BY-NC-ND 2.0
  - https://flic.kr/p/q15pzb
  - https://flic.kr/p/bHSj7D
  - https://flic.kr/p/ixSsfM
  - https://www.flickr.com/photos/waxorian/4228645447/
  - https://www.flickr.com/photos/pennstatelive/8972110324/
  - https://www.flickr.com/photos/dylanstraub/6428496139/
- https://en.wikipedia.org/wiki/Alphabet_Inc.#/media/File:Alphabet_Inc_Logo_2015.svg
- CC BY 3.0
  - https://en.wikipedia.org/wiki/File:Pi_30K.gif