



Theoretical Biophysics

A Computational Approach

Concepts, Models, Methods and Algorithms

Learning

Dieter W. Heermann

April 7, 2020

Heidelberg University

1. Introduction

2. Markov Decision Process

3. Deep Learning

■ Neural Networks

- Backpropagation
- Classification

■ Hidden Markov Models

- Forward algorithm
- Viterbi Algorithm
- n-gram Models
- Forward-Backward Algorithm

■ Reinforcement Learning

- Monte-Carlo Policy Evaluation
- Advantage Function
- Q-Learning
- Actor-Critic policy gradient algorithm
- ϵ Greedy Strategy
- Value Function Fitting
- Continuous State and Action Space

4. Genetic Algorithms


5. Exercises

6. Bibliography


7. Index



Introduction



Neural network
application to gene
finding. Image taken
from [1].



Cell imaging analysis with neural networks.
Image taken from [2].

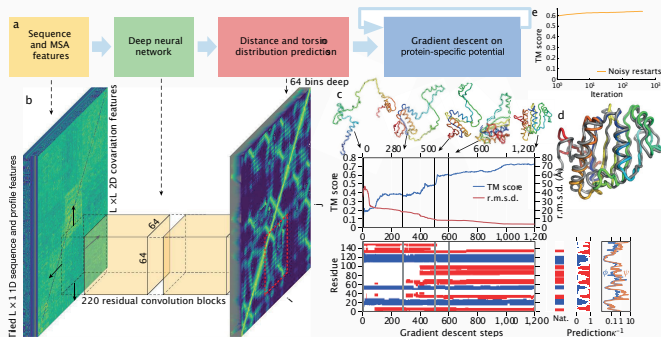


Fig. 2 | The folding process illustrated for CASP13 target T0986. **a**, Steps of structure prediction. The neural network predicts the entire $L \times L$ distogram based on MSA features, accumulating separate predictions for 64×64 -residue regions. **b**, One iteration of gradient descent (1,200 steps) is shown, with the TM score and root mean square deviation (r.m.s.d.) plotted against step number with five snapshots (in grey). **c**, The average (across the test set; 377) TM score of the lowest- ΔG structure. The secondary structure (from SS1) is also shown (helix in blue, potential structure against the number of repeats of gradient descent per strand in red) along with the native secondary structure (Nat.), the secondary target (log scale).

structure prediction probabilities of the network and the uncertainty in torsion angle predictions (σ of the von Mises distributions fitted to the predictions for ϕ and ψ). While each step of gradient descent greedily lowers the potential, large global conformational changes are effected, resulting in a well-packed chain. The final first submission overlaid on the native structure (grey). **d**, The secondary structure (from SS1) is also shown (helix in blue, potential structure against the number of repeats of gradient descent per strand in red) along with the native secondary structure (Nat.), the secondary target (log scale).

Figure 1: Image taken from: Improved protein structure prediction using potentials from deep learning Nature 2020 [3].

(Machine) learning can be roughly categorized into supervised and unsupervised. Typical techniques include:

- supervised methods:
 - Artificial Neural Network,
 - Support Vector Machines and linear classifiers
 - Bayesian Statistics,
 - k-Nearest Neighbors,
 - Hidden Markov Model
 - Decision Trees
- un-supervised methods
 - Autoencoders,
 - Expectation Maximization,
 - Self-Organizing Maps,
 - k-Means
 - Fuzzy clustering
 - Density-based clustering.

Methods developed and applications of machine learning in biophysical problems [1] range from finding genes, as featured in the introduction to the analysis of images such as computer tomography spanning the entire variety of bio-biological problems.



Markov Decision Process

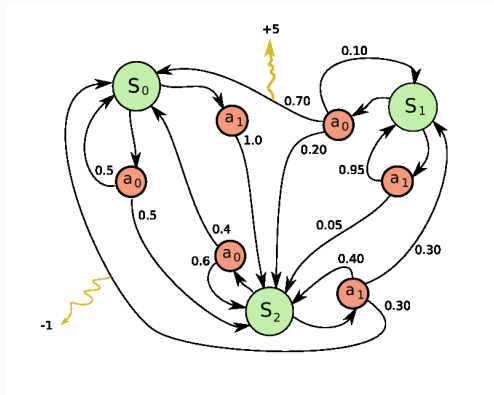


Figure 2: Markov decision process. Image taken from:

<https://towardsdatascience.com/reinforcement-learning-demystified-markov-decision-processes-part-1-bf00dda41690>.

Let S be a state space (a countable non-empty set), A be the action space (countable non-empty set of actions) and O the observation space. Let P_0 be a transition probability kernel that assigns to $(S = s, A = a) \in S \times A$ a probability measure over $S \times \mathbb{R} : P_0(\cdot | s, a)$.

A countable Markov Decision Process (MDP) is defined as a triple $M = (S, A, P_0)$. We further define a reward function

$$R(s, a) = \mathbb{E}[R | S = s, A = a] = \int_{\mathbb{R}} \sum_{s' \in S} R \cdot P_0(s', R | s, a) dR. \quad (1)$$

Furthermore let

$$R(s, a, s') = \mathbb{E}[R | S_t = s, A = a, S_{t+1} = s'] = \int_{\mathbb{R}} R \cdot P_0(s', R | s, a) dR. \quad (2)$$

A Markov Reward Process (MRP) is a Markov process with a reward function. Hence a tuple (S, P, R, γ) . γ is a discount factor, where $\gamma \in [0, 1]$.

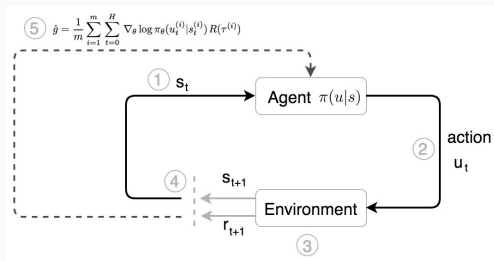


Figure 3: Markov decision process, Figure taken from:
<https://medium.com/@jonathan-hui/rl-policy-gradients-explained-9b13b688b146>.

From the definition it is clear that for a sequence $S_1, A_1 \dots, S_{t-1}, A_{t-1}, S_t, A_t, S_{t+1}, A_{t+1}$ we have

$$\mathbb{P}[S_{t+1}, R_{t+1} | S_t; A_t] = \mathbb{P}[S_{t+1}, R_{t+1} | S_1, A_1 \dots, S_{t-1}, A_{t-1}, S_t; A_t]. \quad (3)$$

Thus the sequence is Markovian

At each step t an agent:

- Receives observation o_t
- Receives (immediate) scalar reward R_t
- Executes action at a_t

The environment:

- Receives action a_t
- Emits observation O_{t+1}
- Emits scalar reward R_{t+1}

If $s_t = o_t$ the environment is fully observable. Let G_t be the total discounted rewards from time step t

$$G_t = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1} . \quad (4)$$



Deep Learning

Recall that the **gradient descent** method computes the gradient of the cost function J w.r.t. to a parameter θ

$$\theta = \theta - \alpha \cdot \nabla_{\theta} J(\theta) . \quad (5)$$

with the rate α . Choosing a proper learning rate can be difficult. Usually a learning rate schedule is used where α progressively decreases with the number of iterations. The challenge is that mostly the function that we want to minimize is highly non-convex so that there is a high probability to get trapped in numerous suboptimal local minima.

Stochastic gradient descent (SGD) performs a parameter update for each training example $x^{(i)}$ and label $y^{(i)}$

$$\theta = \theta - \alpha \cdot \nabla_{\theta} J(\theta; x^{(i)}; y^{(i)}) . \quad (6)$$

SGD performs frequent updates with a high variance that cause the objective function to fluctuate heavily.

Algorithm 1 Stochastic gradient descent (SGD)

```
1: for number of epochs do  
2:   randomly shuffle data  $(x, y)$   
3:   for number of data do  
4:      $\theta = \theta - \alpha \cdot \nabla_{\theta} J(\theta; x^{(i)}; y^{(i)})$   
5:   end for  
6: end for
```

Mini-batch gradient descent: update for every mini-batch of n training examples

$$\theta = \theta - \alpha \cdot \nabla_{\theta} J(\theta; x^{(i:i+n)}; y^{(i:i+n)}) . \quad (7)$$

Algorithm 2 Mini-batch gradient descent

```
1: for number of epochs do  
2:   randomly shuffle data  $(x, y)$   
3:   for batch in take out a batch from data of size  $m$  do  
4:      $\theta = \theta - \alpha \cdot \nabla_{\theta} J(\theta; x^{(i:i+n)}, y^{(i:i+n)})$   
5:   end for  
6: end for
```

Gradient descent optimization algorithms

SGD with momentum [4]. Let $\gamma < 1$ be the resistance, then

$$\begin{aligned}v_t &= \gamma v_{t-1} + \alpha \nabla_{\theta} J(\theta - \gamma v_{t-1}) \\ \theta &= \theta - v_t .\end{aligned}\tag{8}$$

The **Adaptive Moment Estimation (Adam)** [5] is often used in packages like Tensorflow [6] for the gradient descent:

$$\begin{aligned}m_t &= \beta_1 m_{t-1} + (1 - \beta_1) g_t \\v_t &= \beta_2 v_{t-1} + (1 - \beta_2) g_t^2.\end{aligned}\tag{9}$$

Adam stores an exponentially decaying average of past squared gradients v_t and keeps an exponentially decaying average of past gradients m_t . They are estimates of the first moment and the second moment of the gradients respectively producing a bias ($\beta_1 \approx 1$ and $\beta_2 \approx 1$). This being opposed by

$$\begin{aligned}\hat{m}_t &= \frac{m_t}{1 - \beta_1^t} \\ \hat{v}_t &= \frac{v_t}{1 - \beta_2^t}\end{aligned}\tag{10}$$

leading to the actual gradient descent

$$\theta_{t+1} = \theta_t - \frac{\eta}{\sqrt{\hat{v}_t} + \epsilon} \hat{m}_t.\tag{11}$$

The general goal is to create **artificial neural networks** (graphs) (**ANN**) that imitate to some extent the capabilities of the human brain:

- learning
- generalization
- adaptivity
- fault tolerance
- ...

