Autonomous and Adaptive Systems

Policy Gradient Methods First Part

Mirco Musolesi

mircomusolesi@acm.org

Policy Gradient Methods

- In the previous lectures we discussed methods that were based on the calculation of action values. We refer to them as as action-value methods.
- We learned the values of the actions and then select actions based on the estimated action values. The learning can happen with the optimisation of the policy at the same time (e.g., control problem), but we need the values of the actions in the first place.
- We now consider a different type of methods that learn instead a parametrised policy that can select actions without consulting a value function.
- A value function can be used to learn the policy parameter, but it is not required for action selection.

Policy Gradient Methods

- We use the notation $\theta \in \mathbb{R}^{d'}$ for the policy's parameter vector.
- We indicate the probability that action a is taken at time t given that the environment is in state s a time t with parameters θ as follows:

$$\pi(a \mid s, \theta) = Pr\{A_t = a \mid S_t = s, \theta_t = \theta\}$$

▶ If a method uses a learned value function as well, the value function's weight vector is denoted with $\mathbf{w} \in \mathbb{R}^d$.

Policy Gradient Methods

- We consider methods for learning the policy parameter based on the gradient of some scalar performance measure $J(\theta)$ with respect to the policy parameter.
- The goal of these methods is to maximise performance, so their updates approximate gradient ascent in J as follows:

 $\theta_{t+1} = \theta_t + \alpha \nabla \hat{J}(\theta_t)$

where $\nabla \hat{J}(\theta_t) \in \mathbb{R}^d$ is a stochastic estimate whose expectation approximates the gradient of the performance measure with respect to the argument θ_t .

We refer to all the methods that follow this schema as *policy gradient methods*.

Policy Approximation

- In policy gradient methods, the policy can be parameterised in several possible way.
- ▶ The only constraint is that $\pi(a \mid s, \theta)$ is differentiable, i.e., as long as $\nabla \pi(a \mid s, \theta)$ with respect to θ exists and it is finite for all $s \in S$, $a \in \mathscr{A}(s)$ with $\theta \in \mathbb{R}^{d'}$.
- We need to ensure exploration and, therefore, one goal is to make sure that the policy will never become deterministic, i.e. that π(a | s, θ) = (0,1) for all s, a, θ.

Policy Parametrisation and Parametrised Action Preferences

- A standard way for parametrisation of policies is to derive first parametrised numerical preferences $h(a \mid s, \theta) \in \mathbb{R}$ for each state-action pair.
- The actions with the highest preferences in each state are given the highest probabilities of being selected.
- A typical mapping between the preferences $h(a | s, \theta)$ and the probabilities $\pi(a | s, \theta)$ is obtained through the use of an exponential softmax distribution:

$$\pi(a \,|\, s, \theta) \doteq \frac{e^{h(a \mid s, \theta)}}{\sum_{b} e^{h(b \mid s, \theta)}}$$

- ▶ We call this type of policy parametrisation softmax in action preferences.
- ▶ Note that we are not deriving the value functions and then apply a policy (let's say *ε*-greedy). We are deriving directly the probability distributions of the actions given the states, i.e., the policy itself.

Policy Parametrisation and Parametrised Action Preferences

- The action preferences themselves $h(a \mid s, \theta)$ can be parametrised arbitrarily.
- For example, they might be computed by a deep artificial neural network (ANN), where θ is the vector of all the connection weights.
- This is used for example in the AlphaGo system.

ARTICLE

doi:10.1038/nature16961

Mastering the game of Go with deep neural networks and tree search

David Silver¹*, Aja Huang¹*, Chris J. Maddison¹, Arthur Guez¹, Laurent Sifre¹, George van den Driessche¹, Julian Schrittwieser¹, Ioannis Antonoglou¹, Veda Panneershelvam¹, Marc Lanctot¹, Sander Dieleman¹, Dominik Grewe¹, John Nham², Nal Kalchbrenner¹, Ilya Sutskever², Timothy Lillicrap¹, Madeleine Leach¹, Koray Kavukcuoglu¹, Thore Graepel¹ & Demis Hassabis¹

The game of Go has long been viewed as the most challenging of classic games for artificial intelligence owing to its enormous search space and the difficulty of evaluating board positions and moves. Here we introduce a new approach to computer Go that uses 'value networks' to evaluate board positions and 'policy networks' to select moves. These deep neural networks are trained by a novel combination of supervised learning from human expert games, and reinforcement learning from games of self-play. Without any lookahead search, the neural networks play Go at the level of state-of-the-art Monte Carlo tree search programs that simulate thousands of random games of self-play. We also introduce a new search algorithm that combines Monte Carlo simulation with value and policy networks. Using this search algorithm, our program AlphaGo achieved a 99.8% winning rate against other Go programs, and defeated the human European Go champion by 5 games to 0. 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.

