

UCB, Monte Carlo Tree Search, AlphaZero

Milan Straka

■ April 30, 2024

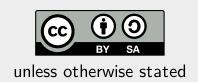








Charles University in Prague Faculty of Mathematics and Physics Institute of Formal and Applied Linguistics





Upper Confidence Bound

Upper Confidence Bound



Revisiting multi-armed bandits with ε -greedy exploration, we note that using the same epsilon for all actions in ε -greedy method seems inefficient.

One possible improvement is to select action according to *upper confidence bound* (instead of choosing a random action with probability ε):

$$A_{t+1} \stackrel{ ext{def}}{=} rg \max_{a} \Bigg[egin{array}{c} Q_t(a) &+ c \sqrt{rac{\ln t}{N_t(a)}} \Bigg], \ & ext{exploitation} \Bigg]$$

where:

- t is the number of times any action has been taken;
- ullet $N_t(a)$ is the number of times the action a has been taken before step t;
- if $N_t(a) = 0$, the right expression is frequently assumed to have a value of ∞ .

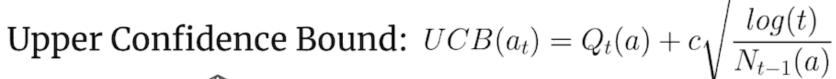
The updates are then performed as before (e.g., using averaging or fixed learning rate α).

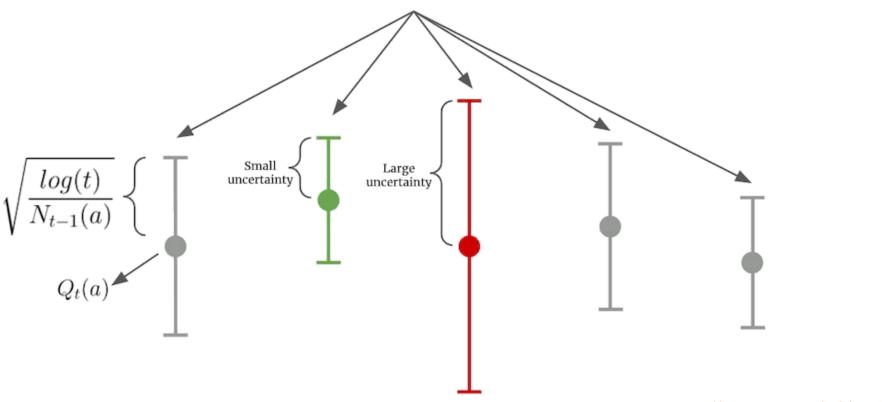
UCB

Alpha*

Upper Confidence Bound







https://miro.medium.com/v2/1*DXdfNLuFDXiKhCkJxtNaow.png

Upper Confidence Bound Derivation



We want to select the upper confidence bound, so that

- the probability of the real expected value being less or equal to the confidence bound gradually converges to 1;
- it is as small as possible.

Assuming that random variables X_i are bounded by [0,1] and denoting $\bar{X}=\frac{1}{N}\sum_{i=1}^N X_i$, (Chernoff-)Hoeffding's inequality states that

$$Pig(ar{X} \leq \mathbb{E}[ar{X}] - \deltaig) \leq e^{-2N\delta^2}, \ ext{ which we can rearrange into } \ Pig(ar{X} + \delta \leq \mathbb{E}[ar{X}]ig).$$

Our goal is to choose δ such that for every action,

$$Pig(Q_t(a) + \delta \leq q_*(a)ig) \leq \left(rac{1}{t}
ight)^lpha.$$

The required inequality will hold if $e^{-2N_t(a)\delta^2} \leq \left(\frac{1}{t}\right)^{\alpha}$, yielding $\delta \geq \sqrt{\alpha/2} \cdot \sqrt{(\ln t)/N_t(a)}$.

NPFL139, Lecture 11

Alpha*

Asymptotic Optimality of UCB



We define **regret** as the difference of maximum of what we could get (i.e., repeatedly using the action with maximum expectation) and what a strategy yields, i.e.,

$$\mathrm{regret}_N \stackrel{ ext{ iny def}}{=} N \max_a q_*(a) - \sum_{i=1}^N \mathbb{E}[R_i].$$

It can be shown that regret of UCB is asymptotically optimal, see Lai and Robbins (1985), Asymptotically Efficient Adaptive Allocation Rules; or the Chapter 8 of the 2018 Bandit Algorithms Book available online at https://banditalgs.com/.