We want this for example for

- pattern classification
- function approximation
- ...

Pioneering work was done by McCulloch and Pitts with the Perceptron [7, 8]. This was extended by Minsky and Papert [9].

McCulloch and Pitts proposed a binary threshold model as a computational model for an artificial neuron. Let x_1, \dots, x_n be the input values and $y = 0, 1$ be the output. The perceptron is defined by

$$y = \begin{cases} 0, & \sum_i x_i w_i \leq b \\ 1, & \sum_i x_i w_i > b \end{cases} . \quad (12)$$

where w_1, \dots, w_n are the synaptical weights that Rosenblatt [8] introduced (see Figure 4). This can be reformulated as

$$y = \Theta\left(\sum_{j=1}^n w_j x_j - b\right) . \quad (13)$$

This generates an output of 1 if the sum is above a certain threshold. Sometimes we include b in the sum and set $w_0 = -b$ and x_0 to a constant input $x_0 = 1$.




Figure 4: Perceptron.

In this setting

- positive weights correspond to **excitatory synapses**
- negative weights correspond to **inhibitory synapses**.

Clearly one can also use other activation function like

- piecewise linear
- sigmoid neuron

- gaussian.

Most often used is the sigmoid function (here the **logistic function**)

$$g(x) = \frac{1}{1 + e^{-\beta x}} . \quad (14)$$

The above constructed node is then the basic unit in a network (graph) of nodes.
Thus the ANN's are weighted directed graphs where

$$\text{neuron} \cong \text{node} \quad (15)$$

$$\text{connection between neuron} \cong \text{directed edge with weights} \quad (16)$$

Example: XOR




Figure 5: Perceptron XOR.

Connectionist models for gene regulation in the form of recurrent Hopfield [10] networks have been proposed by Mjolsness and others [11–13] to describe regulatory networks as directed graphs or matrices of interactions without restrictions on connectivity. These continuous time networks model interphase expression of a cell based on interaction weights that are free to take positive and negative real values.




Figure 6: Example of a neural network topology with input and output layers and one hidden layer.

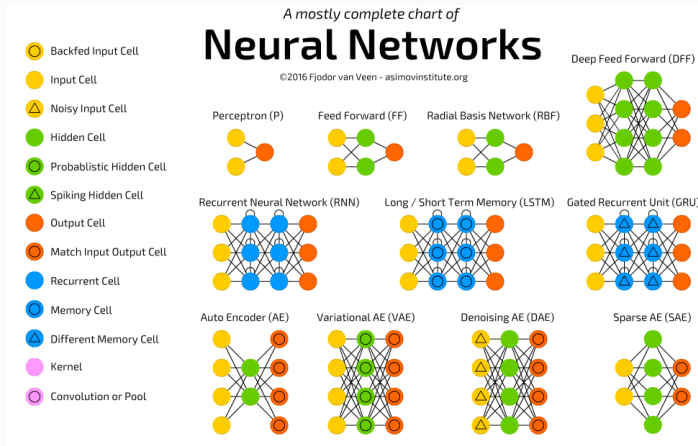


Figure 7: Overview of possible architectural designs of neural networks. Image taken from <https://medium.com/@carynmccarthy15/a-beginners-guide-to-recurrent-neural-networks-bfacb27bddb6>.

Prerequisite: Let $s \odot t$ ($(s \odot t)_j = s_j t_j$) denote the Hadamard product (Schur product), i.e., the elementwise product of the two vectors

Let $C(w, b)$ be a cost function where w is a weight and b is a bias with two assumptions about the form of the cost function:

- a
- b

quadratic cost function

$$C = \frac{1}{2n} \sum_x \|y(x) - a^L(x)\|^2, \quad (26)$$

$$C = \frac{1}{2} \|y - a^L\|^2 = \frac{1}{2} \sum_j (y_j - a_j^L)^2. \quad (27)$$

The goal of backpropagation is to compute the partial derivatives $\partial C / \partial w$ and $\partial C / \partial b$

$$a_j^l = \sigma \left(\sum_k w_{jk}^l a_k^{l-1} + b_j^l \right), \quad (17)$$

$$a^l = \sigma(w^l a^{l-1} + b^l) . \quad (18)$$

to compute a^l we compute $z^l \equiv w^l a^{l-1} + b^l$ weighted input to the neurons in layer l

$$a^l = \sigma(z^l)$$

$$z^l$$

$$z_j^l = \sum_k w_{jk}^l a_k^{l-1} + b_j^l$$

z_j^l is just the weighted input to the activation function for neuron j in layer l .

Let δ_j^l be the error in the j th neuron in the l th layer:

$$\delta_j^l \equiv \frac{\partial C}{\partial z_j^l} . \quad (29)$$

Backpropagation provides a procedure to compute the error δ_j^l . δ^l denotes the vector of errors associated with layer l .

For the error in the output layer L we have

$$\delta_j^L = \frac{\partial C}{\partial a_j^L} \sigma'(z_j^L). \quad (19)$$

Proof: We have

$$\delta_j^L = \frac{\partial C}{\partial z_j^L} \quad (20)$$

and with the chain rule we obtain

$$\delta_j^L = \sum_k \frac{\partial C}{\partial a_k^L} \frac{\partial a_k^L}{\partial z_j^L} \quad (21)$$

$$= \frac{\partial C}{\partial a_j^L} \frac{\partial a_j^L}{\partial z_j^L}. \quad (22)$$

because the sum is over all neurons k in the output layer and the output activation a_k^L of the k th neuron depends only on the weighted input z_j^L for the j th neuron when $k = j$. If $k \neq j$ $\partial a_k^L / \partial z_j^L$ is zero.

Since $a_j^L = \sigma(z_j^L)$ we write $\sigma'(z_j^L)$ we have Equation 22

$$\delta_j^L = \frac{\partial C}{\partial a_j^L} \sigma'(z_j^L) \quad (23)$$

proving our assumption.

$$\delta^L = \nabla_a C \odot \sigma'(z^L) . \quad (24)$$

$$\delta^L = (a^L - y) \odot \sigma'(z^L) . \quad (25)$$

An equation for the error δ^l in terms of the error in the next layer, δ^{l+1}

$$\delta^l = ((w^{l+1})^T \delta^{l+1}) \odot \sigma'(z^l) , \quad (26)$$

Proof: We rewrite $\delta_j^l = \partial C / \partial z_j^l$ in terms of $\delta_k^{l+1} = \partial C / \partial z_k^{l+1}$

$$\delta_j^l = \frac{\partial C}{\partial z_j^l} \quad (27)$$

$$= \sum_k \frac{\partial C}{\partial z_k^{l+1}} \frac{\partial z_k^{l+1}}{\partial z_j^l} \quad (28)$$

$$= \sum_k \frac{\partial z_k^{l+1}}{\partial z_j^l} \delta_k^{l+1}. \quad (29)$$

Note that

$$z_k^{l+1} = \sum_j w_{kj}^{l+1} a_j^l + b_k^{l+1} = \sum_j w_{kj}^{l+1} \sigma(z_j^l) + b_k^{l+1} \quad (30)$$

and taking the derivative

$$\frac{\partial z_k^{l+1}}{\partial z_j^l} = w_{kj}^{l+1} \sigma'(z_j^l). \quad (31)$$

we get

$$\delta_j^l = \sum_k w_{kj}^{l+1} \delta_k^{l+1} \sigma'(z_j^l). \quad (32)$$

$$\frac{\partial C}{\partial b_j^l} = \delta_j^l. \quad (33)$$

$$\frac{\partial C}{\partial b} = \delta, \quad (34)$$

$$\frac{\partial C}{\partial w_{jk}^l} = a_k^{l-1} \delta_j^l. \quad (35)$$

$$\frac{\partial C}{\partial w} = a_{\text{in}} \delta_{\text{out}}, \quad (36)$$


Algorithm 3 Backpropagation Algorithm (single input)

- 1: **repeat**
 - 2: Set the corresponding activation a^1 for the input layer
 - 3: **for** $l = 2, 3, \dots, L$ **do**
 - 4: $z^l = w^l a^{l-1} + b^l$
 - 5: $a^l = \sigma(z^l)$
 - 6: **end for**
 - 7: $\delta^L = \nabla_a C \odot \sigma'(z^L)$
 - 8: **for** $l = L - 1, L - 2, \dots, 2$ **do**
 - 9: $\delta^l = ((w^{l+1})^T \delta^{l+1}) \odot \sigma'(z^l)$
 - 10: **end for**
 - 11: $\frac{\partial C}{\partial w_{jk}^l} = a_k^{l-1} \delta_j^l$
 - 12: $\frac{\partial C}{\partial b_j^l} = \delta_j^l$
 - 13: **until** convergence is reached
-


Algorithm 4 Backpropagation Algorithm (batch input)