All games of perfect information have an optimal value function, $v^*(s)$, which determines the outcome of the game, from every board position or state s under perfect play by all players. These games may be solved

policies^{13–15} or value functions¹⁶ based on a linear combination of input features.

Recently deen convolutional neural networks have achieved unprec-

Advantages of Using Policy Approximation according to Softmax in Action Preferences

- One advantage of parameterising policies according to the softmax in action preferences is that the approximate policy can approach a deterministic policy.
- In fact, with *ϵ*-greedy selection over action values, there is always a probability *ϵ* of selecting a random action.
- One possibility is to use softmax distribution on the action values, but this will not allow to reach a deterministic policy.
 - Action values will always differ and, therefore, there will be always a non-null probability of selecting a different action.
 - Action preferences are different since they do not approach specific values: instead they are driven to produce the optimal stochastic policy.
 - If the optimal policy is deterministic, the preferences of the optimal actions will be driven infinitely higher than all suboptimal actions (the output of the softmax will be then very close to 1, i.e., close to determinism).

Advantages of Using Policy Approximation according to Softmax in Action Preferences

- The second advantage of parameterising policies according to the softmax in action preferences is that it enables the selection of actions with arbitrary probabilities.
- In some situations the best approximate policy might be stochastic, especially in games of imperfect information.
- Action-value methods do not have a natural way of finding stochastic optimal policies; instead, policy approximation methods can.
- It is also worth noting that in some cases, policy approximation might be easier than value approximation.

Advantages of Using Policy Approximation according to Softmax in Action Preferences

- Finally, there is an important "theoretical" advantage. With continuous policy parameterisation, the action probabilities change smoothly as a function of the learned parameter.
- Indeed, in *e*-greedy selection the action probabilities might change dramatically for a small change in the estimated action values, if that change results in a different action having the maximal value.

Policy Gradient Theorem

We consider episodic learning and we define the performance measure as the value at the start of the episode.

We can simplify the notation by assuming that each episode starts in a non-random state s_0 .

Formally:

 $J(\theta) \doteq v_{\pi_{\theta}}(s_0)$

where $v_{\pi_{\theta}}$ is the true value function for π_{θ} , the policy determined by θ .







Policy Gradient Theorem

- Key question: how can we estimate the performance gradient with respect to the policy parameter when the gradient depends on the unknown effect of policy changes on the state distribution?
- Fortunately, the policy gradient theorem provides an analytic expression for the gradient of performance with respect to the policy parameter that does not involve the derivative of the state distributions.
- ▶ In particular, the policy gradient theorem says that:

$$\nabla J(\theta) \propto \sum_{s} \mu(s) \sum_{a} q_{\pi}(s, a) \nabla \pi(a \mid s, \theta)$$

where the gradients are column vectors of partial derivatives with respect to the components of θ and π denotes the policy corresponding to parameter vector θ . $\mu(s)$ is the on-policy distribution under π (i.e., the fraction of time spent in each state normalised to sum to 1).

▶ If you are interested, you can find the proof in Chapter 13 of Sutton and Barto's book.

- REINFORCE is a Monte Carlo Policy Gradient *control* algorithm, i.e., the update leads to the optimal policy.
- ▶ It is based on stochastic gradient ascent as discussed above.
 - We need a way to obtain samples such that the expectation of the sample gradient is proportional to the actual gradient of the performance measure as a function of the parameter.
 - The sample gradients need only be proportional to the gradient because any constant of proportionality can be absorbed into the step size α , which is otherwise arbitrary.
 - ▶ The policy gradient theorem gives an exact expression proportional to the gradient.
 - We need to find a way of sampling whose expectation equals or approximates this expression.

Notice that the right-hand side of the policy gradient theorem is a sum over states weighted by how often the states occur under the target policy π. If π is followed, the states will be encountered in these proportions.

More formally:

$$\nabla J(\theta) \propto \sum_{s} \mu(s) \sum_{a} q_{\pi}(s, a) \nabla \pi(a \mid s, \theta)$$
$$= \mathbb{E}_{\pi} [\sum_{a} q_{\pi}(S_{t}, a) \nabla \pi(a \mid S_{t}, \theta)]$$

▶ In theory we can stop here and instantiate the stochastic gradient algorithm:

$$\theta_{t+1} = \theta_t + \alpha \, \nabla \hat{J}(\theta_t)$$

as follows:

$$\theta_{t+1} \doteq \theta_t + \alpha \sum_a \hat{q}(S_t, a, \mathbf{w}) \nabla \pi(a \mid S_t, \theta)$$

where \hat{q} is a learned approximation to q_{π} .