NPFL139, Lecture 11 UCB MCTS Alpha* AlphaZero A0-MCTS A0-Network A0-Training A0-Evaluation 6/39

Upper Confidence Bound Multi-armed Bandits Results



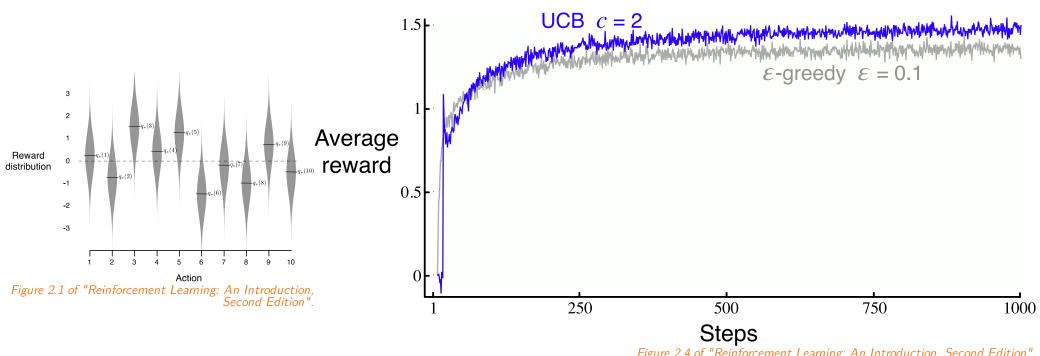
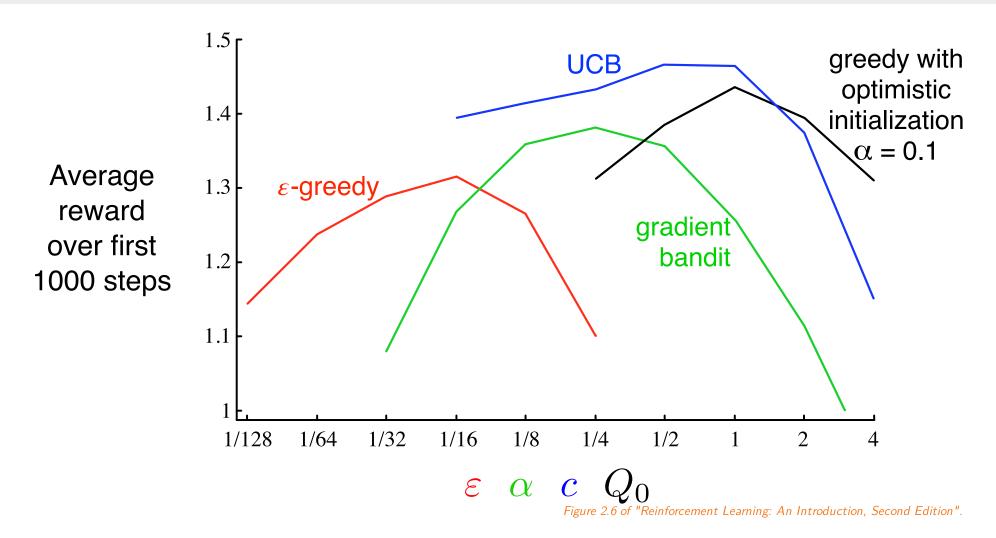


Figure 2.4 of "Reinforcement Learning: An Introduction, Second Edition".

Multi-armed Bandits Comparison





NPFL139, Lecture 11 UCB MCTS Alpha* AlphaZero A0-MCTS A0-Network A0-Training A0-Evaluation 8/39



9/39

Monte Carlo Tree Search

NPFL139, Lecture 11 UCB MCTS Alpha* AlphaZero A0-MCTS A0-Network A0-Training A0-Evaluation

Monte Carlo Tree Search



Assume we want to estimate $Q(s,\cdot)$ for a given state s to decide how to act in s.

We might perform a lot of simulations from the state s and "promising" actions. Many simulations will probably share prefixes of the trajectories; therefore, we might represent these simulations using a tree.

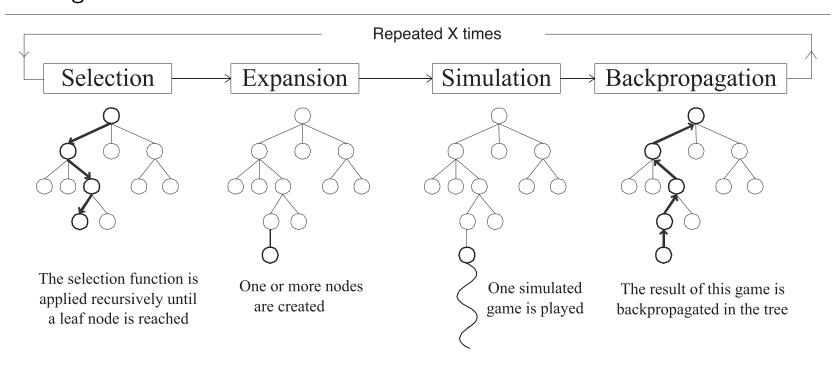


Figure 1 of "Monte-Carlo Tree Search: A New Framework for Game AI" by Guillaume Chaslot et al.