- 1: **repeat**
- 2: **for** each sample x **do**
- 3: Set the corresponding activation $a^{x,1}$ for the input layer
- 4: **for** $l = 2, 3, \dots, L$ **do**
- 5: $z^l = w^l a^{l-1} + b^l$
- 6: $a^l = \sigma(z^l)$
- 7: **end for**
- 8: $\delta^L = \nabla_a C \odot \sigma'(z^L)$
- 9: **for** $l = L - 1, L - 2, \dots, 2$ **do**
- 10: $\delta^l = ((w^{l+1})^T \delta^{l+1}) \odot \sigma'(z^l)$
- 11: **end for**
- 12: **end for**
- 13: **for** $l = L, L - 1, \dots, 2$ **do**
- 14: $w^l \rightarrow w^l - \frac{\eta}{m} \sum_x \delta^{x,l} (a^{x,l-1})^T$
- 15: $b^l \rightarrow b^l - \frac{\eta}{m} \sum_x \delta^{x,l}$
- 16: **end for**
- 17: $\frac{\partial C}{\partial w_{jk}^l} = a_k^{l-1} \delta_j^l$
- 18: $\frac{\partial C}{\partial b_j^l} = \delta_j^l$
- 19: **until** convergence is reached

Let us look at the problem of classifying walks into random walk and self-avoiding walk. Below are two images from a set of generated images.



random walk



Self avoiding walk


```
1 #from __future__ import print_function, division
3 import os
  import numpy as np
5 import tensorflow as tf
  #import matplotlib.pyplot as plt
7
  import matplotlib
9 matplotlib.use('TkAgg')
  import matplotlib.pyplot as plt
11 from skimage import data as dt
  from skimage import transform
13 from skimage.color import rgb2gray
  from scipy import misc
15 import random
17
  # Larger sample for RW and SAW
19 # Size dependence of the classification and recognition
  # What has the network learned
```

./progs/Walkclassifier.py

```
2 def plot_set(images, set):
4     for i in range(len(set)):
6         plt.subplot(1, len(set), i+1)
7         plt.axis('off')
8         plt.imshow(images[set[i]], cmap="gray")
9         plt.subplots_adjust(wspace=0.5)
10
11     plt.show()
12
13     pass
14
15 def load_data(data_directory, labels, images, L):
16     directories = [d for d in os.listdir(data_directory)
17                    if os.path.isdir(os.path.join(data_directory,
18                                                    d))]
19
20     for d in directories:
```

./progs/Walkclassifier.py

```
2 label_directory = os.path.join(data_directory, d)
3 file_names = [os.path.join(label_directory, f)
4               for f in os.listdir(label_directory)
5               if f.endswith(".png")]
6 for f in file_names:
7     img = dt.imread(f, as_gray=True)
8     cropped = img[L/4:3*L/4, L/4:3*L/4]
9     images.append(cropped)
10    labels.append(int(d))
11
12 print(len(images))
13 print(len(labels))
14 pass
15
16 def show_sample_prediction(sample_images, sample_labels):
17
18     fig = plt.figure(figsize=(10, 10))
19     for i in range(len(sample_images)):
20         truth = sample_labels[i]
```

./progs/Walkclassifier.py

```
prediction = predicted[i]
2 plt.subplot(5, 2,1+i)
  plt.axis('off')
4 color='green' if truth == prediction else 'red'
  plt.text(40, 10, "Truth:          {0}\nPrediction: {1}".format(
    truth, prediction),
6         fontsize=12, color=color)
  plt.imshow(sample_images[i], cmap="gray")
8
plt.show()

10
12
tf.set_random_seed(4711)
14
L = 500
16 crop_L = L / 2

18 ROOT_PATH = "/Users/heermann/tensorflow/Prog/Walk/data/"

20 train_data_directory = os.path.join(ROOT_PATH, "Training")
```

./progs/Walkclassifier.py

```
test_data_directory = os.path.join(ROOT_PATH, "Testing")
2
labels = []
4 images = []

6 load_data(train_data_directory, labels, images, L)
images = np.array(images)
8
set = [1, 2, 3, 4]
10 plot_set(images, set)

12
g = tf.Graph()
14
tf.image.per_image_standardization(images)
16

18 # Initialize placeholders
20 with g.as_default():
```

./progs/Walkclassifier.py

```
2   x = tf.placeholder(dtype = tf.float32 , shape = [None, crop_L ,
      crop_L])
3   y = tf.placeholder(dtype = tf.int32 , shape = [None])
4
5   # Flatten the input data
6   images_flat = tf.contrib.layers.flatten(x)
7
8   # Fully connected layer
9   fully_connected1 = tf.contrib.layers.fully_connected(
10  images_flat , 12, tf.nn.relu)
11  fully_connected2 = tf.contrib.layers.fully_connected(
12  fully_connected1 , 6, tf.nn.relu)
13  logits = tf.contrib.layers.fully_connected(fully_connected2 ,
14  12, tf.nn.relu)
15
16  # Define a loss function
17  loss = tf.reduce_mean(tf.nn.
18  sparse_softmax_cross_entropy_with_logits(labels = y,
      logits = logits))
19
20  # Define an optimizer
21  train_op = tf.train.AdagradOptimizer(learning_rate=0.001,
      name="Optimizer").minimize(loss)
22
23  # Convert logits to label indexes
```

```
2 # Define an accuracy metric
3 accuracy = tf.reduce_mean(tf.cast(correct_pred, tf.float32))
4
5
6 print("images_flat: ", images_flat)
7 print("logits: ", logits)
8 print("loss: ", loss)
9 print("predicted_labels: ", correct_pred)
10
11
12 # Add ops to save and restore all the variables.
13 saver = tf.train.Saver()
14 sess = tf.Session()
15
16 sess.run(tf.global_variables_initializer())
17
18 for i in range(201):
19     print('EPOCH', i)
20     _, accuracy_val = sess.run([train_op, accuracy],
21                                feed_dict={x: images, y: labels})
22     if i % 10 == 0:
```

./progs/Walkclassifier.py

```
2         print("Loss: ", loss)
3         print('DONE WITH EPOCH')
4
5         save_path = saver.save(sess, "./my-model.ckpt")
6         saver.save(sess, './my-model')
7         # Display layers
8         layers = {v.op.name: v for v in tf.trainable_variables()}
9         print(layers)
10        #
11
12    # Pick 10 random images
13    sample_indexes = random.sample(range(len(images)), 10)
14    sample_images = [images[i] for i in sample_indexes]
15    sample_labels = [labels[i] for i in sample_indexes]
16
17    # Run the "correct_pred" operation
18    predicted = sess.run([correct_pred], feed_dict={x: sample_images
19        })[0]
20
21    # Print the real and predicted labels
```

./progs/Walkclassifier.py


```
print(sample_labels)
2 print(predicted)

4 # Display the predictions and the ground truth visually.

6 show_sample_prediction(sample_images, sample_labels)

8
# #####

10 # Load the test data
test_labels = []
12 test_images = []


14 load_data(test_data_directory, test_labels, test_images, L)
test_images = np.array(test_images)
16 test_set = [4, 5, 6, 7]
plot_set(test_images, test_set)

18
tf.image.per_image_standardization(test_images)
```


./progs/Walkclassifier.py

```
2 # Run predictions against the full test set.  
   predicted = sess.run([correct_pred], feed_dict={x: test_images}  
   [0])  
4  
   # Calculate correct matches  
6 match_count = sum([int(y == y_) for y, y_ in zip(test_labels,  
   predicted)])  
   print('Match Count = ' + str(match_count) + ' out of = ' + str(len(  
   test_images)))  
8  
10 sess.close()
```


./progs/Walkclassifier.py



The result of the classification after minimal training is already very impressive.



- We would like to identify stretches of sequences that are actually functional (code for proteins or have regulatory functions) from non-coding or junk sequences.
- In prokaryotic DNA we have only two kinds of regions, ignore regulatory sequences which are coding (+) and non-coding (-) and the four letters A,C,G,T.



- This simulates a very common phenomenon [14]:

There is some underlying dynamic system running along according to simple and uncertain dynamics, but we cannot see it.

- All we can see are some noisy signals arising from the underlying system. From those noisy observations we want to do things like predict the most likely underlying system state, or the time history of states, or the likelihood of the next observation

What are Hidden Markov Models good for?

- useful for modeling protein/DNA sequence patterns
- probabilistic state-transition diagrams
- Markov processes - independence from history
- hidden states

- protein families
- DNA patterns
- secondary structure (helix, strand, coil (each has 20x20 table with transition frequencies between neighbors $a_i \rightarrow a_{i+1}$))
- protein fold recognition
- fold classification
- gene silencing
- ...

■ Example: CpG - Islands

The CpG sites or CG sites are regions of DNA where a cytosine nucleotide is followed by a guanine nucleotide in the linear sequence of bases along its 5' → 3' direction. (definition from Wikipedia)

- Regions labeled as CpG - islands → + model
- Regions labeled as non-CpG - islands → - model
- Maximum likelihood estimators for the transition probabilities for each model

$$a_{st} = \frac{c_{st}}{\sum_{t'} c_{st'}}$$

and analogously for the - model. c_{st} is the number of times letter t followed letter s in the labeled region.

- A Hidden Markov Model is a two random variable process, in which one of the random variables is hidden, and the other random variable is observable.
- It has a finite set of states, each of which is associated with a probability distribution.
- Transitions among the states are governed by transition probabilities.
- In a particular state an observation can be generated, according to the associated probability distribution.
- It is only the observation, not the state visible to an external observer, and therefore states are “hidden” from the observer.