> This is called an all-action method because it involves the updates of all the actions.

• However, we are interested in an algorithm whose update at time t involves only A_t , the action taken at time t. This is called REINFORCE algorithm proposed by Williams in 1992.

© 1992 Kluwer Academic Publishers, Boston. Manufactured in The Netherlands.

Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning

RONALD J. WILLIAMS College of Computer Science, 161 CN, Northeastern University, 360 Huntington Ave., Boston, MA 02115

Abstract. This article presents a general class of associative reinforcement learning algorithms for connectionist networks containing stochastic units. These algorithms, called REINFORCE algorithms, are shown to make weight adjustments in a direction that lies along the gradient of expected reinforcement in both immediate-reinforcement tasks and certain limited forms of delayed-reinforcement tasks, and they do this without explicitly computing gradient estimates or even storing information from which such estimates could be computed. Specific examples of such algorithms are presented, some of which bear a close relationship to certain existing algorithms while others are novel but potentially interesting in their own right. Also given are results that show how such algorithms can be naturally integrated with backpropagation. We close with a brief discussion of a number of additional issues surrounding the use of such algorithms, including what is known about their limiting behaviors as well as further considerations that might be used to help develop similar but potentially more powerful reinforcement learning algorithms.

We have the following derivation

$$\begin{aligned} \nabla J(\theta) &\propto \mathbb{E}_{\pi} \left[\sum_{a} q_{\pi}(S_{t}, a) \nabla \pi(a \mid S_{t}, \theta) \right] \\ &= \mathbb{E}_{\pi} \left[\sum_{a} \pi(a \mid S_{t}, \theta) q_{\pi}(S_{t}, a) \frac{\nabla \pi(a \mid S_{t}, \theta)}{\pi(a \mid S_{t}, \theta)} \right] \\ &= \mathbb{E}_{\pi} \left[\mathbb{E}_{\pi} \left[q_{\pi}(S_{t}, A_{t}) \right] \frac{\nabla \pi(A_{t} \mid S_{t}, \theta)}{\pi(A_{t} \mid S_{t}, \theta)} \right] \qquad \text{since } \mathbb{E}_{\pi} \left[\mathbb{E}_{\pi} [x] \right] = \mathbb{E}_{\pi} [x] \\ &= \mathbb{E}_{\pi} \left[q_{\pi}(S_{t}, A_{t}) \frac{\nabla \pi(A_{t} \mid S_{t}, \theta)}{\pi(A_{t} \mid S_{t}, \theta)} \right] \\ &= \mathbb{E}_{\pi} \left[G_{t} \frac{\nabla \pi(A_{t} \mid S_{t}, \theta)}{\pi(A_{t} \mid S_{t}, \theta)} \right] \\ &= \mathbb{E}_{\pi} \left[G_{t} \frac{\nabla \pi(A_{t} \mid S_{t}, \theta)}{\pi(A_{t} \mid S_{t}, \theta)} \right] \end{aligned}$$

where G_t is the return.

Note that the last equality is true because
$$\mathbb{E}_{\pi}[G_t | S_t, A_t] = q_{\pi}(S_t, A_t)$$
 and $\nabla \ln x = \frac{\nabla x}{x}$.

- The final expression in brackets can be used for the update. It is a quantity that can be sampled at each step.
- Think about it in a different way: you are moving around your objective and on average your correction will lead you close to your real objective (the maximum in this case).
- Since you repeat stochastically this correction you end up with a correction that is close to the expectation of the gradient. For this reason we can use this sample to instantiate the generic stochastic gradient ascent.

As in the stochastic gradient descent, in the stochastic gradient ascent, you repeat the "correction" many times in order to reach the maximum.

$$\theta_{t+1} = \theta_t + \alpha \nabla \hat{J}(\theta_t)$$

From the derivation above we have the following REINFORCE update:

$$\theta_{t+1} \doteq \theta_t + \alpha G_t \nabla \ln \pi(A_t | S_t, \theta_t)$$

Remember again that $\nabla J(\theta)$ is proportional to the update value, not equal, but we have the parameter α so it does not really matter.