Monte Carlo Tree Search – Selection



When selecting a child to visit, we apply the **UCT** (Upper Confidence Bounds applied for Trees) method, which chooses an action according to

$$a^* = rg \max_a \left[Q(s,a) + c \sqrt{rac{\ln N(s)}{N(s,a)}}
ight],$$

where:

- ullet Q(s,a) denotes the average result of playing action a in state s in the simulations so far,
- ullet N(s) is the number of times the state s has been visited in previous iterations,
- ullet N(s,a) is the number of times the action a has been sampled in state s.

The Q(s,a) values are often normalized to the [0,1] range, and the constant c is game dependent, but some papers propose to try out $\sqrt{2}$ first.

Monte Carlo Tree Search



The MCTS is by nature asymmetrical, visiting promising actions more often.

Expansion

Once a chosen action does not correspond to a node in the search tree, a new leaf is added to the game tree.

Simulation

The simulation phase is the "Monte Carlo" part of the algorithm, and is heavily game-specific. It reaches a terminal state and produces final game outcome.

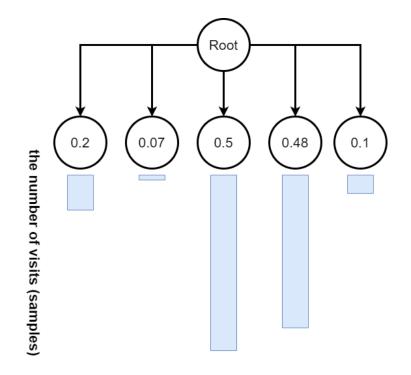


Figure 2 of "Monte Carlo Tree Search: A Review of Recent Modifications and Applications" by M. Świechowski et al.

Backpropagation

The obtained outcome updates the Q(s,a) values in the nodes along the path from the expanded node to the root of the search tree.

Alpha* Gradual Development



- AlphaGo
 - Mar 2016 beat 9-dan professional player Lee Sedol
- AlphaGo Master Dec 2016
 - beat 60 professionals, beat Ke Jie in May 2017
- AlphaGo Zero 2017
 - trained only using self-play
 - surpassed all previous version after 40 days of training
- AlphaZero Dec 2017 (Dec 2018 in Nature)
 - o self-play only, defeated AlphaGo Zero after 30 hours of training
 - o impressive chess and shogi performance after 9h and 12h, respectively
- MuZero Dec 2020 in Nature
 - extends AlphaZero by a trained model
 - later extended into Sampled MuZero capable of handling continuous actions
- AlphaTensor Oct 2022 in Nature
 - automatic discovery of faster algorithms for matrix multiplication



14/39

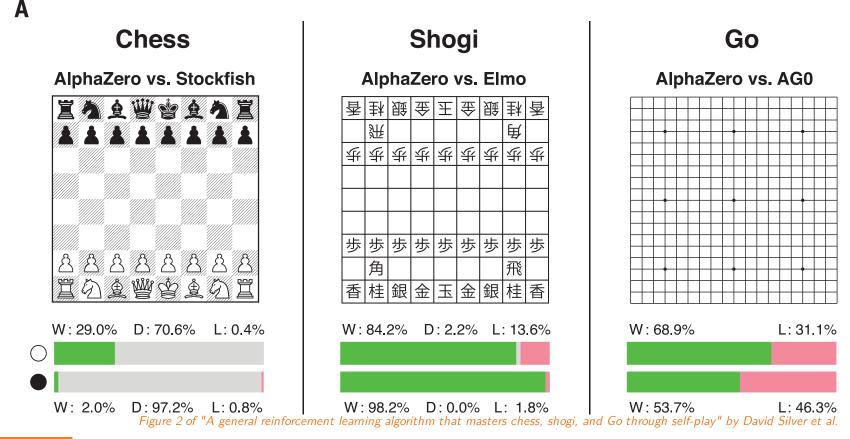
AlphaZero

NPFL139, Lecture 11 UCB MCTS Alpha* ${\sf AlphaZero}$ A0-MCTS A0-Network A0-Training A0-Evaluation

AlphaZero



On 7 December 2018, the AlphaZero paper came out in Science journal. It demonstrates learning chess, shogi and go, *tabula rasa* – without any domain-specific human knowledge or data, only using self-play. The evaluation is performed against strongest programs available.



AlphaZero - Overview



AlphaZero uses a neural network predicting $(\boldsymbol{p}(s),v(s))\leftarrow f(s;\boldsymbol{\theta})$ for a given state s, where:

- ullet policy $oldsymbol{p}(s)$ is a vector of move probabilities, and
- ullet value function v(s) is the expected outcome of the game in range [-1,1].

Instead of the usual alpha-beta search used by classical game playing programs, AlphaZero uses Monte Carlo Tree Search (MCTS).