- ■ For DNA, let + denote coding and - non-coding. Then a possible observed sequence could be

$O = AACCTTCCGCGCAATATAGGTAACCCCGG$


and

$Q = - - + + + + + + + + + + + + + + - - - - - - - - - -$

- Question: How can one find CpG - islands in a long chain of nucleotides?
- Merge both models into one model with small transition probabilities between the chains.

A **Hidden Markov Model (HMM)** $\lambda = \langle Y, S, A, B \rangle$ consists of:

- an output alphabet $Y = \{1, \dots, b\}$
 - a state space $S = \{1, \dots, c\}$ with a unique initial state s_0
 - a transition probability distribution $A(s'|s)$
 - an emission probability distribution $B(y|s, s')$
-
- HMMs are equivalent to (weighted) finite state automata with outputs on transitions.
 - Unlike MMs, constructing HMMs, estimating their parameters and computing probabilities are not so straightforward.



Given a HMM λ and a state sequence $S = (s_1, \dots, s_{t+1})$, the probability of an output sequence $O = (o_1, \dots, o_t)$ is

$$P(O|S, \lambda) = \prod_{i=1}^t P(o_i | s_i, s_{i+1}, \lambda) = \prod_{i=1}^t B(o_i | s_i, s_{i+1}) . \quad (37)$$

Given λ , the probability of a state sequence $S = (s_1, \dots, s_{t+1})$ is

$$P(S|\lambda) = \prod_{i=1}^t P(s_{i+1} | s_i) = \prod_{i=1}^t A(s_{i+1} | s_i) . \quad (38)$$

Of importance is the probability of an output sequence $O = (o_1, \dots, o_t)$ under a given λ . It is easy to show that the straightforward computation yields

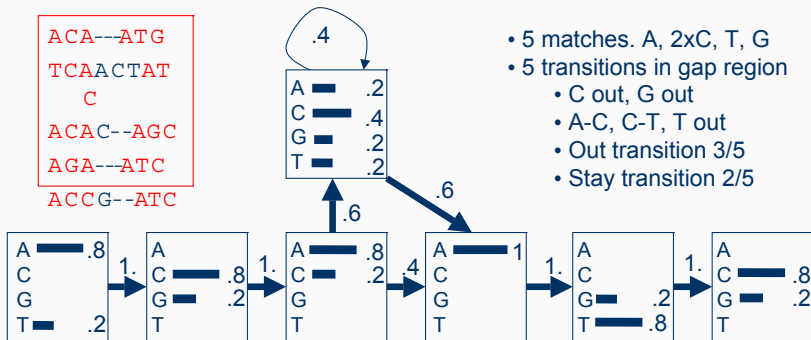
$$P(O|\lambda) = \sum_S \prod_{i=1}^t A(s_{i+1} | s_i) B(o_i | s_i, s_{i+1}) \quad (39)$$

with a computational complexity of $(2c + 1) * c^{t+1}$ multiplications.

Example: Multiple Sequence Alignments

- In theory, making an optimal alignment between two sequences is computationally straightforward (Smith-Waterman algorithm), but aligning a large number of sequences using the same method is almost impossible (e.g. $O(t^N)$).
- The problem increases exponentially with the number of sequences involved (the product of the sequence lengths).
- Statistical Methods:
 - Expectation Maximization Algorithm (deterministic).
 - Gibbs Sampler (stochastic).
 - Hidden Markov Models (stochastic).
- Advantages for HMM: theoretical explanation, no sequence ordering, no insertion and deletion penalties, using prior information.
- Disadvantage for HMM: large number of sequences for training.

ACA--ATG
TCAACTAT
C
ACAC--AGC
AGA--ATC
ACCG--ATC



- 5 matches. A, 2xC, T, G
- 5 transitions in gap region
 - C out, G out
 - A-C, C-T, T out
 - Out transition 3/5
 - Stay transition 2/5

$$ACA--ATG \quad 0.8 \times 1 \times 0.8 \times 1 \times 0.8 \times 0.4 \times 1 \times 0.8 \times 1 \times 0.2 = 3.3 \times 10^{-2}$$

- There are three basic problems:
 - 1 Given a model, how likely is a specific sequence of observed values (**evaluation problem**).
 - 2 Given a model and a sequence of observations, what is the most likely state sequence in the model that produces the observations (**decoding problem**).
 - 3 Given a model and a set of observations, how should the model's parameters be updated so that it has a high probability of generating the observations (**learning problem**).

- We define $\alpha_s(i)$ as the probability being in state s at position i :

$$\alpha_s(i) = P(o_1, \dots, o_i, s_i = s | \lambda) . \quad (40)$$

- Base case: $\alpha_s(1)$ if $s = s_0$ and $\alpha_s(0) = 0$ otherwise

- Induction:

$$\alpha_s(i+1) = \max_{s' \in S} A(s|s')B(o_i|s', s)\alpha_{s'}(i) . \quad (41)$$

- Finally, at the end:

$$P(o_1, \dots, o_k | \lambda) = \sum_{s \in S} \alpha_s(k) . \quad (42)$$

- Partial sums could as well be computed right to left (backward algorithm), or from the middle out

- In general, for any position i :

$$P(O|\lambda) = \sum_{s \in S} \alpha_s(i)\beta_s(i) . \quad (43)$$

- This algorithm could be used, e.g. to identify which λ is most likely to have produced an output sequence O .

What is the most probable path given observations (decoding problem)?

- Given o_1, \dots, o_t what is

$$\operatorname{argmax}_S P(s, o_1, \dots, o_t | \lambda) ? \quad (44)$$

- Slow and stupid answer:

$$\operatorname{argmax}_S \frac{P(o_1, \dots, o_t | s) P(s)}{P(o_1, \dots, o_t)} . \quad (45)$$

- We define $\delta_s(i)$ as the probability of the most likely path leading to state s at position i :

$$\delta_s(i) = \max_{s_1, \dots, s_{i-1}} P(s_1, \dots, s_{i-1}, o_1, \dots, o_{i-1}, s_i = s | M). \quad (46)$$

- Base case: $\delta_s(1)$ if $s = s_0$ and $\delta_s(0) = 0$ otherwise
- Again we proceed recursively:

$$\delta_s(i+1) = \max_{s' \in S} A(s|s')B(o_i|s', s)\delta_{s'}(i) \quad (47)$$

and since we want to know the identity of the best state sequence and not just its probability, we also need

$$\Psi(i+1) = \operatorname{argmax}_{s' \in S} A(s|s')B(o_i|s', s)\delta_{s'}(i). \quad (48)$$

- Finally, we can follow Ψ backwards from the most likely final state.
- The Viterbi algorithm efficiently searches through $|S|^T$ paths for the one with the highest probability in $O(T|S|^2)$ time.

- In practical applications, use log probabilities to avoid underflow errors.
- Can be easily modified to produce the n best paths.
- A beam search can be used to prune the search space further when $|S|$ is very large (n -gram models).

- Predicting the next state s_n depending on s_1, \dots, s_{n-1} results in

$$P(s_n | s_1, \dots, s_{n-1}) . \quad (49)$$

- Markov Assumption ($n - 1$)th order : last $n - 1$ states are in the same equivalence class.

- Given an HMM with a fixed architecture, how do we estimate the probability distributions A and B?
- If we have labeled training data, this is not any harder than estimating non-Hidden Markov Models (supervised training):

$$A(s'|s) = \frac{C(s \rightarrow s')}{\sum_{s''} C(s \rightarrow s'')} \quad (50)$$

$$B(o|s, s') = \frac{C(s \rightarrow s', o)}{C(s \rightarrow s')} \quad (51)$$

- Also known as the **Baum-Welch algorithm**.
- Instance of the **Expectation Maximization (EM) algorithm**:
 - 1 Choose a model at random.
 - 2 E: Find the distribution of state sequences given the model.
 - 3 M: Find the most likely model given those state sequences.
 - 4 Go back to 2.
- Our estimate of A is:

$$A(s'|s) = \frac{E[C(s \rightarrow s')]}{E[C(s \rightarrow ?)]} \quad (52)$$

- We estimate $E[C(s \rightarrow s')]$ via $\tau_t(s, s')$, the probability of moving from state s to state s' at position t given the output sequence O :

$$\tau_t(s, s') = P(s_t = s, s_{t+1} = s' | O, \lambda) \quad (53)$$

$$= \frac{P(s_t = s, s_{t+1} = s', O | \lambda)}{P(O | \lambda)} \quad (54)$$

$$= \frac{\alpha_s(t) A(s | s') B(o_{t+1} | s, s') \beta_{s'}(t+1)}{\sum_{s''} \alpha_{s''}} \quad (55)$$

- This lets us estimate A :

$$A(s' | s) = \frac{\sum_t \tau_t(s, s')}{\sum_t \sum_{s''} \tau_t(s, s'')} \quad (56)$$

- We can estimate B along the same lines, using $\sigma_t(o, s, s')$, the probability of emitting o while moving from state s to state s' at position t given the output sequence O .
- Alternate re-estimating A from τ , then τ from A , until estimates stop changing.
- If the initial guess is close to the right solution, this will converge to an optimal solution.

The fundamental idea of reinforcement learning is to interact with the environment and learn from this interaction. Let \mathcal{S} denote the states that the environment can be in and \mathcal{A} the actions that an **agent** interacting with the environment can take [15–17]. For each interaction the agent gets a **return** or **reward** $r \in \mathcal{R}$. The agent is trained maximizing the cumulative reward. The actions are chosen according to a policy π .