Input: a differentiable policy parametrisation $\pi(a \mid s, \theta)$ Algorithm parameter: step size $\alpha > 0$ Initialise policy parameter $\theta \in \mathbb{R}^{d'}$ Loop forever (for each episode): Generate an episode $S_0, A_0, R_1, \ldots, S_{T-1}, A_{T-1}, R_T$ following π Loop for each step of the episode t = 0, 1, ..., T - 1: $G_t \leftarrow \sum_{k=1}^{T} R_k$ k=t+1 $\theta \leftarrow \theta + \alpha G_t \nabla ln \pi(A_t | S_t, \theta)$

Deep Neural Networks



Autonomous and Adaptive Systems 2024-2025

Gradient-based Optimisation



Stochastic Gradient **Descent**

- Recall the stochastic gradient *descent:*
 - > Draw a batch of training example x and corresponding targets $y_{\textit{target}}$
 - Run the network on \mathbf{X} (forward pass) to obtain predictions \mathbf{y}_{pred} .
 - Compute the loss of the network on the batch, a measure of the mismatch between \mathbf{y}_{pred} and \mathbf{y}_{target} .
 - Compute the gradient of the loss with regard to the network's parameters (backward pass).

Stochastic Gradient **Descent**

Move the parameters in the opposite direction from the gradient with:

$$\theta_j \leftarrow \theta_j - \Delta \theta_j = \theta_j - \eta \frac{\partial J}{\partial \theta_j}$$

where J is the loss (cost) function.

• If you have a batch of samples of dimension k:

$$\theta_j \leftarrow \theta_j - \Delta \theta_j = \theta_j - \eta \ average(\frac{\partial J_k}{\partial \theta_j})$$

for all the k samples of the batch.

Stochastic Gradient Ascent

Similarly in our case, we move the parameters in the same direction as the gradient with:

$$\theta_j \leftarrow \theta_j + \Delta \theta_j = \theta_j + \eta \frac{\partial J}{\partial \theta_j}$$

where J is the loss (cost) function.

If you have a batch of samples of dimension k you can think about moving towards your actual objective using this formula:

$$\theta_j \leftarrow \theta_j + \Delta \theta_j = \theta_j + \eta \ average\left(\frac{\partial J_k}{\partial \theta_j}\right)$$

for all the k samples of the batch.

Implementing REINFORCE with Artificial Neural Networks/Deep Learning

- We said that an artificial neural network can be used to implement REINFORCE. But how can we do this in practice in a package like TensorFlow or Keras?
- ▶ The key problem is to select a loss function that can be mapped to REINFORCE.
 - We can talk about *compatibility* between Artificial Neural Networks and REINFORCE in a sense (see also theory in Williams 1992).
- From a practical point of view, we want to find a way of exploiting backpropagation for updating the weights using the "machinery" that is offered for example by the existing frameworks.
- We want to adapt a "machinery" that is built for stochastic gradient descent to a stochastic gradient ascent. As you can imagine the trick would be to do a correction in the opposite direction.

Implementing REINFORCE with Artificial Neural Networks/Deep Learning

• Let us consider again our formula: $\theta_{t+1} = \theta_t + \alpha \nabla \hat{J}(\theta_t)$

- Our loss function is $J(\theta)$. Our goal is to correct each weight by $\nabla J(\theta)$ using the correction $\nabla J(\theta)$ for each weight.
- The value can be considered against the value 0, i.e., the difference $G_t ln\pi 0$.
- However, since this is stochastic gradient *ascent* we will take the opposite of this quantity. $J'(\theta)$, the loss function we are going to use in Tensorflow, is the opposite of $J(\theta)$.
- ► More formally:

$$J'(\theta) = -(G_t ln\pi - 0) = -G_t ln\pi$$

Implementing REINFORCE with Artificial Neural Networks/Deep Learning

And of course the gradient of this quantity is

 $\nabla J'(\theta) = -G_t \nabla \ln \pi$

- Essentially, for example in Keras (TensorFlow) it would be sufficient to set the loss function to a custom loss function equal to $-G_t ln\pi$.
- Then Keras will take care of the stochastic gradient ascent, i.e., optimisation problem.

Issues with REINFORCE

- As a stochastic gradient method, REINFORCE has good theoretical convergence properties.
- By construction, the expected update over an episode in the same direction as the performance gradient.
- This assures an improvement in expected performance for sufficient small α and convergence to a local optimum under standard stochastic approximation conditions for decreasing α .
- However, since it is a Monte Carlo method, REINFORCE suffers from high variance and this might lead to slow learning.
- One way of dealing with this problem is to use baselines and actor-critic methods.

