

Applications of Model-Based Deep Reinforcement Learning

Johanni Brea

23 April 2024

Artificial Neural Networks CS-456

Table of Contents

1. Model-Based Deep Reinforcement Learning

Model-Based Deep RL with Variational State Tabulation (VaST)
Other Model-Based Deep RL Methods

2. Monte Carlo Tree Search

3. AlphaZero

4. MuZero

5. Other Success Stories of Deep RL

Model-Based Deep Reinforcement Learning

A model-based reinforcement learning method estimates the transition dynamics (e.g. $P_{s \rightarrow s'}^a$) and reward structure (e.g. $R_{s \rightarrow s'}^a$).

Today you will learn

1. how the model can be used for **planning**,
2. how we can use standard planning methods in deep RL with variational state tabulation (**VaST**).
3. why we may prefer **decision time planning** over **background planning**,
4. how **AlphaZero** learns to play games with **Monte Carlo Tree Search**,
5. how **MuZero** reaches the same level as AlphaZero without explicitly knowing the rules of the game.

Suggested: Sutton & Barto Ch. 8,

[Corneil et al., 2018, Silver et al., 2016, Silver et al., 2018, Schrittwieser et al., 2019]

Model-Based Tabular and Deep RL

- ▶ Tabular RL: finite number of states and action (without structure).
 - ▶ Use samples (experiences) (s_t, a_t, r_t, s_{t+1}) to update the estimate of the model:
increment counts $N_{s \rightarrow s'}^a \rightarrow N_{s \rightarrow s'}^a + \delta_{s, s_t} \delta_{a, a_t} \delta_{s', s_{t+1}}$ and
cumulative returns $W_{s \rightarrow s'}^a \rightarrow W_{s \rightarrow s'}^a + r_t \delta_{s, s_t} \delta_{a, a_t} \delta_{s', s_{t+1}}$ and
compute model estimates $\hat{P}_{s \rightarrow s'}^a = \frac{N_{s \rightarrow s'}^a}{\sum_{s'} N_{s \rightarrow s'}^a}$ and $\hat{R}_{s \rightarrow s'}^a = \frac{W_{s \rightarrow s'}^a}{N_{s \rightarrow s'}^a}$.
 - ▶ Use any MDP (Markov Decision Process) solver to compute a policy given the estimated model.
- ▶ Deep RL: infinite (or very high) number of states (generalization desired).
 - ▶ How should samples be used to estimate a model?
 - ▶ How should the model be used to compute a policy?

Model-Based Deep RL with Variational State Tabulation (VaST)



Raw visual observations (in everyday life and in video games) are almost always at least a tiny bit different from one another, but the value of similar observations may be similar.

⇒ we want model-based RL that generalizes well.

Corneil, Gerstner, Brea (2018), ICML

Notes



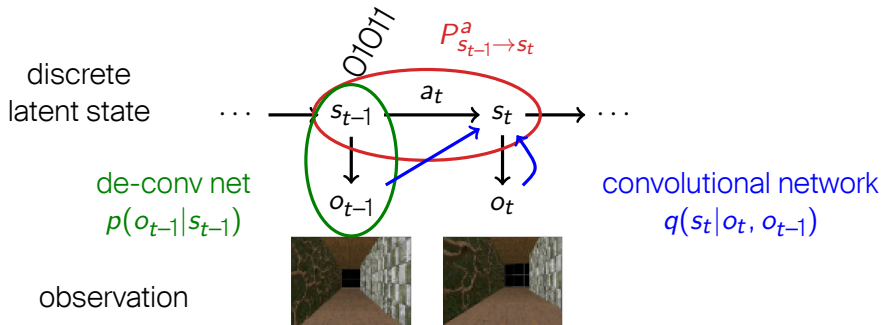
Raw visual observations (in everyday life and in video games) are almost always at least a tiny bit different from one another, but the value of similar observations may be similar.
→ we want model-based RL that generalizes well.

Cornel, Geometric Deep Learning (2018), ©

- What is a discrete state in the environment of animals, humans or in video games? Every raw observation looks at least a tiny bit different from all other observations, but the value of states that look similar (e.g. observations that are only one footstep separated) are likely to be similar.
- Naively assigning each raw observation to a new, discrete state and learning the value for each state is unreasonable: there are many states, most states are visited only once or a few times and in standard tabular reinforcement learning there is no generalization across states.
- Idea: learn a mapping to abstract discrete states such that neighbouring observations are mapped to the same discrete state and define a similarity metric in the space of abstract discrete states such that similar observations are mapped to similar discrete states.
- With discrete abstract states we can use again standard planning methods.


Model-Based Deep RL with Variational State Tabulation (VaST)

Learn to encode images as discrete states (similar images o , same discrete state s).
Whenever the agent learns something new about the environment ($P_{s \rightarrow s'}^a$ or $R_{s \rightarrow s'}^a$ changes), update offline the policy by running (a smart version of) value iteration.





Corneil, Gerstner, Brea (2018), ICML

Notes



- The abstract states (discrete latent states) are binary words of length K . The similarity metric is the Hamming distance, i.e. the number of bits that are different between any pair of binary words.
- The binary words can be represented as integers. These integers are the states that can be used to estimate the model, i.e. the transition probabilities and the reward table. The model is used to find a good policy by running a smart version of value iteration that prioritizes updating the values of states, whose next states had recently an update of the value.
- The function that maps observations to binary words (encoder) is a convolutional network with K binary output neurons. To deal with moving objects we use the same trick as in DQN for Atari games and give as input multiple subsequent observations.
- To make sure that the discrete latent state contains relevant information about the state we use a deconvolutional network that takes latent states as input and produces raw observations. The only purpose of this decoder is to enforce the encoder to learn a latent representation that discriminates different states; otherwise it may happen that the encoder maps every observation to a single latent state s_0 and $P_{s_0 \rightarrow s_0}^a = 1$, $R_{s_0 \rightarrow s_0}^a = \text{average reward}$, but no useful policy can be found.