Image taken from: Kenji

Doya (DOI:

10.2976/1.2732246)



State cycle for the
reinforcement
learning

Figure 8

Let π be a **policy** mapping a state to an action

$$\pi : \mathcal{S} \rightarrow \mathcal{A} \tag{57}$$

$$s \mapsto a \tag{58}$$

The state cycle (c.f. Figure 8) is

$$t : s_t \rightarrow a_t \rightarrow s_{t+1} \tag{59}$$

$$(s_t, a_t, s_{t+1}) \rightarrow r_{t+1} \tag{60}$$

Let τ be a sequence or trajectory under π

$$\tau : (s_1, a_1, r_1, \dots, s_T, a_T, r_T) \sim \pi \tag{61}$$

where T is the **horizon**. We have $a_i \sim \pi_\theta(a_i|s_i)$ and $s_i \sim P(s_i|s_{i-1}, a_{i-1})$. Since the next state in the trajectory depends only on the immediate predecessor we have

$$P_\theta(s_1, a_1, r_1, \dots, s_T, a_T, r_T) = \mu(s_1) \prod_{i=2}^T \pi(a_i|s_i) P(s_i|s_{i-1}, a_{i-1}) \quad (62)$$

that the probability P to get the trajectory τ is split into individual transitions and represents the dynamics (Markov chain). μ is the starting state distribution. Such a sequence can be obtained using Monte Carlo methods.

The policy function is usually parameterized with a parameter θ :

$$\pi_\theta(s) = \pi(a | s, \theta) = p(a | s; \theta) . \quad (63)$$

Our objective is to maximize the return

$$\max_{\theta} \mathbb{E}_{\tau \sim \pi_\theta} [R(\tau)] = \max_{\theta} \int \pi_\theta(\tau) R(\tau) d\tau \quad (64)$$

with the return function R that usually is a function of (s_i, a_i, s_{i+1}) . Hence we want to find

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [\sum_t^T r(s_t, a_t)] \quad (65)$$

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{(s,a) \sim \pi_{\theta}(s,a)} [r(s, a)] \quad (66)$$

for finite horizon and infinite horizon respectively. Maximizing could be done by taking the derivative with respect to the return. However, the return may not be differentiable. This could be rectified by using a neural network (see later). We will take the approach of the **policy gradient**, i.e., taking the derivative with respect to the policy (parameter θ).

We have assumed that the policy is stochastic, i.e., mapping state s under the condition of parameter value θ to a with probability p . We distinguish between deterministic and stochastic policies:

- deterministic policy: $\pi(a | s, \theta) = 1$,
- Stochastic policy: $\pi(a | s, \theta) = p(a | s; \theta)$.

We have the choice to optimize values or actions:

- **Values policy:** Learn the interaction between states, actions and subsequent rewards.
- **Action policy:** Determine which is the best action to choose given the above.

Let $V^\pi(s)$ be the value of state s following policy π (**value-state function**):

$$V^\pi(s) = \mathbb{E}_{a \sim \pi}[G_t | S_t = s] \quad (67)$$

where

$$G_t = \sum_{k=0}^T \gamma^k r_{t+k+1} \quad (68)$$

is the **cumulative discounted return** with **discount parameter** γ ($0 < \gamma \leq 1$).

Further, let $Q^\pi(s, a)$ be the **action-value function**

$$Q^\pi(s, a) = \mathbb{E}_{a \sim \pi}[G_t | S_t = s, A_t = a] \quad (69)$$

and

$$A^\pi(s, a) = Q^\pi(s, a) - V^\pi(s) \quad (70)$$

the **advantage** telling us how much better or worse the action a is. Note that in complex formulae, for clarity, we are dropping the subscript θ .

We can define the **reward function** in terms of the state-value or action-state function as

$$J(\theta) = \sum_{s \in \mathcal{S}} d^\pi(s) V^\pi(s) = \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \pi_\theta(a|s) Q^\pi(s, a) \quad (71)$$

where $d^\pi(s) = \lim_{t \rightarrow \infty} P(s_t = s | s_0, \pi_\theta)$, i.e., the stationary distribution of the Markov chain of the policy π .

We want to take the gradient of the reward function to maximize the return with respect to the policy parameterized by θ . Before we do this, consider the following:

$$\nabla_{\theta} \mathbb{E}_{X \sim p(X|\theta)} [f(X)] = \nabla_{\theta} \left(\int_{\mathcal{X}} f(X) p(X | \theta) dX \right) \quad (72)$$

$$= \int_{\mathcal{X}} f(X) \nabla_{\theta} (p(X | \theta)) dX \quad (73)$$

$$= \int_{\mathcal{X}} f(X) p(X | \theta) \frac{\nabla_{\theta} (p(X | \theta))}{p(X | \theta)} dX \quad (74)$$

$$= \int_{\mathcal{X}} f(X) p(X | \theta) \nabla_{\theta} (\log p(X | \theta)) dX \quad (75)$$

$$= \mathbb{E}_{X \sim p(X|\theta)} [f(X) \nabla_{\theta} (\log p(X | \theta))] . \quad (76)$$

Hence, for the policy this implies

$$\nabla_{\theta} \pi_{\theta}(s, a) = \pi_{\theta}(s, a) \frac{\nabla_{\theta} \pi_{\theta}(s, a)}{\pi_{\theta}(s, a)} = \pi_{\theta}(s, a) \nabla_{\theta} \log \pi_{\theta}(s, a) . \quad (77)$$

We define the score function to be

$$\nabla_{\theta} \log \pi_{\theta}(s, a) . \quad (78)$$

It describes how sensitive the stochastic policy π to is to changes in θ , i.e. how likely the trajectory is under the current policy.

Example: Policy function

Let us look at the following example of the policy function for a linear model for the unnormalized log-probability: $\phi(s, a)^T \theta$ i.e. weighting of the actions using a linear combination of features $\phi(s, a)$

The score function for a softmax policy is:

$$\pi_{\theta}(s, a) = \frac{e^{\phi(s, a)^T \theta}}{\sum_{a' \in \mathcal{A}} e^{\phi(s, a')^T \theta}} \quad (79)$$

$$\nabla_{\theta} \log \pi_{\theta}(s, a) = \phi(s, a) - \mathbb{E}_{\pi_{\theta}}[\phi(s, \cdot)] \quad (80)$$

$$h_{\theta}(x) = \frac{1}{1 + e^{-\theta^T x}} \cdot \quad (81)$$

Stochastic Gradient Policy Theorem

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\pi_{\theta}} [\nabla_{\theta} \log \pi_{\theta}(s, a) Q^{\pi_{\theta}}(s, a)] . \quad (82)$$

Hence, the computation of the policy gradient reduces to a simple expectation. Thus, we are looking for sampling algorithms where trajectories are generated, the action value or value state function being evaluated along the trajectory (sometimes called **play out** or **episode**) and the gradient of the log of the policy computed.

To show the above statement, we take the following steps (see Lilian Weng <https://lilianweng.github.io/lil-log/2018/04/08/policy-gradient-algorithms.html>):

$$\nabla_{\theta} V^{\pi}(s) = \nabla_{\theta} \left[\sum_{a \in \mathcal{A}} \pi_{\theta}(a|s) Q^{\pi}(s, a) \right] \quad (83)$$

$$= \sum_{a \in \mathcal{A}} \left[\nabla_{\theta} \pi_{\theta}(a|s) Q^{\pi}(s, a) + \pi_{\theta}(a|s) \nabla_{\theta} Q^{\pi}(s, a) \right] \quad (84)$$

$$= \sum_{a \in \mathcal{A}} \left[\nabla_{\theta} \pi_{\theta}(a|s) Q^{\pi}(s, a) + \pi_{\theta}(a|s) \nabla_{\theta} \sum_{s', r} P(s', r|s, a) (r + V^{\pi}(s')) \right] \quad (85)$$

$$= \sum_{a \in \mathcal{A}} \left[\nabla_{\theta} \pi_{\theta}(a|s) Q^{\pi}(s, a) + \pi_{\theta}(a|s) \sum_{s', r} P(s', r|s, a) \nabla_{\theta} V^{\pi}(s') \right] \quad (86)$$

$$= \sum_{a \in \mathcal{A}} \left[\nabla_{\theta} \pi_{\theta}(a|s) Q^{\pi}(s, a) + \pi_{\theta}(a|s) \sum_{s'} P(s'|s, a) \nabla_{\theta} V^{\pi}(s') \right]. \quad (87)$$

$$\nabla_{\theta} V^{\pi}(s) = \phi(s) + \sum_a \pi_{\theta}(a|s) \sum_{s'} P(s'|s, a) \nabla_{\theta} V^{\pi}(s') \quad (88)$$

$$= \phi(s) + \sum_{s'} \sum_a \pi_{\theta}(a|s) P(s'|s, a) \nabla_{\theta} V^{\pi}(s') \quad (89)$$

$$= \phi(s) + \sum_{s'} \rho^{\pi}(s \rightarrow s', 1) \nabla_{\theta} V^{\pi}(s') \quad (90)$$

$$= \phi(s) + \sum_{s'} \rho^{\pi}(s \rightarrow s', 1) \nabla_{\theta} V^{\pi}(s') \quad (91)$$

$$= \phi(s) + \sum_{s'} \rho^{\pi}(s \rightarrow s', 1) [\phi(s') + \sum_{s''} \rho^{\pi}(s' \rightarrow s'', 1) \nabla_{\theta} V^{\pi}(s'')] \quad (92)$$

$$= \phi(s) + \sum_{s'} \rho^{\pi}(s \rightarrow s', 1) \phi(s') + \sum_{s''} \rho^{\pi}(s \rightarrow s'', 2) \nabla_{\theta} V^{\pi}(s'') \quad (93)$$

