

UCT-Enhanced Deep Convolutional Neural Networks For Move Recommendation in Go

Sarun Paisarnsrisomsuk
Pitchaya Wiratchotisation




MTA SZTAKI
COMPUTER AND AUTOMATION
RESEARCH INSTITUTE
HUNGARIAN ACADEMY OF SCIENCES

Outline

- The Game of Go
- Goal
- Deep Convolutional Neural Networks (Deep CNN)
- Upper Confidence Bounds applied to tree (UCT)
- Methodology
- Experiments and Results
- Conclusion and Future Work


Go: An Ultimate Challenge for AI

- An ancient board game
- Two-player
- Zero-sum
- Deterministic
- Perfect-information



Simple Rules, but Complex Strategies

- Place stones in turn on a 19x19 board
- Basic goal: secure more territories than the opponent
- Enormous combinatorial complexity
- Long-term influence of a move




Goal

To enhance move recommendation in Go
using DCNN and UCT


Deep Convolutional Neural Networks (deep CNN)



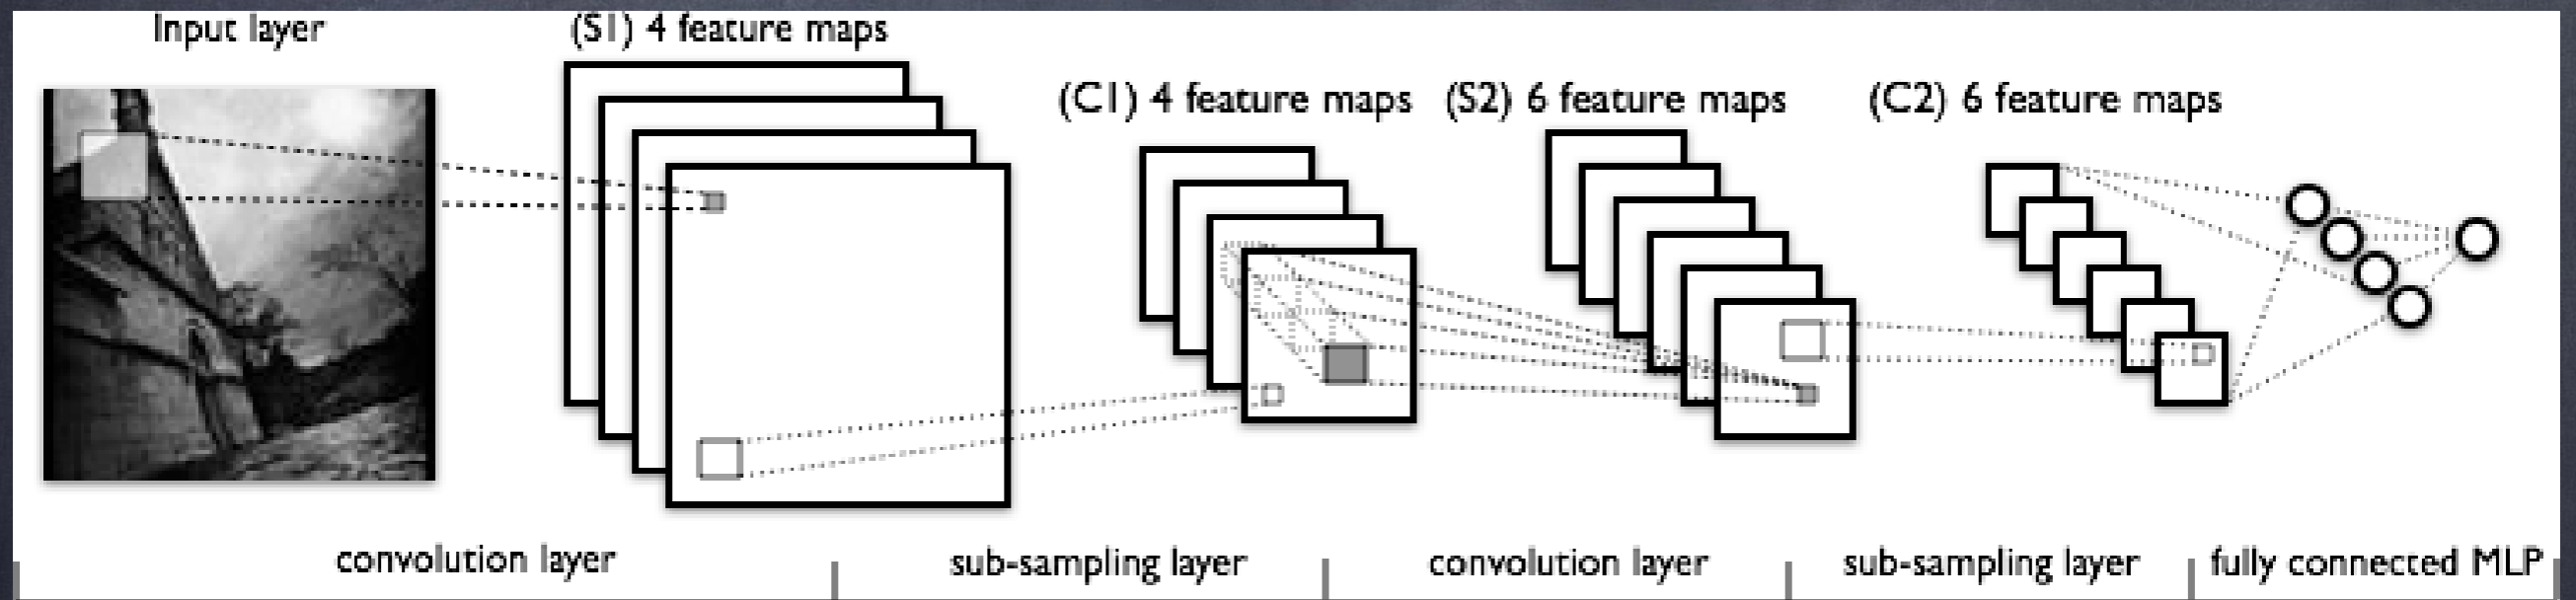
Neural Network




Neural Network



Deep Convolutional Neural Networks




Deep Convolutional Neural Networks

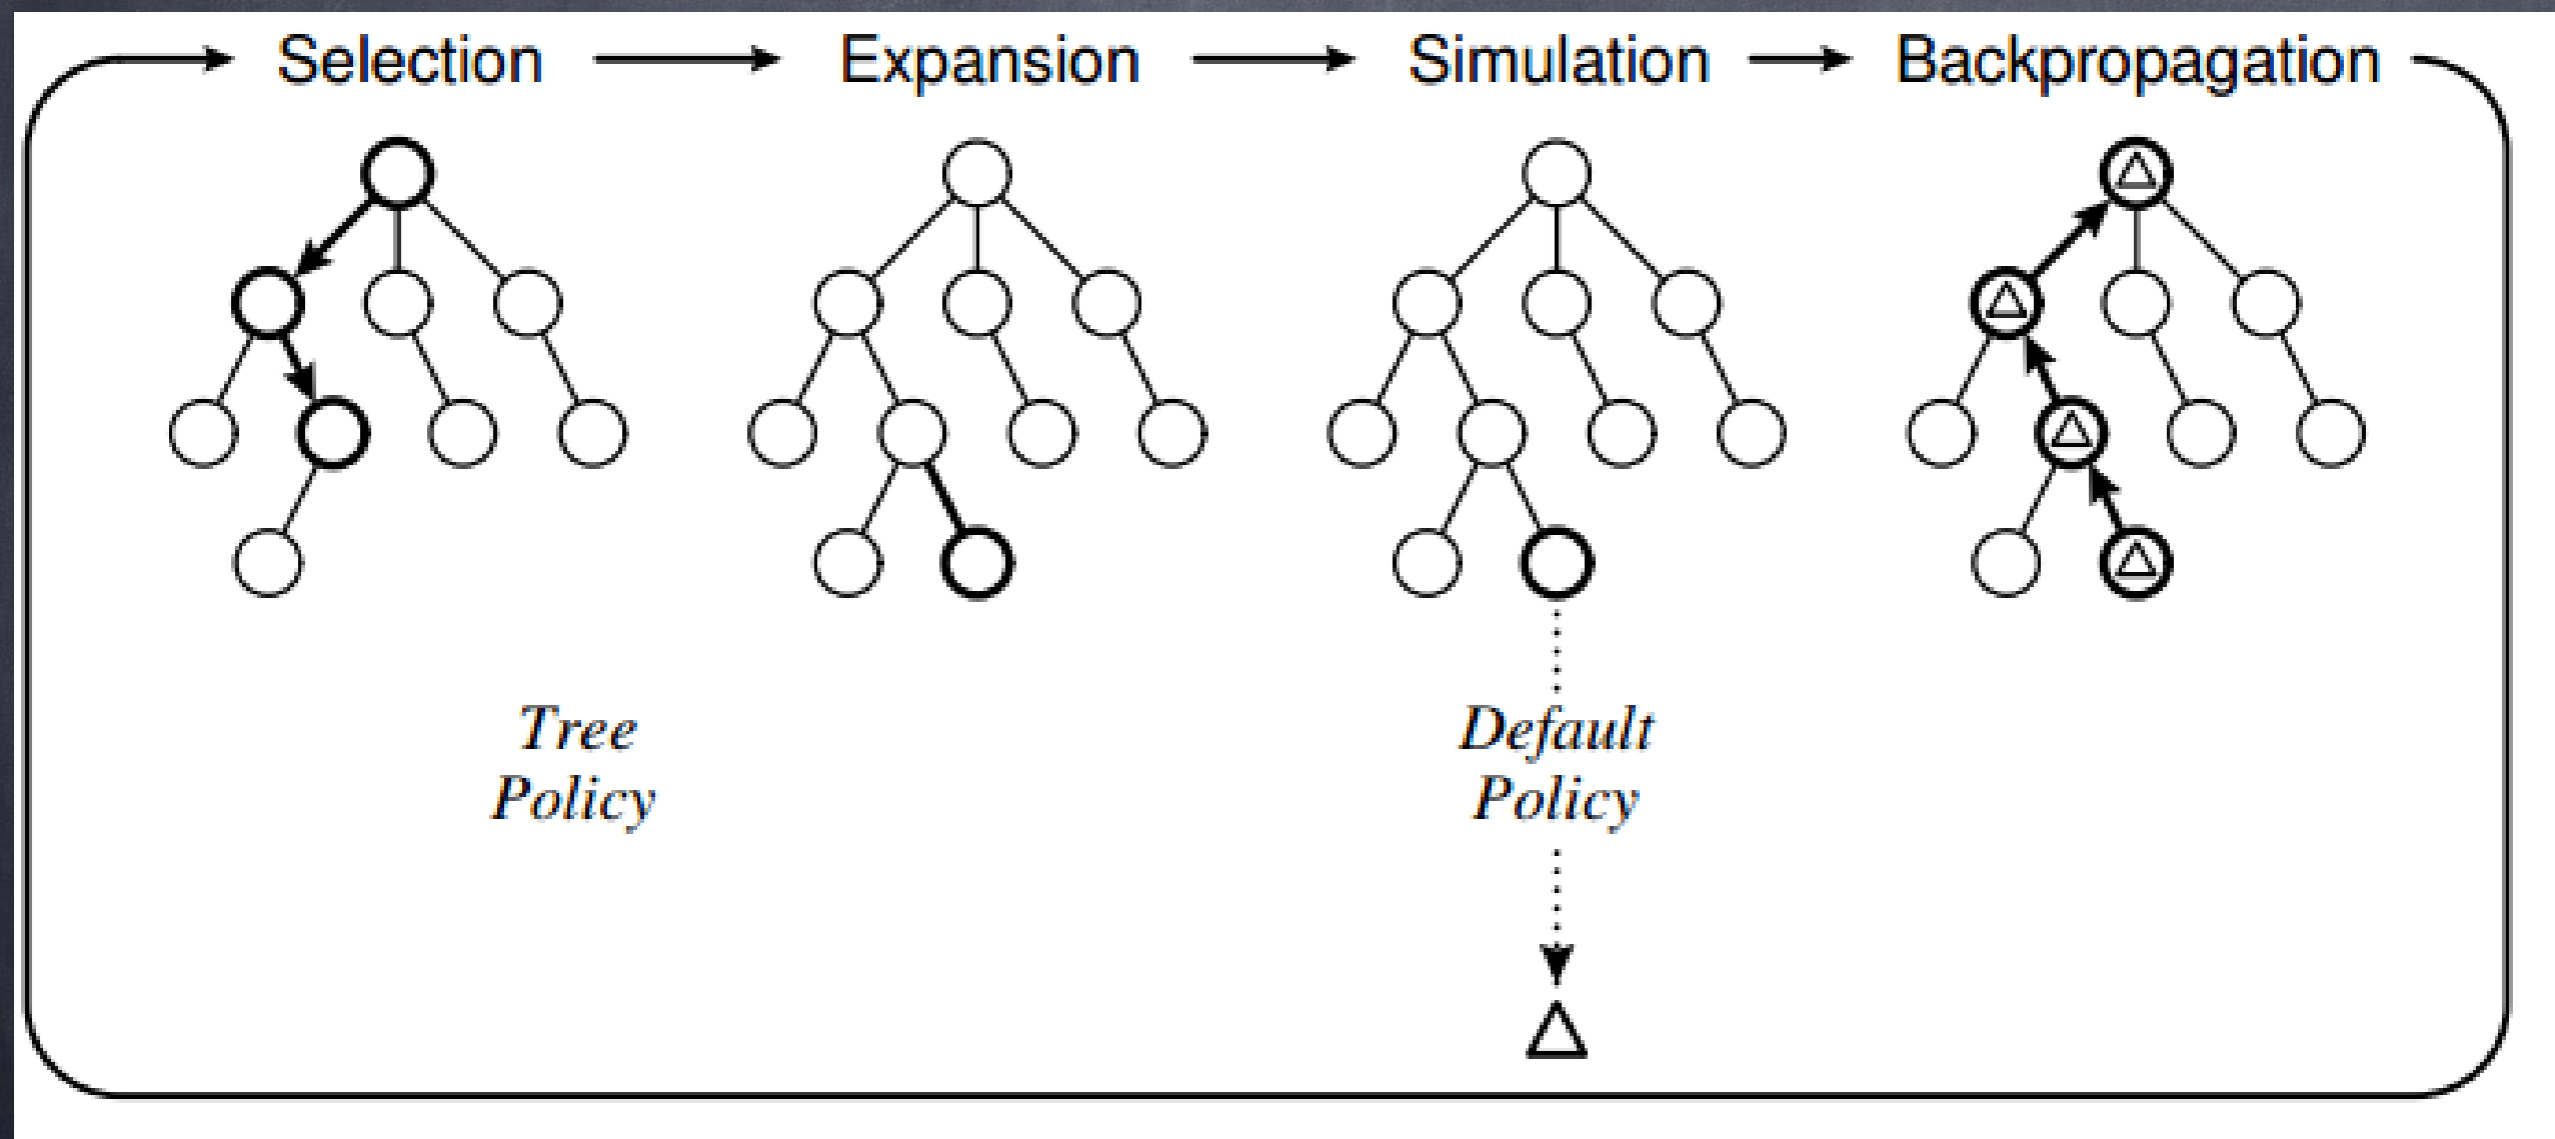


Upper Confident Bounds applied to tree (UCT)


Applied bandit-based method
to guide Monte-Carlo planning



Monte Carlo Tree Search



Upper Confident Bounds applied to tree (UCT)



Selection Policy:

$$UCT = \bar{X}_j + 2C_p \sqrt{\frac{2 \ln n}{n_j}}$$

Exploitation

Exploration

Choose node which has maximum value of UCT


Methodology



Go Data




- 170,000 complete games from the KGS Server (Kiseido Go database)
- Extracted features from each individual move