WorldModels, SimPLe, Dreamer



- ▶ The model can consist of video-frame predictions (e.g. SimPLe).
- ▶ The model can consist of abstract latent states with learned transition dynamics (e.g. WorldModels, DreamerV2).
- ▶ The model can be used as a simulator to train a standard model-free RL method (e.g. PPO, Actor-Critic.)

Other approaches to model-based deep RL exist.

Most of them

interact with the environment to collect data,

build a model (e.g. with the help of a neural network)

and sample from the model to learn a policy.


[Ha and Schmidhuber, 2018, Kaiser et al., 2019, Hafner et al., 2020, Hafner et al., 2023]

Notes




- Alternatively to classic planning methods, the world model can be used to simulate experiences and train a model-free reinforcement learning policy on the samples from the world model.
- In video-frame prediction a neural network learns to approximate directly $p(o_{t+1}|a_t, o_t, o_{t-1}, \dots)$ (SimPLe = simulated policy learning).
- WorldModels and DreamerV2 learn an abstract representation with an autoencoder approach similar to VaST, but instead of using planning methods in the latent state, they sample (dreamed) trajectories to train a policy.
- It may sound silly to replace a simulator (like the Atari video game environment) by an approximate simulator that is learned with a neural network and then use this approximate simulator to train a policy instead of training a policy directly on the original simulator. The goal is, however, to show that training with the learned simulator allows to reduce the sample complexity, i.e. the number of interactions with the actual environment. This can be useful for environments that are costly to interact. For example, imagine a robot that can cook. If it has a somewhat accurate model of the things that can happen in a kitchen, this robot can simulate many cooking episodes until it has found a policy that leads to a great meal, without having to waste any resources during the episodes that are used to train the policy.
- Currently, model-based deep reinforcement learning is still a very active field of research, where the goal is to show that using a model leads to policies that are as good as those obtained with model-free reinforcement learning while requiring far less interactions with the environment than the model-free methods.

DreamerV3



(a) World Model Learning





(b) Actor Critic Learning

- ▶ Discrete latent representation z_t : variational auto-encoder

- ▶ World-model: $h_{t+1} = f_\phi(h_t, z_t, a_t)$


- ▶ Policy: Actor-Critic trained with samples from world model

- ▶ Performs well on an impressive list of tasks!



[Hafner et al., 2023]

Notes



- DreamerV3 is a particularly impressive model in this class.
- As in VaST a discrete latent representation z_t is learned in a variational auto-encoder style.
- In contrast to VaST, this discrete latent state is not directly used as a state in a tabular RL method, but it is given to a recurrent neural network ($h_{t+1} = f_{\phi}(h_t, z_t, a_t)$) that is trained to predict the next latent state z_{t+1} . Using a recurrent neural network helps to deal with environments that are not fully Markovian. For fully Markovian environments one could also just use a neural network to predict directly the next state $z_{t+1} = f_{\phi}(z_t, a_t)$.
- The biggest difference to VaST is, that the learned model is not solved with any direct optimization method, like value iteration. Instead, the model is used to sample (“dream”) trajectories (right hand side in schematic). Importantly, sampling does not require interactions with the environment. These samples are used to train an actor-critic.
- Note that one assumes that sampling is cheaper than actually interacting with the environment; therefore, to train the policy, tricks like replay buffers (DQN-style) or surrogate losses (PPO-style) are not needed, but one could, of course, sample multiple trajectories in parallel (A2C-style). However, a replay buffer of actual experiences is still needed to train the model!



Summary: Background Planning and Variational State Tabulation

- ▶ In model-based reinforcement learning we can use background planning/learning to update the Q- and V-values efficiently (e.g. by value iteration), if the number of states and actions is not too large.
- ▶ Alternatively, the model can be used to sample experiences offline and train a model-free agent with these artificial samples.
- ▶ Learned abstractions may be useful to map high dimensional states (like images) to discrete states that can be used with standard planning methods.
- ▶ Model-based RL is typically more sample efficient than model-free RL, in the sense that fewer interactions with the real environment are needed.
- ▶ Model-based RL is typically computationally more expensive than model-free RL, because of training the world model and using it to improve the policy.

Value Iteration Does Not Scale!


- ▶ There are presumably less than 10^8 intersections of roads in Europe. Value iteration on current compute hardware is feasible for this problem.
- ▶ But there are more than 10^{170} possible positions for the game Go. (The number of atoms on earth is roughly 10^{50}).

Decision Time Planning



1. Start where you are e.g. in L
2. Expand the decision tree, i.e. move to every state s that is a successor of L ($P_{L \rightarrow s}^a$).
3. Iteratively move to the successors of the successors up to a certain depth or until a terminal state is reached.
4. Propagate values backwards along the decision tree like in backward planning (but only along the decision tree).

Notes



- When full value iteration is not feasible, one can start from the current state and expand the decision tree.
- For example, if we want to find the shortest path from Lausanne to Rome, we can start in Lausanne, expand the decision tree by moving to every successor state (neighbouring city) and from each successor state iteratively continue expanding the decision tree up to a certain depth or until a terminal state is reached. Then we propagate values backwards along the decision tree. In this example we see that the optimal solution can be found at depth 2 of the decision tree. Thus, the update equation of value iteration (Eq. 1 on slide 3) needs to be applied only three times (for Nice, Milano and Lausanne; the branch following Zurich has not yet reached any terminal node).
- In general, it is not a good idea to ignore branches that have not yet reached a terminal node (if there were a teleporter in Frankfurt that takes zero cost to travel to Rome, the optimal path would only be discoverable at depth 3 of the decision tree and thus ignoring the node Zurich at depth 2 would not be a good idea), but it may be too costly to expand the decision tree to large depth. One strategy to deal with this is Monte Carlo Tree Search.