$$= \phi(s) + \sum_{s'} \rho^{\pi}(s \rightarrow s', 1) \phi(s') + \sum_{s''} \rho^{\pi}(s \rightarrow s'', 2) \phi(s'') + \dots \quad (94)$$

$$= \dots; \text{ Repeatedly unrolling the part of } \nabla_{\theta} V^{\pi}(\cdot) \quad (95)$$

$$= \sum_{x \in \mathcal{S}} \sum_{k=0}^{\infty} \rho^{\pi}(s \rightarrow x, k) \phi(x) . \quad (96)$$

$$\nabla_{\theta} J(\theta) = \nabla_{\theta} V^{\pi}(s_0) \quad (97)$$

$$= \sum_s \sum_{k=0}^{\infty} \rho^{\pi}(s_0 \rightarrow s, k) \phi(s) \quad (98)$$

$$= \sum_s \eta(s) \phi(s) \quad (99)$$

$$= \left(\sum_s \eta(s) \right) \sum_s \frac{\eta(s)}{\sum_s \eta(s)} \phi(s) \quad (100)$$

$$\propto \sum_s \frac{\eta(s)}{\sum_s \eta(s)} \phi(s) \quad (101)$$

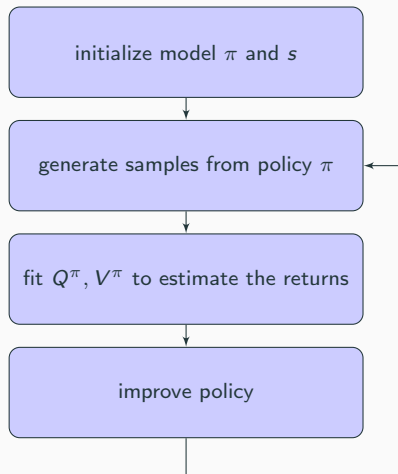
$$= \sum_s d^{\pi}(s) \sum_a \nabla_{\theta} \pi_{\theta}(a|s) Q^{\pi}(s, a) . \quad (102)$$

$\sum_s \eta(s)$ is the average length of the episode in the continuous case. And further

$$\nabla_{\theta} J(\theta) \propto \sum_{s \in \mathcal{S}} d^{\pi}(s) \sum_{a \in \mathcal{A}} Q^{\pi}(s, a) \nabla_{\theta} \pi_{\theta}(a|s) \quad (103)$$

$$= \sum_{s \in \mathcal{S}} d^{\pi}(s) \sum_{a \in \mathcal{A}} \pi_{\theta}(a|s) Q^{\pi}(s, a) \frac{\nabla_{\theta} \pi_{\theta}(a|s)}{\pi_{\theta}(a|s)} \quad (104)$$

$$= \mathbb{E}_{\pi} [Q^{\pi}(s, a) \nabla_{\theta} \ln \pi_{\theta}(a|s)] . \quad (105)$$



Actually, the question is how to compute the score function

$$\nabla_{\theta} \log \pi_{\theta}(s, a)$$

(106)

specifically under light that gradients can be very noisy. They suffer from high variance and low convergence. We have

$$\nabla_{\theta} \log P_{\theta}(\tau) = \nabla \log \left(P(s_1) \prod_{t=1}^T \pi_{\theta}(a_t | s_t) P(s_{t+1} | s_t, a_t) \right) \quad (107)$$

$$= \nabla_{\theta} \left[\log \mu(s_1) + \sum_{t=1}^T (\log \pi_{\theta}(a_t | s_t) + \log P(s_{t+1} | s_t, a_t)) \right] \quad (108)$$

$$= \nabla_{\theta} \sum_{t=1}^T \log \pi_{\theta}(a_t | s_t) . \quad (109)$$

Using gradient ascend

$$\theta \leftarrow \theta + \alpha \nabla f(x) \quad (110)$$

we can write the generic algorithm is as follows:

Algorithm 5 Gradient Policy

1: **repeat**

2: $\nabla_{\theta} J(\theta) = \frac{1}{N} \sum_{i=1}^N (\sum_{t=1}^T \nabla_{\theta} \log \pi_{\theta}(s_{i,t}, a_{i,t})) (\sum_{t=1}^T R(s_{i,t}, a_{i,t}))$

3: $\theta \leftarrow \theta + \alpha \nabla_{\theta} J(\theta)$

4: **until** finished

α is the **learning rate** determining the rate of convergence.

The **Monte Carlo policy gradient** or **REINFORCE** estimates (learns) the value state function V^π from episodes under policy π . Hence, we generate episodes

$$s_1, a_1, \dots, s_T, a_T \sim \pi. \quad (111)$$

The Monte-Carlo policy evaluation uses empirical mean return instead of expected return. Note that for the cumulated discounted return we have

$$Q^\pi(s_t, a_t) = \mathbb{E}_\pi[G_t | a_t, s_t] \quad (112)$$

and hence we can write

$$\nabla_\theta J(\theta) = \mathbb{E}_\pi[Q^\pi(s, a) \nabla_\theta \ln \pi_\theta(a|s)] \quad (113)$$

$$= \mathbb{E}_\pi[G_t \nabla_\theta \ln \pi_\theta(a_t | s_t)] \quad (114)$$

and sample the return. Our goal is find the policy, i.e. the value of θ maximizing the return

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{\pi} \left[\sum_{t=1}^T \gamma^t r_t \right]. \quad (115)$$

A generic version is listed in Algorithm 6.

Algorithm 6 Generic Monte-Carlo Policy Evaluation: REINFORCE

- 1: Initialize the policy parameter θ
 - 2: **repeat**
 - 3: Generate episode using $\pi_\theta \sim (s_1, a_1, \dots, a_T, s_T)$
 - 4: **for** t in range $(1, T)$ **do**
 - 5: Evaluate G_t
 - 6: $\theta \leftarrow \theta + \alpha \gamma^t G_t \nabla_\theta \ln \pi_\theta(a_t | s_t)$
 - 7: **end for**
 - 8: **until** false
-




Figure 9: Balancing of a stick

- The task at hand is to balance a stick in a one-dimensional setting. The stick is mounted frictionless on a rail and can be moved to left and to the right. The stick can rotate and is subject to gravitation (for info on the gym environment implementing the balance stick see <https://github.com/openai/gym/wiki/CartPole-v0>).
- Neglecting friction, the equations of motion are [18]:

$$\ddot{x} = \frac{F + m_p l (\alpha^2 \sin \alpha - \ddot{\alpha} \cos \alpha)}{m_c + m_p} \quad (116)$$

$$\ddot{\alpha} = \frac{g \sin \alpha + \cos \alpha \left(\frac{-F - m_p l \alpha^2 \sin \alpha}{m_c + m_p} \right)}{l \left(\frac{4}{3} - \frac{m_p \cos \alpha^2}{m_c + m_p} \right)} \quad (117)$$

- In this example we assume that there are two action $a = 0, 1$ or $a = -1, +1$ corresponding to **left** and **right**. The state of the system is given by a state vector with the components $s = (\text{position } (x), \text{velocity } (v), \text{stick angle } (\alpha), \text{velocity at tip } (v_T))$.

Table 1: openai CartPole v0 states (<https://openai.com/resources/>)

Num	Observation	Min	Max
0	Cart Position	-2.4	2.4
1	Cart Velocity	-Inf	Inf
2	Pole Angle	$\sim -41.8^\circ$	$\sim 41.8^\circ$
3	Pole Velocity At Tip	-Inf	Inf

- Since our action is binary, we can choose the logistic function as part of the policy π

$$L(x) = \frac{1}{1 + e^{-x}}. \quad (118)$$

- We can define the policy π as

$$\pi_\theta(s, a = 0) = 1 - L(s^T \theta) \quad (119)$$

$$\pi_\theta(s, a = 1) = L(s^T \theta). \quad (120)$$

- Our task is to estimate the state action function

$$Q^\pi(s, a) \tag{121}$$

from the discounted return function

$$G_t = r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots + \gamma^{T-t} r_T \tag{122}$$

where the reward is 1 for every step taken, including the termination step.

$$J(\theta) \approx \sum_{t=1}^T \pi(a_t | s_t, \theta) A_t . \tag{123}$$

- For one episode we have

$$\nabla_\theta J(\theta) \approx \sum_{t=1}^T G_t \nabla_\theta \log \pi_\theta(s, a) \tag{124}$$

$$\frac{d}{dx} \text{sigmoid}(x) = \text{sigmoid}(x)(1 - \text{sigmoid}(x)) . \tag{125}$$

- The problem is considered solved when the average reward is greater than or equal to 195.0 over 100 consecutive trials.

```
1 def episode(theta, max_episode_length=1000):
3     observation = env.reset()
5
6     actions      = []
7     states      = []
8     rewards     = []
9     done        = False
11
12     i = 0
13
14     while not done:
15         i+=1
16
17         action = get_action(theta, observation)
18         states.append(observation)
19         actions.append(action)
20         observation, reward, done, info = env.step(action)
21         rewards.append(reward)
22
23         if i > max_episode_length:
24             break
25
26     return np.array(rewards), np.array(states), np.array(actions)
```



```
2 def discounted_sum_of_rewards(rewards, gamma):
4     cum = np.zeros_like(rewards)
      c   = 0.0
6     for i, r in enumerate(rewards[::-1]):
          c   = r + gamma * c
8         cum[i] = c
      return cum[::-1]
```

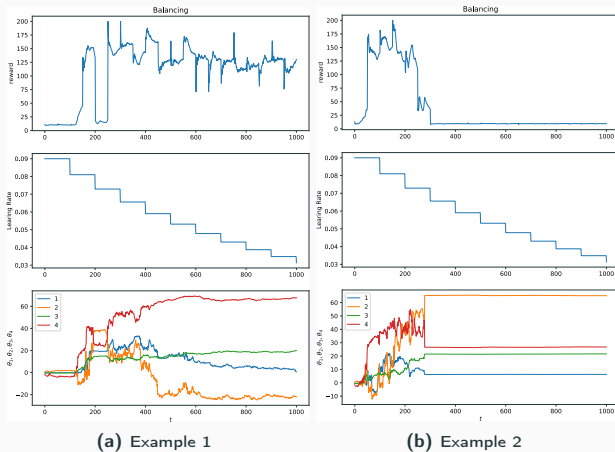


Figure 10: Two examples of balancing of a stick using Monte Carlo policy gradient reinforcement learning. Max play out length was 1000 and 1000 episodes were calculated.