By a sequence of simulated self-play games, the search can improve the estimate of \boldsymbol{p} and v, and can be considered a powerful policy improvement operator – given a network f predicting policy \boldsymbol{p} and value estimate v, MCTS produces a more accurate policy $\boldsymbol{\pi}$ and better value estimate w for a given state s:

$$egin{aligned} (oldsymbol{p}(s),v(s)) &\leftarrow f(s;oldsymbol{ heta}), \ (oldsymbol{\pi}(s),w(s)) &\leftarrow ext{MCTS}(s,f) \end{aligned}$$

AlphaZero – Overview



The network is trained from self-play games.

A game is played by repeatedly running MCTS from a state s_t and choosing a move $a_t \sim \pmb{\pi}_t$, until a terminal position s_T is encountered, which is then scored according to game rules as $z \in \{-1,0,1\}$. (Note that the range [0,1] is also often used, for example in MuZero.)

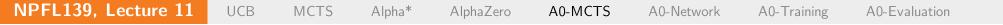
Finally, the network parameters are trained to minimize the error between the predicted outcome v and the simulated outcome z, and maximize the similarity of the policy vector p and the search probabilities π (in other words, we want to find a fixed point of the MCTS):

$$\mathcal{L} \stackrel{ ext{ iny def}}{=} (z-v)^2 - oldsymbol{\pi}^T \log oldsymbol{p} + c \|oldsymbol{ heta}\|^2.$$

The loss is a combination of:

- a mean squared error for the value functions;
- a crossentropy for the action distribution;
- L^2 -regularization.







MCTS keeps a tree of currently explored states from a fixed root state. Each node corresponds to a game state and to every non-root node we got by performing an action a from the parent state. Each state-action pair (s,a) stores the following set of statistics:

- visit count N(s,a),
- ullet total action-value W(s,a),
- ullet mean action value $Q(s,a)\stackrel{ ext{def}}{=} W(s,a)/N(s,a)$, which is usually not stored explicitly,
- ullet prior probability P(s,a) of selecting the action a in the state s.

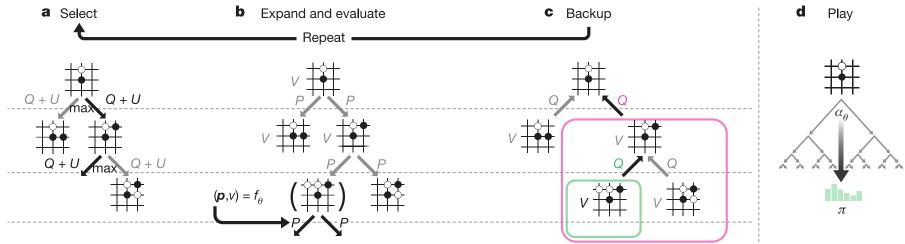


Figure 2 of "Mastering the game of Go without human knowledge" by David Silver et al.



Each simulation starts in the root node and finishes in a leaf node s_L . In a state s_t , an action is selected using a variant of PUCT algorithm as

$$a_t = rg \max_a ig(Q(s_t, a) + U(s_t, a)ig),$$

where

$$U(s,a) \stackrel{ ext{ iny def}}{=} C(s) P(s,a) rac{\sqrt{N(s)}}{1+N(s,a)},$$

with
$$C(s) = \log\left(rac{1+N(s)+c_{ ext{base}}}{c_{ ext{base}}}
ight) + c_{ ext{init}}$$
 being slightly time-increasing exploration rate.

The paper uses $c_{
m init}=1.25$, $c_{
m base}=19652$ without any (public) supporting experiments.

The reason for the modification of the UCB formula was never discussed in any of the AlphaGo/AlphaZero/MuZero papers and is not obvious. However, in June 2020 a paper discussing the properties of this modified formula was published.



Additionally, exploration in the root state $s_{
m root}$ is supported by including a random sample from Dirichlet distribution,

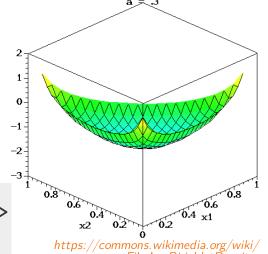
$$P(s_{\mathrm{root}},a) = (1-arepsilon)p_a + arepsilon\operatorname{Dir}(lpha),$$

with arepsilon=0.25 and lpha=0.3,0.15,0.03 for chess, shogi and go, respectively.

Note that using $\alpha < 1$ makes the Dirichlet noise non-uniform, with a smaller number of actions with high probability.