Table of Contents

1. Model-Based Deep Reinforcement Learning

Model-Based Deep RL with Variational State Tabulation (VaST)
Other Model-Based Deep RL Methods

2. Monte Carlo Tree Search


3. AlphaZero

4. MuZero


5. Other Success Stories of Deep RL

Monte Carlo Tree Search(MCTS)


Selection




Expansion



Simulation



Backpropagation




root node = node on top of the tree, e.g. current position of a game.

internal node = node where each action has already been taken at least once.

leaf node = node with untaken actions.

terminal node = node at which the tree expansion ends, e.g. end of a game.

Notes



- In MCTS the tree gets expanded in an iterative fashion.
- In the figure you see a tree starting at a certain position in a two player game. It is known that white won 12 out of 21 simulated games from the current position. If white chooses the first action the game reaches a position (where it is the black player's turn), where white won 7 out of 10 simulated games, etc.
- In MCTS, one of the leaf nodes is selected for expansion (in this case it is a position where it is black's turn and white won 3 out of 3 simulated games). One of black's allowed actions is taken and a full game is simulated with a random policy (wiggly line). The random game from the expanded position led to a loss (0/1). This value is back-propagated. At the end of this iteration white one 12 out of 22 simulated games from the root position.
- Thanks to the simulation step that always ends at a terminal node, every leaf node has an estimate of the actual value. However, this estimate can be terrible, because the simulation is run with a random policy, which may fail to reflect that some positions lead to a win with certainty, if the actions are selected carefully.

Monte Carlo Tree Search

- 1: Initialize root node s_0 with visitation count $N(s_0) = 0$, state-action count $N(s_0, a) = 0$, total return $W(s_0, a) = 0$ and expected return $Q(s_0, a) = 0 \forall a$.
- 2: (**Selection**) Start at root node $s = s_0$.
- 3: **repeat**
- 4: Select action $a^* = \arg \max_a Q(s, a) + \sqrt{\frac{2 \log(N(s))}{N(s, a)}}$
- 5: Increment $N(s) \leftarrow N(s) + 1$, $N(s, a^*) \leftarrow N(s, a^*) + 1$ and go to next state s
- 6: **until** s is a leaf node.
- 7: (**Expansion**) Take a new action a_N that has never been taken in s , increment $N(s, a_N)$, go to next state s_N and set $N(s_N) = N(s_N, a) = W(s_N, a) = Q(s_N, a) = 0 \forall a$.
- 8: (**Simulation**) Simulate tree expansion with default policy from s_N (e.g. taking actions uniformly at random or some smart heuristic) and observe value v at terminal node.
- 9: (**Backpropagation**) Increment $W(s, a) \leftarrow W(s, a) + v$ for all nodes visited during the selection and recompute $Q(s, a) = W(s, a) / N(s, a)$.
- 10: Iterate these four steps (lines 2-9) for as long as you want.
- 11: Select actual action $\arg \max_a Q(s_0, a)$ to perform in state s_0 .

Notes


1. Initialize root node s_0 with visitation count $N(s_0) = 0$, state-action count $N(s_0, a) = 0$, total return $W(s_0, a) = 0$ and expected return $Q(s_0, a) = 0$ via
2. (Selection) Start at root node $s = s_0$.
3. repeat
 4. Select action $a' = \arg \max_a Q(s, a) + \sqrt{\frac{2 \ln N(s)}{N(s, a)}}$
 5. Increment $N(s) \leftarrow N(s) + 1$, $N(s, a') \leftarrow N(s, a') + 1$ and go to next state s
 6. until s is a leaf node.
 7. (Expansion) Take a new action a_N that has never been taken in s , increment $N(s, a_N)$ or to next state s_N and set $W(s_N) \leftarrow W(s, a) + Q(s_N, a) = 0$ via
 8. (Simulation) Simulate tree expansion with default policy from s_N (e.g. taking actions uniformly at random or some smart heuristic) and observe value v at terminal node.
 9. (Backpropagation) Increment $W(s, a) \leftarrow W(s, a) + v$ for all nodes visited during the selection and recompute $Q(s, a) \leftarrow W(s, a) / N(s, a)$.
 10. Break from loop state (line 2-6) for as long as you want.
 11. Select actual action $\arg \max_a Q(s_0, a)$ to perform in state s_0 .

- 1: If one has already visitation counts etc. from previous runs of MCTS, one can use them to initialize the root node with other values than 0.
- 2: Every iteration starts at the root node, e.g. the current game position.
- 4: Here we show a variant of the UCB1 exploration-exploitation criterion. $Q(s, a)$ drives exploitation, the second term drives exploration. The exploration term should be such that it decreases with $N(s, a)$ (actions with low $N(s, a)$ should be explored) and increases slowly with $N(s)$ (if s is visited often, we want to be really sure that none of the less taken actions would in fact be optimal; an increase in $N(s)$ drives occasional re-exploration). Other explorations policies are possible, e.g. other functions of $N(s)$ and $N(s, a)$, ϵ -greedy exploration, Thompson sampling.
- 7: If a node has actions that have never been taken before, one of this actions is taken and the landing state s_N is initialized.
- 8: The random game is played to have a very crude approximation of the value of state s_N .
- 9: Despite its name, backpropagation in MCTS is completely different from the error backpropagation algorithm for deep neural networks! In MCTS, backpropagation means that the Q-values along the path that was taken during selection and expansion should be updated with the value that was obtained at the terminal node.
- 11: After many iterations the tree has grown quite a bit and we can select at the root node the action that has the highest Q-value. We do not need to perform exploration in the actual environment, because the model is perfectly known and MCTS would converge to the correct Q-value in the limit of $N(s_0) \rightarrow \infty$.


Monte Carlo Tree Search: An Example

An example with two possible actions a_1 and a_2 in every state.