Feature	Planes	Description
Black / white / empty	3	Stone colour
Liberties	4	Number of liberties (empty adjacent points)
Liberties after move	6	Number of liberties after this move is played
Legality	1	Whether point is legal for current player
Turns since	5	How many turns since a move was played
Capture size	7	How many opponent stones would be captured
Ladder move	1	Whether a move at this point is a successful ladder capture
KGS rank	0	Rank of current player




New Feature

- Board pattern at the end of game, or “final board pattern”





Deep Convolutional Neural Network

- We implemented a small deep CNN
 - 1 hidden layer; no pooling
 - 10 kernels



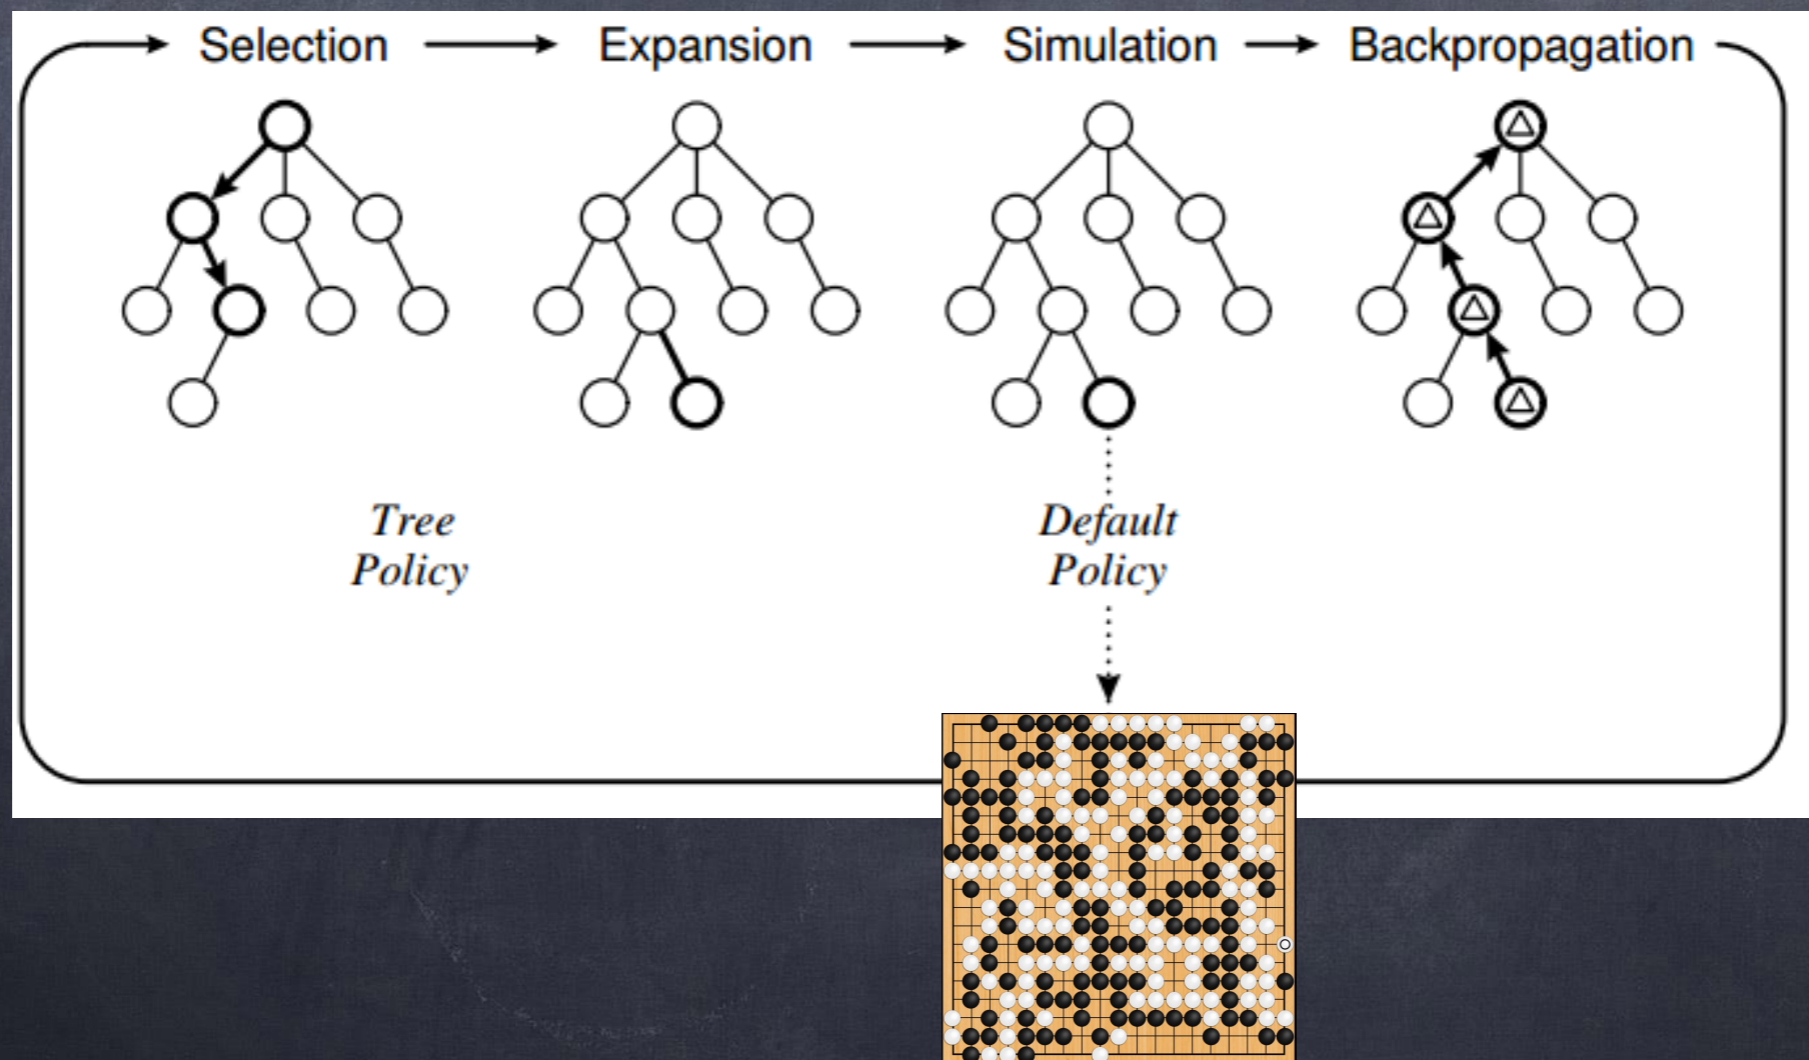
Result from Adding Final Board Pattern to input