Algorithm 7 Generic Monte-Carlo Policy Evaluation

- 1: Given π the policy to be evaluated
 - 2: Initialize V randomly
 - 3: Returns(s) \leftarrow empty list for all $s \in S$
 - 4: **repeat**
 - 5: Generate episode using π
 - 6: **for** s in trial **do**
 - 7: $R \leftarrow$ return following the first occurrence of s
 - 8: Append R to Returns(s)
 - 9: $V(s) \leftarrow$ average(Returns(s))
 - 10: **end for**
 - 11: **until** false
-

Bellman Equation:

$$V^\pi(s) = \sum_a \pi(a | s) \left(R_s^a + \gamma \sum_{s' \in S} P_{s,s'}^a V^\pi(s') \right). \quad (126)$$

- This works well because the output is a probability over available actions.
- If we feed it with a neural network, we will get higher values and thus we will be more likely to choose the actions that we learned produce a better reward.
- In the long-run, this will trend towards a deterministic policy, $\pi(a | s, \theta) = 1$, but it will continue to explore as long as one of the probabilities does not dominate the others (which will likely take some time).

For the algorithm we are going to assume

- a differentiable policy parameterization $\pi(a | s, \theta)$
- and define the step-size $\alpha > 0$.

Algorithm 8 Generic Monte-Carlo Policy Evaluation Neural Network

```
1: Initialize policy parameters  $\theta$ 
2: repeat
3:   Generate episode using  $\pi$ 
4:   for  $N$  batches do
5:     Generate an episode  $s_0, a_0, r_1, \dots, s_{T-1}, a_{T-1}, r_T$ , following  $\pi(a | s, \theta)$ 
6:     for  $t = 0, \dots, T - 1$  do
7:        $G_t \leftarrow$  from step  $t$ 
8:     end for
9:     Calculate the loss  $L(\theta) = -\frac{1}{N} \sum_t^T \ln(\gamma^t G_t \pi(a_t | s_t, \theta))$ 
10:    Update policy parameters through backpropagation:  $\theta := \theta + \alpha \nabla_{\theta} L(\theta)$ 
11:  end for
12: until  $n$  episodes
```

We are going to apply the neural network approach to the balancing of a stick problem defined above. We will be using a fully connected neural network as shown in Figure 11. The layer size is halved from one layer to the next. The last layer essentially represents a binary decision to move left or right.




Figure 11: Principle design of the neural network to be used in the learning of the policy gradient in our example. All layers are fully connected. Only the last layers in this graph show the actual connectivity.



```
1 import gym
2 import tensorflow as tf
3 from tensorflow.contrib.layers import fully_connected
4 import warnings
5
6 class policy_estimator(object):
7
8     def __init__(self, sess, env):
9         # Pass TensorFlow session object
10        self.sess = sess
11
12        # Get number of inputs and outputs from environment
13        self.n_inputs = env.observation_space.shape[0]
14        self.n_outputs = env.action_space.n
15        self.learning_rate = 0.01
16
17        # Define number of hidden nodes
18        self.n_hidden_nodes = 256
19
20        # Set graph scope name
21        self.scope = "policy_estimator"
```

```
2      # Create network
3      with tf.variable_scope(self.scope):
4          initializer = tf.contrib.layers.xavier_initializer()
5
6          # Define placeholder tensors for state, actions,
7          # and rewards
8          self.state = tf.placeholder(tf.float32,
9                                     [None, self.n_inputs], name='state')
10         self.rewards = tf.placeholder(tf.float32,
11                                       [None], name='rewards')
12         self.actions = tf.placeholder(tf.int32,
13                                      [None], name='actions')
14
15         layer_1 = fully_connected(self.state,
16                                   self.n_hidden_nodes,
17                                   activation_fn=tf.nn.swish,
18                                   weights_initializer=initializer)
19         layer_2 = fully_connected(layer_1,
20                                   int(self.n_hidden_nodes/2),
```



```
2         activation_fn=tf.nn.swish,  
3         weights_initializer=initializer)  
4     layer_3 = fully_connected(layer_2,  
5         int(self.n_hidden_nodes/4),  
6         activation_fn=tf.nn.swish,  
7         weights_initializer=initializer)  
8     layer_4 = fully_connected(layer_3,  
9         int(self.n_hidden_nodes/8),  
10        activation_fn=tf.nn.swish,  
11        weights_initializer=initializer)  
12    layer_5 = fully_connected(layer_4,  
13        int(self.n_hidden_nodes/16),  
14        activation_fn=tf.nn.swish,  
15        weights_initializer=initializer)  
16    layer_6 = fully_connected(layer_5,  
17        int(self.n_hidden_nodes/32),  
18        activation_fn=tf.nn.swish,  
19        weights_initializer=initializer)  
20    output_layer = fully_connected(layer_6,  
21        self.n_outputs,
```

```
2         activation_fn=None,
3         weights_initializer=initializer)
4
5     # Get probability of each action
6     self.action_probs = tf.squeeze(
7         tf.nn.softmax(output_layer -
8             tf.reduce_max(output_layer)))
9
10    # Get indices of actions
11    indices = tf.range(0, tf.shape(output_layer)[0]) \
12        * tf.shape(output_layer)[1] + self.actions
13
14    selected_action_prob = tf.gather(
15        tf.reshape(self.action_probs, [-1]), indices)
16
17    # Define loss function
18    self.loss = -tf.reduce_mean(
19        tf.log(selected_action_prob) * self.rewards)
20
21    # Get gradients and variables
```

```
2 self.tvars = tf.trainable_variables(self.scope)
3 self.gradient_holder = []
4 for j, var in enumerate(self.tvars):
5     self.gradient_holder.append(
6         tf.placeholder(tf.float32,
7             name='grads' + str(j)))
8
9     self.gradients = tf.gradients(self.loss,
10         self.tvars)
11
12     # Minimize training error
13     self.optimizer = tf.train.AdamOptimizer(
14         self.learning_rate)
15     self.train_op = self.optimizer.apply_gradients(
16         zip(self.gradient_holder, self.tvars))
17
18 def predict(self, state):
19     probs = self.sess.run([self.action_probs],
20         feed_dict={self.state: state})[0]
```

```
    return probs
2
3
4 def update(self, gradient_buffer):
    feed = dict(zip(self.gradient_holder, gradient_buffer))
6     self.sess.run([self.train_op], feed_dict=feed)
7
8
9
10 def get_vars(self):
    net_vars = self.sess.run(
11         tf.trainable_variables(self.scope))
12     return net_vars
13
14
15
16 def get_grads(self, states, actions, rewards):
    grads = self.sess.run([self.gradients],
17         feed_dict={
18             self.state: states,
19             self.actions: actions,
20             self.rewards: rewards
```




2


```
} [0]  
return grads
```

Stability:


Reinforcement learning is known to be unstable or even to diverge when a nonlinear function approximator such as a neural network is used to represent the action-value (also known as Q) function. This instability has several causes: the correlations present in the sequence of observations, the fact that small updates to Q may significantly change the policy and therefore change the data distribution, and the correlations between the action-values and the target values [19].




(a) first



(b) second



(c) third



(d) forth

Figure 12: Demonstration of the variance involved in the REINFORCE algorithm. Here results from the application of neural network learning of the policy is shown. Shown are results for a neural network where the first layer consists of 128 fully connected nodes as shown schematically in Figure 11. The episode length was a maximum of 100.




Figure 13: Shown are results for a neural network where the first layer consists of 256 fully connected nodes as shown schematically in Figure 11. The sample used an episode length of a maximum of 200.