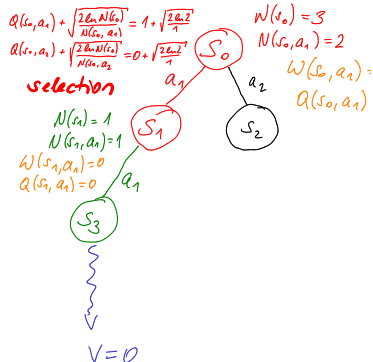
iteration 1



iteration 2




iteration 3



An alternative example: <https://youtu.be/UXW2yZndl7U>

Notes



iteration 1 The root node is a leaf node, because it has actions that have never been explored. Therefore we **expand** with one of the untaken actions (a_1) and increment the counts. Next we **simulate** a random game and observe a win. This is **backpropagated** to keep track of the fact that taking action a_1 in s_0 led to a win.

iteration 2 The root node is still a leaf node, because action a_2 has not yet been explored. We **expand** the tree with action a_2 , **simulate** a random game and observe a loss. This is **backpropagated** to keep track of the fact that taking action a_2 in s_0 led to a loss.

iteration 3 Now the root node is not a leaf node anymore and we need to select. Because both actions have been taken once and therefore have the same exploration bonus, the Q-value determines the action and a_1 is **selected**. State s_1 is a leaf node, therefore it gets **expanded** etc.

iteration 4 (not shown on the slide). At the root node we need now to compare $\frac{1}{2} + \sqrt{\frac{2 \log(3)}{2}}$ (for action a_1) to $0 + \sqrt{\frac{2 \log(3)}{1}}$ (for action a_2) to decide which action to take. Despite the larger exploration bonus for action a_2 , the total value for action a_1 is larger and therefore a_1 would be selected. Next, s_1 would be expanded with action a_2 and thus become an internal node.

Quiz

Which statement is correct?

- With background planning, the computation at decision time is very cheap, since one can use the stored Q -values to find the optimal action.
- In Monte Carlo Tree Search, the computation at decision time is very cheap, since one can use the stored Q -values to find the optimal action.
- Monte Carlo Tree Search plays games with a random policy to estimate the value of a position.
- The model in model-based deep RL can be used to “dream up” sequences that were never actually experienced.
- Typically, model-based deep RL methods are computationally heavier but more sample efficient than model-free deep RL methods.

Table of Contents

1. Model-Based Deep Reinforcement Learning

Model-Based Deep RL with Variational State Tabulation (VaST)

Other Model-Based Deep RL Methods

2. Monte Carlo Tree Search

3. AlphaZero

4. MuZero

5. Other Success Stories of Deep RL

The Game Go



- ▶ $19 \times 19 = 361$ positions.
- ▶ Decision tree grows very quickly!
- ▶ Traditional Go programs use human expert knowledge to focus MCTS on relevant branches.
- ▶ AlphaZero learns focus through self-play.

Notes



- 19 x 19 = 361 positions.
- Decision tree grows very quickly!
- Traditional Go programs use human expert knowledge to focus MCTS on relevant branches.
- AlphaZero learns focus through self-play.

- Go is usually played on a board with $19 \times 19 = 361$ positions.
- For the first few actions there are thus more than 300 legal positions to place the stone; even in later stages of the game, there are many legal actions. Expanding the decision tree in MCTS is therefore very costly ($300^3 = 27$ million).
- Traditional Go computer programs (like CrazyStone) use MCTS with heuristics to exclude moves that human experts would never play.
- AlphaZero does not use any expert knowledge, but learns to focus on relevant parts of the decision tree thanks to some neural networks that are trained with self-play.

AlphaZero

- 1: Initialize root node s_0 with $N(s_0) = 0, N(s_0, a) = 0, W(s_0, a) = 0, Q(s_0, a) = 0$ and prior action probability $P(s_0, a) = \pi_\theta(s_0)_a \forall a$ (π_θ is a neural network).
- 2: (**Selection**) Start at root node $s = s_0$.
- 3: **repeat**
- 4: Select action $a^* = \arg \max_a Q(s, a) + P(s, a)C(s) \frac{\sqrt{N(s)}}{N(s,a)+1}$
- 5: Increment $N(s) \leftarrow N(s) + 1, N(s, a^*) \leftarrow N(s, a^*) + 1$ and go to next state s
- 6: **until** $N(s) = 0$.
- 7: (**Expansion**) Take a new action a_N that has never been taken in s , increment $N(s, a_N)$, go to s_N and set $N(s_N) = N(s_N, a) = W(s_N, a) = Q(s_N, a) = 0$ and $P(s_N, a) = \pi_\theta(s_N)_a$.
- 8: (**No Simulation**) Set $v = v_\theta(s_N)$ (v_θ is a neural network).
- 9: (**Backpropagation**) Increment $W(s, a) \leftarrow W(s, a) + v$ for all nodes visited during the selection and recompute $Q(s, a) = W(s, a)/N(s, a)$.
- 10: Iterate these four steps (lines 2-9) for as long as you want.
- 11: **Select actual action** $\arg \max_a N(s_0, a)$ to perform in state s_0 during evaluation or **sample action** a from $\pi(a|s_0) = (N(s_0, a) + \alpha)/(\sum_b N(s_0, b) + \alpha)$ during training.

Notes

```

1. Initialize root node  $s_0$  with  $N(s_0) = C, N(s_0, a) = 0, W(s_0, a) = 0, Q(s_0, a) = 0$  and
   prior action probabilities  $P(s_0, a) = \frac{1}{|A(s_0)|}$  for all  $a \in A(s_0)$ .
2. Selection: Start at root node  $s = s_0$ .
3. Expand:
   a. Select action  $a' = \arg \max_a Q(s, a) + P(s, a) \sqrt{\frac{c_1}{N(s, a) + c_2}}$ .
   b. Increment  $N(s) \leftarrow N(s) + 1, N(s, a') \leftarrow N(s, a') + 1$  and go to next state  $s'$ .
   c. Leaf Node: If  $s'$  is a terminal state, increment  $W(s', a')$  and  $Q(s', a') \leftarrow W(s', a')$ .
   d. Backpropagation: Increment  $W(s, a) \leftarrow W(s, a) + v$  for all nodes visited during the selection and recompute  $Q(s, a) \leftarrow W(s, a) / N(s, a)$ .
4. Evaluate: Iterate these first three steps (lines 2-4) for as long as you want.
5. Select actual action  $a$  from  $\pi(s_0, a) \leftarrow \frac{N(s_0, a) + \alpha}{\sum_b N(s_0, b) + \alpha}$  during evaluation or
   sample action  $a$  from  $\pi(s_0, a) \leftarrow \frac{N(s_0, a) + \alpha}{\sum_b N(s_0, b) + \alpha}$  during training.