The Dirichlet distribution can be seen as a limit of the Pólya's urn scheme, where in each step we sample from a bowl of balls (with initial counts α) and return an additional ball of the same color to the bowl.

sorted(np.random.dirichlet([0.3] * 9), reverse=True) ->
0.496 0.139 0.124 0.113 0.066 0.035 0.027 0.000 0.000
0.324 0.250 0.151 0.133 0.068 0.044 0.024 0.005 0.001
0.526 0.222 0.111 0.061 0.035 0.026 0.014 0.005 0.000



https://commons.wikimedia.org/wiki/ File:LogDirichletDensityalpha 0.3 to alpha 2.0.gif



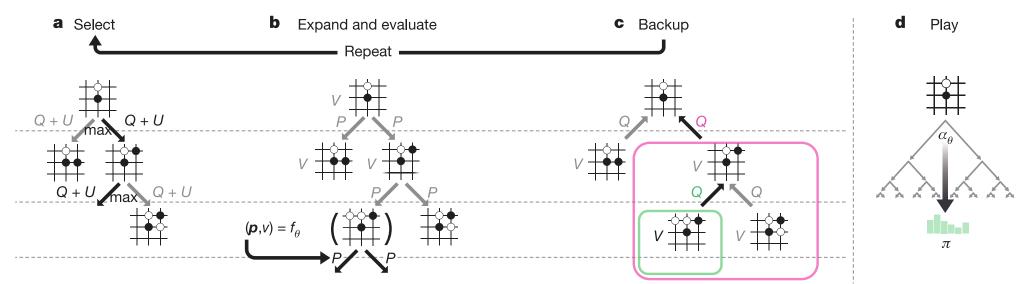


Figure 2 of "Mastering the game of Go without human knowledge" by David Silver et al.

When reaching a leaf node s_L , we:

- evaluate it by the network, generating (\boldsymbol{p},v) ,
- add all its children with N=W=0 and the prior probability ${m p}$,
- in the backward pass for all $t \leq L$, we update the statistics in nodes by performing
 - $\circ~N(s_t,a_t) \leftarrow N(s_t,a_t) + 1$, and
 - $\circ W(s_t, a_t) \leftarrow W(s_t, a_t) \pm v$, depending on the player on turn.



The Monte Carlo Tree Search runs usually several hundreds simulations in a single tree. The result is a distribution proportional to exponentiated visit counts $N(s_{\mathrm{root}},a)^{\frac{1}{\tau}}$ using a temperature τ ($\tau=1$ is mostly used), together with the predicted value function.

The next move is chosen as either:

ullet proportional to visit counts $N(s_{\mathrm{root}},\cdot)^{rac{1}{ au}}$:

$$m{\pi}_{
m root}(a) \propto N(s_{
m root},a)^{rac{1}{ au}},$$

deterministically as the most visited action

$$oldsymbol{\pi}_{ ext{root}} = rg\max_{a} N(s_{ ext{root}}, a).$$

During self-play, the stochastic policy is used for the first 30 moves of the game, while the deterministic is used for the rest of the moves. (This does not affect the internal MCTS search, there we always sample according to PUCT rule.)

AlphaZero – Monte Carlo Tree Search Example



Visualization of the 10 most visited states in a MCTS with a given number of simulations. The displayed numbers are predicted value functions from the white's perspective, scaled to [0, 100]range. The border thickness is proportional to a node visit count.

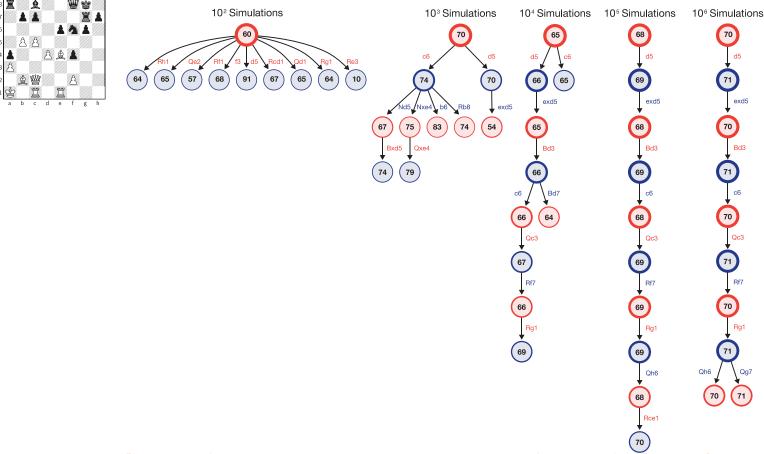


Figure 4 of "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play" by David Silver et al.



AlphaZero – Network Architecture

NPFL139, Lecture 11 UCB MCTS Alpha* AlphaZero A0-MCTS A0-Network A0-Training

A0-Evaluation

AlphaZero – Network Architecture



The network processes game-specific input, which consists of a history of 8 board positions encoded by several $N \times N$ planes, and some number of constant-valued inputs.