- Without final board pattern, accuracy 6%
- With final board pattern, accuracy 18%
- The actual final board pattern is not possible to be obtained



UCT-Simulated Final Board Pattern

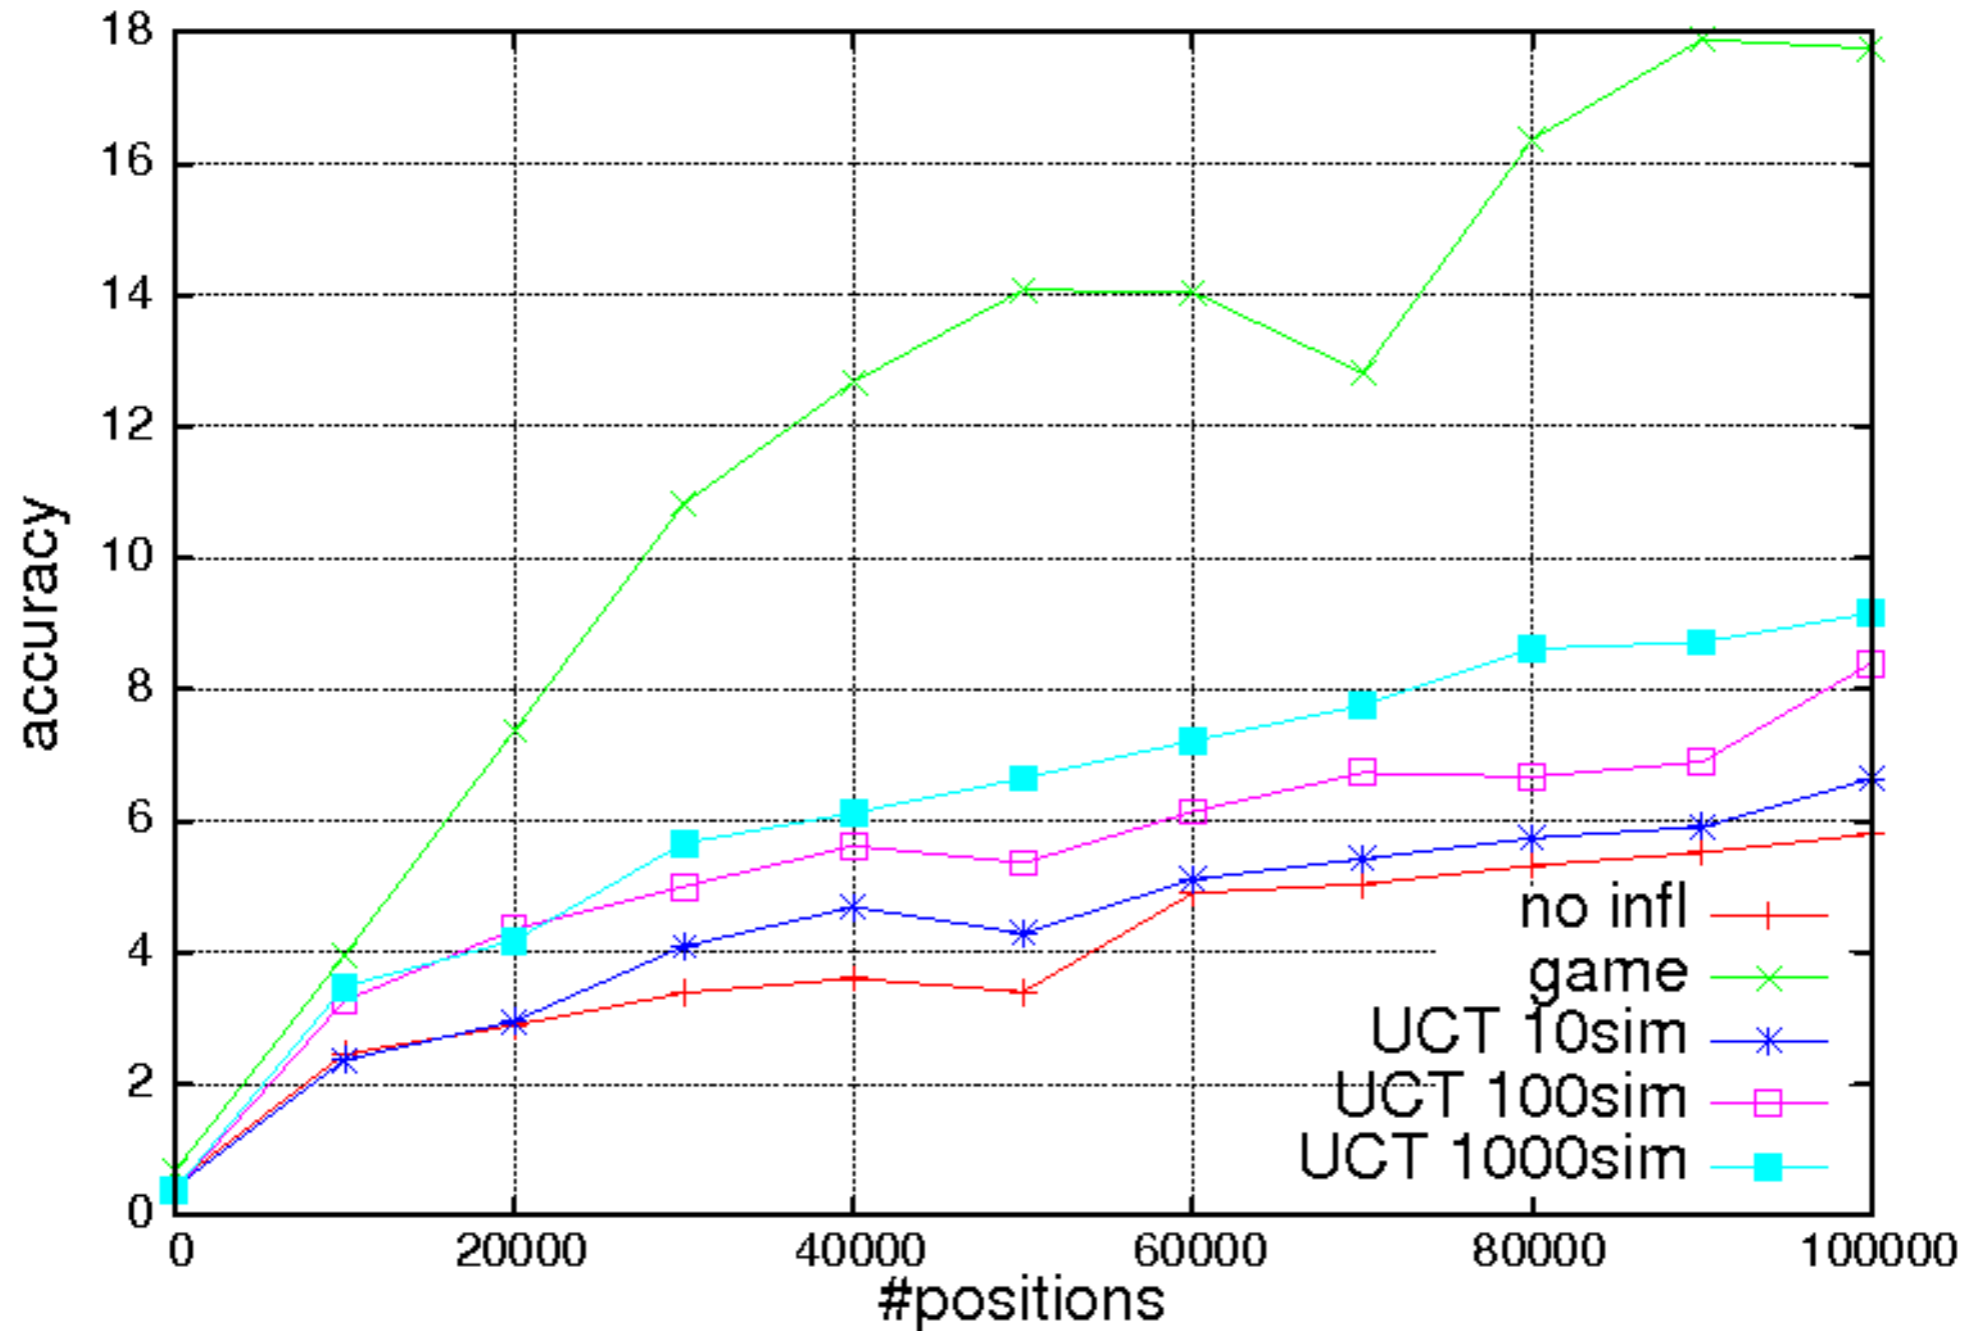
- Collect final board pattern in each simulation during UCT




Experiments and Results



Result



Result



Conclusion and Future Work



Conclusion

- Adding final board pattern to inputs of the deep CNN improves the accuracy of the network.
- In practice, we can collect statistics in each simulation of UCT to approximate final board pattern.

Future Work

- Deep network's size
- More Training Data
- Additional Features
- Combining deep network with UCT



Acknowledgements

- Levente Kocsis, Project Advisor and SZTAKI liaison
- Gabor Sarkozy, MQP Advisor
- Worcester Polytechnic Institute
- MTA-SZTAKI
- SZTAKI Colleagues
- Pachi and Gnu Go Development Team



03.12.2

Questions?





Köszönöm szépen!