```

- AlphaZero modifies MCTS in two important ways: a neural network to compute prior action probabilities helps in the selection process to focus on relevant parts of the decision tree and another neural network helps to estimate the value at leaf nodes, making simulations of random games unnecessary.
- The prior action probabilities for root node s_0 are computed with a neural network.
 - In the selection process, these prior action probabilities influence exploration. For example, “silly actions” that lead almost certainly to a loss will not be explored with $P(s, a) \approx 0$ (usually one does not need to consider how a chess game continues after sacrificing the queen for no reason), whereas promising actions with large $P(s, a)$ are heavily explored. Side note: the exploration rate $C(s) = c_1 + \log((1 + N(s) + c_2)/c_2)$ grows slowly with search time, but is essentially constant if $N(s)$ is small (in MuZero $c_1 = 1.25$, $c_2 = 19652$; I didn’t find the parameter values for AlphaZero). I do not know why the factor $C(s)\sqrt{N(s)}/(N(s, a) + 1)$ is chosen instead of the UCB1 exploration term (it may have better convergence bounds). But this term satisfies the criteria we discussed on line 4 in slide 12.
 - In the expansion step the prior action probabilities are computed using the neural network π_θ .
 - Instead of simulating a random game, a value network is used to estimate the value (win probability) of a given position.
 - Greedy actions are selected during evaluation. The argmax of the counts is apparently less sensitive to outliers than the argmax of the Q-values. During training the actions are sampled from $\pi(a|s_0) = (N(s_0, a) + \alpha) / (\sum_b N(s_0, b) + \alpha)$ where $\alpha = \{0.3, 0.15, 0.03\}$ for chess, shogi and go, respectively. Thanks to non-zero α random actions are occasionally taken.

AlphaZero: How to train the networks

Loss function

$$L(\theta) = \sum_{s_i} (z(s_i) - v_{\theta}(s_i))^2 - \sum_{a_i} \pi(a_i|s_i) \log \pi_{\theta}(s_i)_a + c \|\theta\|_2^2$$

where $z(s_i)$ is the final outcome of the game in which position s_i was encountered, and $\pi(a_i|s_i)$ is the action selection policy of MCTS during training.

Sample positions s_i are generated during self-play, i.e. the same networks and MCTS are used to select actions for both players.

After some games of self-play, the gradient of $L(\theta)$ is computed to update the parameters; only positions s_i since the last update are used to evaluate the gradient.

Notes

Loss function

$$\mathcal{L}(\theta) = \sum_k (v_\theta(s_k) - v(s_k))^2 + \sum_k \pi_\theta(a_k|s_k) \log \pi_\theta(a_k|s_k) + \pi_\theta(a_k|s_k)$$

where s_k is the final outcome of the game in which position s_k was encountered, and $\pi_\theta(a_k|s_k)$ is the action selection policy of MCTS during training.


Sample positions s_k are generated during self-play, i.e. the same networks and MCTS are used to select actions for both players.

After some games of self-play, the gradient of $\mathcal{L}(\theta)$ is computed to update the parameters, only positions s_k since the last update are used to evaluate the gradient.


- With the squared error loss the value $v_\theta(s_j)$ learns to approximate the win probability from a given position.
- The prior action probabilities are trained with the cross-entropy loss to approximate the policy with which MCTS selects actions. Like this, exploration will be biased towards actions that were in the past selected by MCTS. This may look like a vicious loop, where MCTS is biased towards actions deemed useful by π_θ and π_θ learns to imitate the MCTS policy. I guess the sampling step in line 11 of AlphaZero during training, assures that π_θ does not get stuck at a sub-optimal policy where the probability of some actions would be zero and does never explored.
- The prior action probabilities π_θ resemble the policy of a model-free RL agent. However, it is not trained by a standard model-free RL method, but instead learns to imitate the policy of MCTS. We are thus in a funny situation, where we use model-free reinforcement learning to drive exploration in a setting where the full model is known, but the state space is so large that standard model-based approaches, like value iteration, fail to solve the problem.
- Why does MCTS still make sense, once we have π_θ ? There are so many possible states that π_θ remains a crude approximation of the MCTS policy, even after extensive training. A useful analogy may be that π_θ is the policy that a professional player uses when forced to make a move without time to think (e.g. in blitz-chess). Given time to think, however, the player may use this policy to explore promising actions in the mind, evaluate resulting positions with v_θ and take a more informed decision than without thinking.

AlphaZero: representations and architecture

Policy network




Value network



- ▶ input $2MT + L$ planes of $N \times N$ positions (M = number of pieces types, factor 2: own pieces & opponent pieces, T subsequent positions, L extra information)
- ▶ output: for each of K action types one plane.
- ▶ deep conv nets with residual blocks and batch normalization.


Notes



- For games with $N \times N$ positions and M different types of pieces the input consists of $2MT + L$ planes of size $N \times N$.
- M planes indicate with 0 and 1 the positions of own pieces, M planes the positions of opponent pieces. The previous T positions are stacked to $2MT$ planes.
- L additional planes (constant values across the $N \times N$ positions) are used to represent the color of the player, the move number, the availability of special rules (e.g. castling in chess).
- Outputs of the policy networks are also represented as planes. For go there is only one plane that indicates where to place the piece. For chess there are 73 planes: one plane for each type of move (e.g. move 1 square north, move 3 squares south, move 5 square south-west etc.). A large value at a certain position in one of these planes, indicates that the piece at that position should be taken to perform the action indicated by the plane identity. Illegal moves are masked to have 0 probability.
- π_{θ} and v_{θ} use the same network body: a conv layer with batch-normalization followed by 19 residual blocks. Policy and value heads add some additional layers on top.