Output is considered to be a categorical distribution of possible moves. For chess and shogi, for each piece we consider all possible moves (56 queen moves, 8 knight moves, and 9 underpromotions for chess).

The input is processed by:

- initial convolution block with CNN with 256 3×3 kernels with stride 1, batch normalization, and ReLU activation,
- 19 residual blocks, each consisting of two CNN with 256 3×3 kernels with stride 1, batch normalization, ReLU activation, and a residual connection around them,
- policy head, which applies another CNN with batch normalization, followed by a convolution with 73/139 filters for chess/shogi, or a linear layer of size 362 for go,
- value head, which applies another CNN with one 1×1 kernel with stride 1, followed by a ReLU layer of size 256, and a final \tanh layer of size 1.

AlphaZero – Network Inputs



Go		Chess		Shogi			
Feature	Planes	Feature	Planes	Feature	Planes		
P1 stone	1	P1 piece	6	P1 piece	14		
P2 stone	1	P2 piece	6	P2 piece	14		
		Repetitions	2	Repetitions	3		
				P1 prisoner count	7		
				P2 prisoner count	7		
Colour	1	Colour	1	Colour	1		
		Total move count	1	Total move count	1		
		P1 castling	2				
		P2 castling	2				
		No-progress count	1				
Total	17	Total	119	Total	362		

Table S1 of "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play" by David Silver et al.

AlphaZero – Network Outputs



Chess		Shogi			
Feature Planes		Feature	Planes		
Queen moves	56	Queen moves	64		
Knight moves	8	Knight moves	2		
Underpromotions	9	Promoting queen moves	64		
		Promoting knight moves	2		
		Drop	7		
Total	73	Total	139		

Table S2 of "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play" by David Silver et al.

Alpha*



AlphaZero – Training

NPFL139, Lecture 11

UCB

MCTS

Alpha*

AlphaZero

A0-MCTS

A0-Network

A0-Training

A0-Evaluation

AlphaZero - Training

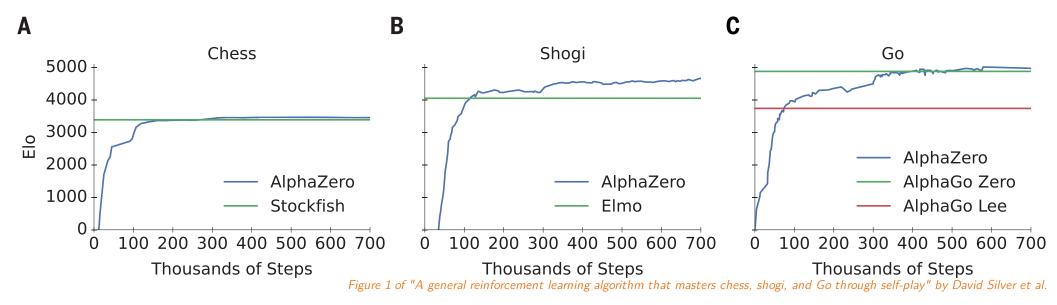


Training is performed by running self-play games of the network with itself. Each MCTS uses 800 simulations. A replay buffer of one million most recent games is kept.

During training, 5000 first-generation TPUs are used to generate self-play games. Simultaneously, network is trained using SGD with momentum of 0.9 on batches of size 4096, utilizing 16 second-generation TPUs. Training takes approximately 9 hours for chess, 12 hours for shogi and 13 days for go.

AlphaZero – Training





	Chess	Shogi	Go
Mini-batches	700k	700k	700k
Training Time	9h	12h	13d
Training Games	44 million	24 million	140 million
Thinking Time	800 sims	800 sims	800 sims
-	$\sim 40~\mathrm{ms}$	$\sim 80~\mathrm{ms}$	$\sim 200~\mathrm{ms}$

Table S3 of "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play" by David Silver et al.

AlphaZero – Training



According to the authors, training is highly repeatable.

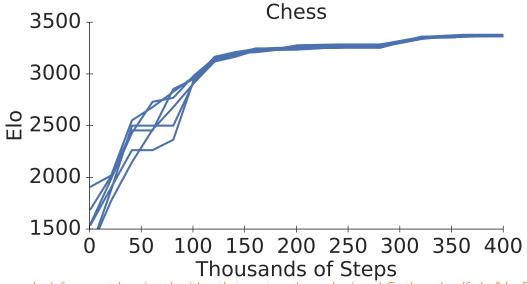


Figure S3 of "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play" by David Silver et al.

Alpha*

AlphaZero – Symmetries



In the original AlphaGo Zero, symmetries (8 in total, using rotations and reflections) were explicitly utilized, by

- randomly sampling a symmetry during training,
- randomly sampling a symmetry during MCTS evaluation.

However, AlphaZero does not utilize symmetries in any way (because chess and shogi do not have them), so it generates only approximately 1/8 positions per training step.