The policy gradient theorem can be generalised to include a comparison of the action value to an arbitrary baseline b(s):

$$\nabla J(\theta) \propto \sum_{s} \mu(s) \sum_{a} \left(q_{\pi}(s, a) - b(s) \right) \nabla \pi(a \,|\, s, \theta)$$

- The baseline can be any function, even a random variable, as long as it does not vary with a.
- The equation remains valid because the subtracted quantity is zero:

$$\sum_{a} b(s) \nabla \pi(a \mid s, \theta) = b(s) \nabla \sum_{a} \pi(a \mid s, \theta) = b(s) \nabla 1 = 0$$

- The policy gradient theorem with baseline can be used to derive an update rule using similar steps.
- ▶ The update for REINFORCE with baseline is as follows:

$$\theta_{t+1} \doteq \theta_t + \alpha(G_t - b(S_t)) \frac{\nabla \pi(A_t | S_t, \theta_t)}{\pi(A_t | S_t, \theta_t)}$$
$$= \theta_t + \alpha(G_t - b(S_t)) \nabla \ln \pi(A_t | S_t, \theta_t)$$

- The use of the baseline leaves the expected value unchanged, but it can have a large effect on its variance.
- The value of b(s) can be a random number but it is not ideal.
 - In some states all actions have high values and we need a high baseline to differentiate the higher valued actions from the less highly values ones. In other states all actions will have low values and a low baseline is appropriate.

A natural choice is to learn the state value at the same time;

 $b(S_t) = \hat{v}(S_t, \mathbf{w})$

where $\mathbf{w} \in \mathbb{R}^m$ is a weight vector learned using, for example, another deep neural network.

- Because REINFORCE is a Monte Carlo method it makes to use a Monte Carlo method also for learning the state-value weights w.
- This method will be based on two step-size denoted α^{θ} and α^{w} .

Input: a differentiable policy parametrisation $\pi(a \mid s, \theta)$ and a differentiable state-value function parametrisation $\hat{v}(s, \mathbf{w})$.

Algorithms parameters: steps size $\alpha^{\theta} > 0$ and $\alpha^{w} > 0$

Initialise policy parameter
$$\theta \in \mathbb{R}^{d'}$$
 and state-value weights $\mathbf{w} \in \mathbb{R}^{d}$

Loop forever (for each episode):

Generate an episode $S_0, A_0, R_1, \ldots, S_{T-1}, A_{T-1}, R_T$ following π

Loop for each step of the episode t = 0, 1, ..., T - 1:

$$\begin{aligned} G_t &\leftarrow \sum_{k=t+1}^T R_k \\ \delta &\leftarrow G_t - \hat{v}(S_t, \mathbf{w}) \\ \mathbf{w} &\leftarrow \mathbf{w} + \alpha^{\mathbf{w}} \delta \,\nabla \, \hat{v}(S_t, \mathbf{w}) \\ \theta &\leftarrow \theta + \alpha^{\theta} (G_t - \hat{v}(S_t, \mathbf{w})) \,\nabla \ln \pi (A_t \,|\, S_t, \theta) \end{aligned}$$

Input: a differentiable policy parametrisation $\pi(a \mid s, \theta)$ m a differentiable state-value function parametrisation $\hat{v}(s, \mathbf{w})$.

Algorithms parameters: steps size $\alpha^{\theta} > 0$ and $\alpha^{w} > 0$

Initialise policy parameter
$$\theta \in \mathbb{R}^{d'}$$
 and state-value weights $\mathbf{w} \in \mathbb{R}^{d}$

Loop forever (for each episode):

Generate an episode $S_0, A_0, R_1, \ldots, S_{T-1}, A_{T-1}, R_T$ following π

Loop for each step of the episode t = 0, 1, ..., T - 1:

$$\begin{split} G_t &\leftarrow \sum_{k=t+1}^T R_k \\ \delta &\leftarrow G_t - \hat{v}(S_t, \mathbf{w}) \\ \mathbf{w} &\leftarrow \mathbf{w} + \alpha^{\mathbf{w}} \delta \nabla \hat{v}(S_t, \mathbf{w}) \\ \theta &\leftarrow \theta + \alpha^{\theta} (G_t - \hat{v}(S_t, \mathbf{w})) \nabla \ln \pi (A_t | S_t, \theta) \end{split}$$

References

Chapter 13 of Barto and Sutton. Introduction to Reinforcement Learning. Second Edition. MIT Press 2018.