AlphaZero: Success Story

Games: Silver et al. 2018




TESLA Autopilot:

<https://youtu.be/j0z4FweCy4M?t=4802>

Physics: Optimizing quantum annealing schedules with Monte Carlo tree search enhanced with neural networks

<https://www.nature.com/articles/s42256-022-00446-y>

Notes



- Except of input and output representations, the same architecture and approach can be applied to reach better better performance than any computer program or human player in chess, shogi and go.
- Similar approaches can be applied to search problems that need to be solved efficiently. In a video from last year, Tesla has an example, where an AlphaZero approach was applied in a car's autopilot, to plan the trajectory from current position to a parking lot while avoiding many obstacles.
- In research we see the same approach applied to search problems in physics.

AlphaZero: Summary

- ▶ In each position of the game (for both players), MCTS runs for some time.
- ▶ The sub-tree below the actual next state is retained as the initial MCTS tree.
- ▶ In contrast to traditional Go/Chess/Shogi software there is no hand-crafting of promising tree expansions or position evaluations: selection of MCTS is focused on promising actions (according to learned prior action probability π_θ) and the value of leaf nodes v_θ is based on a learned position evaluation function.
- ▶ Prior action probabilities π_θ are trained with the cross-entropy loss to match the actual policy π of MCTS and the values v_θ are trained with the squared error loss to the actual outcome z of the game.
- ▶ AlphaZero does not need to learn the reward and transition model $R_{s \rightarrow s'}^a, P_{s \rightarrow s'}^a$; for MCTS it relies on the hard coded rules of the game.
What if the model is unknown?

Table of Contents

1. Model-Based Deep Reinforcement Learning

Model-Based Deep RL with Variational State Tabulation (VaST)
Other Model-Based Deep RL Methods

2. Monte Carlo Tree Search

3. AlphaZero

4. MuZero


5. Other Success Stories of Deep RL

What If the Model Is Unknown?

In high dimensional settings, like in the game Go, it seems hopeless to learn explicitly $R_{s \rightarrow s'}^a$ and $P_{s \rightarrow s'}^a$.

MuZero = AlphaZero with a learned latent representation of observations and learned dynamics model.


How MuZero uses its model to plan



- ▶ Root node of MCTS: encoded state $s_0 = h_\theta(o)$ instead of raw observation o .
- ▶ MCTS runs with learned dynamics $s', r = g_\theta(s, a)$ instead of known rules.
- ▶ Otherwise, MuZero works like AlphaZero with learned prior action probabilities p and values v , computed from latent states.


Schrittwieser et al. 2019

Notes



- The model consists of three connected components for representation, dynamics and prediction. The initial hidden state $s^{(0)} \in \mathbb{R}^N$ is obtained by passing current (and potentially a stack of past) observations (e.g. the Go board or Atari screen) into a representation function h_θ .
- Given a previous hidden state $s^{(k-1)}$ and a candidate action $a^{(k)}$, the dynamics function g_θ produces an immediate reward $r^{(k)}$ and a new hidden state $s^{(k)}$. The prediction of the immediate reward r is necessary for environments with intermediate rewards, like the Atari games.
- The policy $p^{(k)}$ and value function $v^{(k)}$ are computed from the hidden state $s^{(k)}$ by a prediction function $f_\theta = (\pi_\theta, v_\theta)$, as in AlphaZero.
- In contrast to VaST, DreamerV3 etc. the latent states s do not need to contain all information to reproduce the raw observations. “There is no direct constraint or requirement for the hidden state to capture all information necessary to reconstruct the original observation, drastically reducing the amount of information the model has to maintain and predict; nor is there any requirement for the hidden state to match the unknown, true state of the environment; nor any other constraints on the semantics of state. Instead, the hidden states are free to represent state in whatever way is relevant to predicting current and future values and policies. Intuitively, the agent can invent, internally, the rules or dynamics that lead to most accurate planning.” This is an advantage when the raw observation contains information that is irrelevant for the task.

How MuZero trains its model



Model

$$\left. \begin{aligned} s^0 &= h_{\theta}(o_1, \dots, o_t) \\ r^k, s^k &= g_{\theta}(s^{k-1}, a^k) \\ \mathbf{p}^k, v^k &= f_{\theta}(s^k) \end{aligned} \right\} \mathbf{p}^k, v^k, r^k = \mu_{\theta}(o_1, \dots, o_t, a^1, \dots, a^k)$$

Search

$$\nu_t, \pi_t = MCTS(s_t^0, \mu_{\theta})$$

$$a_t \sim \pi_t$$

Learning Rule

$$\mathbf{p}_t^k, v_t^k, r_t^k = \mu_{\theta}(o_1, \dots, o_t, a_{t+1}, \dots, a_{t+k})$$

$$z_t = \begin{cases} u_T & \text{for games} \\ u_{t+1} + \gamma u_{t+2} + \dots + \gamma^{n-1} u_{t+n} + \gamma^n v_{t+n} & \text{for general MDPs} \end{cases}$$

$$l_t(\theta) = \sum_{k=0}^K l^r(u_{t+k}, r_t^k) + l^v(z_{t+k}, v_t^k) + l^p(\pi_{t+k}, p_t^k) + c \|\theta\|^2$$