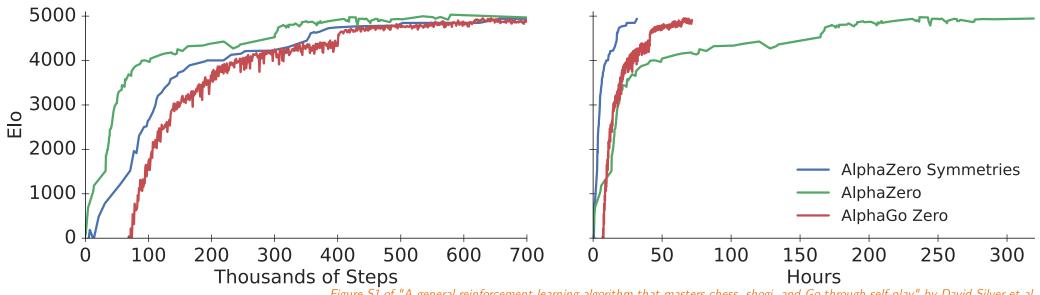


Figure S1 of "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play" by David Silver et al.

Alpha*

AlphaZero – Inference



During inference, AlphaZero utilizes much less evaluations than classical game playing programs.

Program	Chess	Shogi	Go
AlphaZero Stockfish	63k (13k) 58,100k (24,000k)	58k (12k)	16k (0.6k)
Elmo		25,100k (4,600k)	
AlphaZero	1.5 GFlop	1.9 GFlop	8.5 GFlop

Table S4 of "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play" by David Silver et al.

AlphaZero – Main Results



Evaluation against strongest program available.

Α Shogi Chess Go AlphaZero vs. Stockfish AlphaZero vs. Elmo AlphaZero vs. AG0 置为食业等食为置 香|掛|銀|金|玉|金|銀|科|香 | 찬| 찬| 찬| 찬| 찬| 찬| 찬 |歩|歩|歩|歩|歩|歩|歩|歩 角 罩勾食豐堂桌交 |桂|銀|金|玉|金|銀|桂|香 W: 29.0% D: 70.6% L: 0.4% W: 84.2% D: 2.2% L: 13.6% W: 68.9% L: 31.1% W: 2.0% D: 97.2% L: 0.8% W: 98.2% D: 0.0% L: 1.8% W:53.7% L: 46.3% Figure 2 of "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play" by David Silver et al.



UCB

MCTS

Alpha*

AlphaZero

A0-MCTS

A0-Network

A0-Training

A0-Evaluation

AlphaZero – Ablations



		AlphaZero			Opponent				
Fig.	Match	Start Position	Book	Main	Inc	Book	Main	Inc	Program
2A	Main	Initial Board	No	3h	15s	No	3h	15s	Stockfish 8
2B	1/100 time	Initial Board	No	108s	0.15s	No	3h	15s	Stockfish 8
2B	1/30 time	Initial Board	No	6min	0.5s	No	3h	15s	Stockfish 8
2B	1/10 time	Initial Board	No	18min	1.5s	No	3h	15s	Stockfish 8
2B	1/3 time	Initial Board	No	1h	5s	No	3h	15s	Stockfish 8
2C	latest Stockfish	Initial Board	No	3h	15s	No	3h	15s	Stockfish 2018.01.1
2C	Opening Book	Initial Board	No	3h	15s	Yes	3h	15s	Stockfish 8
2D	Human Openings	Figure 3A	No	3h	15s	No	3h	15s	Stockfish 8
2D	TCEC Openings	Figure S4	No	3h	15s	No	3h	15s	Stockfish 8

Table S8 of "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play" by David Silver et al.

			AlphaZero		Opponent			ent	
Fig.	Match	Start Position	Book	Main	Inc	Book	Main	Inc	Program
2A	Main	Initial Board	No	3h	15s	Yes	3h	15s	Elmo
2B	1/100 time	Initial Board	No	108s	0.15s	Yes	3h	15s	Elmo
2B	1/30 time	Initial Board	No	6min	0.5s	Yes	3h	15s	Elmo
2B	1/10 time	Initial Board	No	18min	1.5s	Yes	3h	15s	Elmo
2B	1/3 time	Initial Board	No	1h	5s	Yes	3h	15s	Elmo
2C	Aperyqhapaq	Initial Board	No	3h	15s	No	3h	15s	Aperyqhapaq
2C	CSA time control	Initial Board	No	10min	10s	Yes	10min	10s	Elmo
2D	Human Openings	Figure 3B	No	3h	15s	Yes	3h	15s	Elmo

Table S9 of "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play" by David Silver et al.