Losses

$$l^r(u, r) = \begin{cases} 0 & \text{for games} \\ \phi(u)^T \log \mathbf{r} & \text{for general MDPs} \end{cases}$$

$$l^v(z, q) = \begin{cases} (z - q)^2 & \text{for games} \\ \phi(z)^T \log \mathbf{q} & \text{for general MDPs} \end{cases}$$

$$l^p(\pi, p) = \pi^T \log \mathbf{p}$$

Schrittwieser et al. 2019




Notes


- A trajectory is sampled from the replay buffer. For the initial step, the representation function h receives as input the past observations $\sigma_1, \dots, \sigma_t$ from the selected trajectory.
- The model is subsequently unrolled recurrently for K steps. At each step k , the dynamics function g receives as input the hidden state $s^{(k-1)}$ from the previous step and the real action a_{t+k} .
- The parameters of the representation, dynamics and prediction functions are jointly trained, end-to-end by backpropagation-through-time, to predict three quantities: the prior action selection policy $p^{(k)} \approx \pi_{t+k}$ should approximate the MCTS policy π_{t+k} in time-step $t+k$, the value $v^{(k)} \approx z_{t+k}$ should approximate the return z_{t+k} (u_T indicates win or loss in a board game and u_t the intermediate rewards in general), and the immediate reward $r^{(t+k)} \approx u_{t+k}$ should approximate the actual immediate reward u_{t+k} .

MuZero: Success Story


Chess







Shogi



Go




Atari



Schrittwieser et al. 2019

Notes



Benchmark lines for Chess, Shogi & Go: AlphaZero.

For Atari: R2D2, a model-free agent that was state-of the art in 2019 (solid lines: mean, dashed lines: median)

MuZero: Summary

- ▶ Instead of relying on a given reward and transition model, MuZero learns an encoding of observations to states h_θ and a transition function from states to next states g_θ .
- ▶ Like AlphaZero, it relies on MCTS with a learned function f_θ that estimates the action probability π_θ and value v_θ of a state.
- ▶ In contrast to VaST, WoldModels and SimPLe, the latent state of MuZero is not trained in an autoencoder approach, i.e. o_t is not necessarily reconstructable from s_t .

Replay Buffers, Models, Policies


1. **Model-Free RL:** use raw experiences (s, a, r, s') to train policy $\pi(a|s)$ or $Q(s, a)$.
2. **Model-Free RL with Replay Buffer:** store raw experiences (s, a, r, s') in buffer M and train policy $\pi(a|s)$ or $Q(s, a)$ with samples from buffer.
3. **Model-Based RL with Background Learning:** use raw experiences (s, a, r, s') to update model $P_{s \rightarrow s'}^a$ and $R_{s \rightarrow s'}^a$ and use samples from model to train policy $\pi(a|s)$ or $Q(s, a)$.
4. **Model-Based RL with Replay Buffer and Background Learning:** store raw experiences (s, a, r, s') in replay buffer M and train a model (simulator) with samples from M ; use samples from the model to train policy $\pi(a|s)$ or $Q(s, a)$.
5. **Model-Based RL with Decision Time Planning:** use raw experiences or a replay buffer to update a model; use the model for planning at decision time.

Notes

1. e.g. standard policy gradient or Q-learning
2. e.g. DQN or ACER (policy gradient with replay buffer).
3. e.g. Dyna-Q (see Sutton and Barto Chapter 8.2)
4. e.g. VaST, Dreamer, etc.
5. e.g. MuZero

1. **Model-Free RL**: use raw experiences (s, a, r, s') to train policy $\pi(a|x)$ or $Q(x, a)$.
2. **Model-Free RL with Replay Buffer**: store raw experiences (s, a, r, s') in buffer M and train policy $\pi(a|x)$ or $Q(x, a)$ with samples from buffer.
3. **Model-Based RL with Background Learning**: use raw experiences (s, a, r, s') to update model $P_{s'|s, a}$ and $R_{s, s'}$ and use samples from model to train policy $\pi(a|x)$ or $Q(x, a)$.
4. **Model-Based RL with Replay Buffer and Background Learning**: store raw experiences (s, a, r, s') in replay buffer M and train a model (permeated with samples from M) use samples from the model to train policy $\pi(a|x)$ or $Q(x, a)$.
5. **Model-Based RL with Decision Time Planning**: use raw experiences or a replay buffer to update a model, use the model for planning at decision time.

Dyna Architecture



Notes




Figure 8.1 from Sutton and Barto: The general Dyna Architecture. Real experience, passing back and forth between the environment and the policy, affects policy and value functions in much the same way as does simulated experience generated by the model of the environment.

Quiz

Which statement is correct?

- In AlphaZero's MCTS the actual actions are sampled from the prior action probability $P(s, a)$.
- In AlphaZero's MCTS the prior action probability is trained to be as close as possible to the actual policy of MCTS.
- MCTS and value iteration serve the same purpose in model-based RL: both methods are used to find a policy.

Table of Contents

1. Model-Based Deep Reinforcement Learning

Model-Based Deep RL with Variational State Tabulation (VaST)

Other Model-Based Deep RL Methods


2. Monte Carlo Tree Search

3. AlphaZero

4. MuZero

5. Other Success Stories of Deep RL


Agent 57



- ▶ Super-human on 57 Atari games.
- ▶ Prioritized replay for efficient learning from the replay buffer.
- ▶ LSTM to deal with partial observability; similar to Recurrent Replay Distributed QDN (R2D2).
- ▶ Intrinsic motivation for efficient exploration; similar to Never Give Up (NGU).
 $Q(s, a) = Q^r(s, a) + \beta Q^i(s, a)$, where Q^r is trained with standard reward and the intrinsic part Q^i is given a novelty signal as reward (think of $1/(\text{visitation count})$).
- ▶ Meta-controller to dynamically adjust the level of exploration and discount factor.

[Badia et al., 2020]

AlphaStar



- ▶ Grandmaster level in StarCraft II.
- ▶ Supervised pre-training to imitate human players.
- ▶ Transformer & LSTM.
- ▶ V-trace (version of actor-critic with off-policy corrections)
- ▶ distributed training (multiple TPUs)

[Vinyals et al., 2019]


OpenAI Five



- ▶ Defeated world champion in Dota 2
- ▶ LSTM with shared weights for all 5 “heros”
- ▶ PPO
- ▶ distributed training (thousands of GPUs for many months)

[OpenAI et al., 2019a]

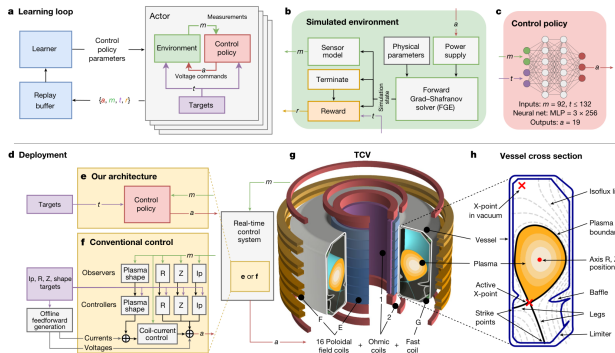
Solving Rubik's Cube with a Robot Hand



- ▶ Solves Rubik's cube with a robot hand
- ▶ Vision model based on ResNet50 & LSTM controller
- ▶ PPO
- ▶ Automated Domain Randomization: training on simulator with a distribution of environments (slightly different physics).

[OpenAI et al., 2019b]

Magnetic control of tokamak plasmas



- ▶ Controls nuclear fusion plasma
- ▶ Maximum a posteriori policy optimization (MPO); an actor-critic algorithm.
- ▶ Training on simulator.

[Degrave et al., 2022]

References I



Badia, A. P., Piot, B., Kapturowski, S., Sprechmann, P., Vitvitskyi, A., Guo, Z. D., and Blundell, C. (2020).

Agent57: Outperforming the Atari human benchmark.

In Ili, H. D. and Singh, A., editors, [Proceedings of the 37th International Conference on Machine Learning](#), volume 119 of [Proceedings of Machine Learning Research](#), pages 507–517. PMLR.



Cornel, D., Gerstner, W., and Brea, J. (2018).

Efficient model-based deep reinforcement learning with variational state tabulation.

In Dy, J. and Krause, A., editors, [Proceedings of the 35th International Conference on Machine Learning](#), volume 80 of [Proceedings of Machine Learning Research](#), pages 1057–1066, Stockholmsmässan, Stockholm Sweden. PMLR.



Degrave, J., Felici, F., Buchli, J., Neunert, M., Tracey, B., Carpanese, F., Ewalds, T., Hafner, R., Abdolmaleki, A., de las Casas, D., Donner, C., Fritz, L., Galperti, C., Huber, A., Keeling, J., Tsimpoukelli, M., Kay, J., Merle, A., Moret, J.-M., Noury, S., Pesamosca, F., Pfau, D., Sauter, O., Sommariva, C., Coda, S., Duval, B., Fasoli, A., Kohli, P., Kavukcuoglu, K., Hassabis, D., and Riedmiller, M. (2022).

Magnetic control of tokamak plasmas through deep reinforcement learning.

[Nature](#), 602(7897):414–419.



Ha, D. and Schmidhuber, J. (2018).

World Models.

[ArXiv e-prints](#).

References II



Hafner, D., Lillicrap, T., Norouzi, M., and Ba, J. (2020).
Mastering Atari with Discrete World Models.
[arXiv e-prints](#), page arXiv:2010.02193.



Hafner, D., Pasukonis, J., Ba, J., and Lillicrap, T. (2023).
Mastering Diverse Domains through World Models.
[arXiv e-prints](#), page arXiv:2301.04104.



Kaiser, L., Babaeizadeh, M., Milos, P., Osinski, B., Campbell, R. H., Czechowski, K., Erhan, D., Finn, C., Kozakowski, P., Levine, S., Mohiuddin, A., Sepassi, R., Tucker, G., and Michalewski, H. (2019).
Model-Based Reinforcement Learning for Atari.
[arXiv e-prints](#), page arXiv:1903.00374.



OpenAI, :, Berner, C., Brockman, G., Chan, B., Cheung, V., Dębiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Józefowicz, R., Gray, S., Olsson, C., Pachocki, J., Petrov, M., Pinto, H. P. d. O., Raiman, J., Salimans, T., Schlatter, J., Schneider, J., Sidor, S., Sutskever, I., Tang, J., Wolski, F., and Zhang, S. (2019a).
Dota 2 with Large Scale Deep Reinforcement Learning.
[arXiv e-prints](#), page arXiv:1912.06680.

References III



OpenAI, Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., Paino, A., Plappert, M., Powell, G., Ribas, R., Schneider, J., Tezak, N., Tworek, J., Welinder, P., Weng, L., Yuan, Q., Zaremba, W., and Zhang, L. (2019b). Solving Rubik's Cube with a Robot Hand. [arXiv e-prints](#), page arXiv:1910.07113.



Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., Guez, A., Lockhart, E., Hassabis, D., Graepel, T., Lillicrap, T., and Silver, D. (2019). Mastering Atari, Go, Chess and Shogi by Planning with a Learned Model. [arXiv e-prints](#), page arXiv:1911.08265.



Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., and et al. (2016). Mastering the game of go with deep neural networks and tree search. [Nature](#), 529(7587):484–489.



Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., and et al. (2018). A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. [Science](#), 362(6419):1140–1144.

References IV



Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., Oh, J., Horgan, D., Kroiss, M., Danihelka, I., Huang, A., Sifre, L., Cai, T., Agapiou, J. P., Jaderberg, M., Vezhnevets, A. S., Leblond, R., Pohlen, T., Dalibard, V., Budden, D., Sulsky, Y., Molloy, J., Paine, T. L., Gulcehre, C., Wang, Z., Pfaff, T., Wu, Y., Ring, R., Yogatama, D., Wünsch, D., McKinney, K., Smith, O., Schaul, T., Lillicrap, T., Kavukcuoglu, K., Hassabis, D., Apps, C., and Silver, D. (2019).
Grandmaster level in starcraft ii using multi-agent reinforcement learning.
[Nature](#), 575(7782):350–354.