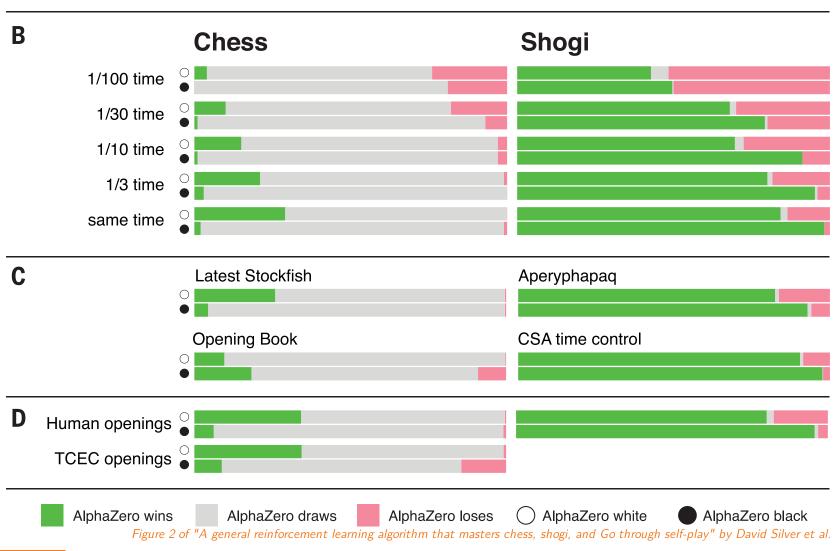
UCB

MCTS

Alpha*

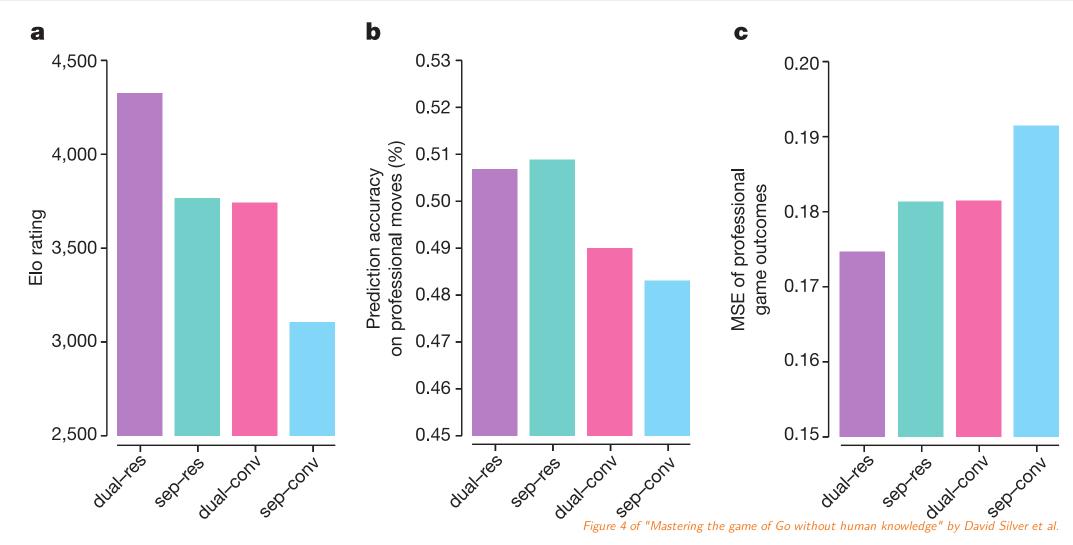
AlphaZero – Ablations





AlphaZero – Ablations





NPFL139, Lecture 11 UCB MCTS Alpha* AlphaZero A0-MCTS A0-Network A0-Training A0-Evaluation 38/39

AlphaZero – Preferred Chess Openings



39/39

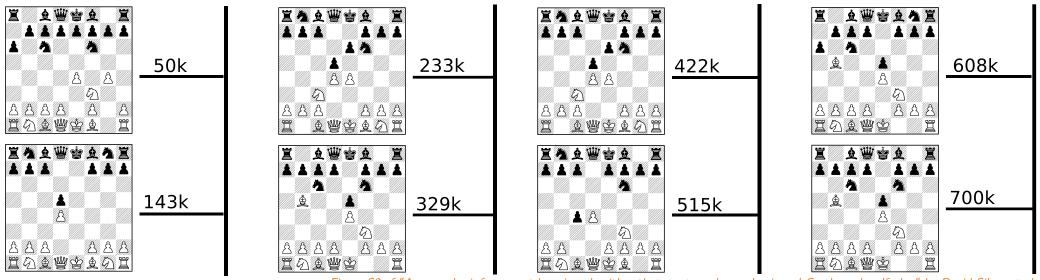


Figure S2 of "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play" by David Silver et al.

NPFL139, Lecture 11 UCB MCTS Alpha* AlphaZero A0-MCTS A0-Network A0-Training A0-Evaluation