To reduce the variance, the standard is to introduce a function $b(s_t)$ inside the expectation on which we are computing the gradient. b is supposed to be an expected return. Let $R(\tau) = \sum_{t=0}^{T-1} r_t$ where we have set the discount parameter equal to one. We can write

$$\nabla_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [R(\tau)] = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\left(\sum_{t=0}^{T-1} r_t \right) \cdot \nabla_{\theta} \left(\sum_{t=0}^{T-1} \log \pi_{\theta}(a_t | s_t) \right) \right] \quad (127)$$

$$= \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t'=0}^{T-1} r_{t'} \sum_{t=0}^{t'} \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \right] \quad (128)$$

$$= \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^{T-1} \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \left(\sum_{t'=t}^{T-1} r_{t'} \right) \right]. \quad (129)$$

With this we can introduce the baseline function b

$$\nabla_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [R(\tau)] = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^{T-1} \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \left(\sum_{t'=t}^{T-1} r_{t'} - b(s_t) \right) \right]. \quad (130)$$

If γ is not one than

$$\nabla_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [R(\tau)] = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^{T-1} \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \left(\sum_{t'=t}^{T-1} r_{t'} - b(s_t) \right) \right] \quad (131)$$

$$\approx \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^{T-1} \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \left(\sum_{t'=t}^{T-1} \gamma^{t'-t} r_{t'} - b(s_t) \right) \right] \quad (132)$$

with the baseline

$$b(s_t) \approx \mathbb{E}[r_t + \gamma r_{t+1} + \dots + \gamma^{T-1-t} r_{T-1}]. \quad (133)$$

The REINFORCE Algorithm with baseline is shown in Algorithm 9. Let

$$\theta_p := \theta_p + \alpha_p \gamma^t \delta \nabla_{\theta_p} \ln(\pi(a_t | s_t, \theta_p)) \quad (134)$$

where δ is the difference between the actual value and the predicted value at that given state:

$$\delta = G_t - v(S_t, \theta_v) . \tag{135}$$

Note that the subscripts p and v to differentiate between the policy estimation function and the value estimation function. Thus, we assume a differentiable policy parameterization $\pi(a | s, \theta_p)$ and a differentiable value function parameterization $v(s, \theta_v)$

Algorithm 9 REINFORCE with baseline: Monte-Carlo policy gradient

- 1: Define step-size $\alpha_p > 0$, $\alpha_v > 0$
 - 2: Initialize policy parameters θ_p , θ_v
 - 3: **repeat**
 - 4: **for** N batches **do**
 - 5: Generate an episode $s_0, a_0, r_1, \dots, s_{T-1}, a_{T-1}, r_T$, following $\pi(a | s, \theta_p)$
 - 6: **for** $t = 0, \dots, T - 1$ **do**
 - 7: $G_t \leftarrow$ from step t
 - 8: **end for**
 - 9: $\delta \leftarrow G_t - v(s, \theta_v)$
 - 10: Calculate the loss $L(\theta_v) = \frac{1}{N} \sum_t^T (\gamma^t G_t - v(s_t, \theta_v))^2$
 - 11: Calculate the loss $L(\theta_p) = -\frac{1}{N} \sum_t^T \ln(\gamma^t \delta \pi(a_t | S_t, \theta_p))$
 - 12: Update policy parameters through backpropagation: $\theta_p := \theta_p + \alpha_p \nabla_{\theta}^p L(\theta_p)$
 - 13: Update policy parameters through backpropagation: $\theta_v := \theta_v + \alpha_v \nabla_{\theta}^v L(\theta_v)$
 - 14: **end for**
 - 15: **until** n episodes
-

Recall the definition of the action state function

$$Q^\pi(s, a) = \mathbb{E}_{\mathcal{T} \sim \pi_\theta} \left[\sum_{t=0}^{T-1} r_t \mid s_0 = s, a_0 = a \right] \quad (136)$$

and the value-state function

$$V^\pi(s) = \mathbb{E}_{\mathcal{T} \sim \pi_\theta} \left[\sum_{t=0}^{T-1} r_t \mid s_0 = s \right] \quad (137)$$

and the advantage function

$$A^\pi(s, a) = Q^\pi(s, a) - V^\pi(s). \quad (138)$$

We have [20]

$$\nabla_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} [R(\tau)] = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^{T-1} \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \left(\sum_{t'=t}^{T-1} r_{t'} - b(s_t) \right) \right] \quad (139)$$

$$= \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^{T-1} \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \cdot \left(Q^{\pi}(s_t, a_t) - V^{\pi}(s_t) \right) \right] \quad (140)$$

$$= \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^{T-1} \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \cdot A^{\pi}(s_t, a_t) \right] \quad (141)$$

$$\approx \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^{T-1} \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \cdot A^{\pi, \gamma}(s_t, a_t) \right] . \quad (142)$$

- Deep Q network and the epsilon-greedy policy.
- Q learning is a value based method of supplying information to inform which action an agent should take.
- In tabular Q-learning, for example, you are selecting the action that gives the highest expected reward ($\max_a Q(s', a')$), possibly also in an ϵ -greedy fashion) which means if the values change slightly, the actions and trajectories may change radically.

The Q learning rule

$$Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a') - Q(s, a)] \quad (143)$$

$$0 \leq \gamma \leq 1 \quad (144)$$

$$\alpha \quad (145)$$

with α being the learning rate.

Both α and the $Q(s, a)$ subtraction are not required to be explicitly defined in deep Q learning, as the neural network will take care of that during its optimized learning

process, i.e., deep Q-learning applies the Q-learning updating rule during the training process. A neural network is created which takes the state s as its input, and then the network is trained to output appropriate $Q(s, a)$ values for each action in state s .

Actor-critic methods consist of two models, which may optionally share parameters:

- Critic updates the value function parameters w and depending on the algorithm it could be action-value $Q_w(a|s)$ or state-value $V_w(s)$.
- Actor updates the policy parameters θ for $\pi_\theta(a|s)$ in the direction suggested by the critic.

Let α_θ and α_w be two learning rates. predefined for policy and value function parameter updates respectively. The actor-critic Monte-Carlo policy gradient algorithm is shown in Algorithm 10.

Algorithm 10 Actor-Critic: Monte-Carlo policy gradient

- 1: Initialize s, θ, w at random; sample $a \sim \pi_\theta(a|s)$
 - 2: **for** $t (1, \dots, T)$ **do**
 - 3: Sample reward $r_t \sim R(s, a)$ and next state $s' \sim P(s'|s, a)$
 - 4: Then sample the next action $a' \sim \pi_\theta(a'|s')$
 - 5: Update the policy parameters: $\theta \leftarrow \theta + \alpha_\theta Q_w(s, a) \nabla_\theta \ln \pi_\theta(a|s)$
 - 6: Compute the correction (TD error) for action-value at time t :
 - $\delta_t = r_t + \gamma Q_w(s', a') - Q_w(s, a)$
 - $w \leftarrow w + \alpha_w \delta_t \nabla_w Q_w(s, a)$
 - 7: Update $a \leftarrow a'$ and $s \leftarrow s'$
 - 8: **end for**
-

Let r be a uniform random number. ϵ -greedy strategy





Figure 14: ϵ -greedy strategy. r is a uniform random number.

Algorithm 11 ϵ -greedy strategy

```
1: for  $i$  (1, ..., samples) do  
2:    $r \sim \text{uniform}(0,1)$   
3:   if  $r < \epsilon$  then  
4:     choose random action  
5:   else  
6:     choose best action  
7:   end if  
8: end for
```



Double pendulum

```
1 '''
3 Source:
4 https://adventuresinmachinelearning.com/reinforcement-learning-tensorflow/
5 '''
7
8
9 import gym
10
11 import tensorflow as tf
12
13 import numpy as np
14 import time
15 import seaborn as sns
16 import matplotlib as mpl
17 import matplotlib.pyplot as plt
18 from scipy.stats import norm
19 import random as random
20 import math
```

```
BATCH_SIZE = 100
2 MAX_EPSILON = 1.0
  MIN_EPSILON = 0.0
4 LAMBDA = 0.001
  GAMMA = 0.99
6
8 class Model:
10
12     def __init__(self, num_states, num_actions, batch_size):
14
16         self._num_states = num_states
17         self._num_actions = num_actions
18         self._batch_size = batch_size
19
20         # define the placeholders
21         self._states = None
22         self._actions = None
```



```
2      # the output operations
3      self._logits = None
4      self._optimizer = None
5      self._var_init = None
6
7      # now setup the model
8      self._define_model()
9
10     def _define_model(self):
11
12         self._states = tf.placeholder(shape=[None, self._num_states
13 ], dtype=tf.float32)
14         self._q_s_a = tf.placeholder(shape=[None, self.
15 _num_actions], dtype=tf.float32)
16
17         # create a couple of fully connected hidden layers
18         fc1          = tf.layers.dense(self._states, 50, activation
19 =tf.nn.relu)
20         fc2          = tf.layers.dense(fc1, 50, activation=tf.nn.
21 relu)
22         self._logits = tf.layers.dense(fc2, self._num_actions)
23
24         loss          = tf.losses.mean_squared_error(self._q_s_a,
25 self._logits)
```

```
2     self._optimizer = tf.train.AdamOptimizer().minimize(loss)
3     self._var_init = tf.global_variables_initializer()
4
5
6     def predict_one(self, state, sess):
7         return sess.run(self._logits, feed_dict={self._states:
8                                                     state.reshape
9                                                     (1, self._num_states)})
10
11
12
13
14     def predict_batch(self, states, sess):
15         return sess.run(self._logits, feed_dict={self._states:
16                                                     states})
17
18
19
20     def train_batch(self, sess, x_batch, y_batch):
21         sess.run(self._optimizer, feed_dict={self._states: x_batch,
22                                               self._q_s_a: y_batch})
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
```

```
2     self._max_memory = max_memory
3     self._samples = []
4
5     def add_sample(self, sample):
6
7         self._samples.append(sample)
8         if len(self._samples) > self._max_memory:
9             self._samples.pop(0)
10
11     def sample(self, no_samples):
12
13         if no_samples > len(self._samples):
14             return random.sample(self._samples, len(self._samples))
15         else:
16             return random.sample(self._samples, no_samples)
17
18 class GameRunner:
```



```
def __init__(self, sess, model, env, memory, max_eps, min_eps,
    decay, render=True):
    2
        self._sess = sess
    4
        self._env = env
        self._model = model
    6
        self._memory = memory
        self._render = render
    8
        self._max_eps = max_eps
        self._min_eps = min_eps
    10
        self._decay = decay
        self._eps = self._max_eps
    12
        self._steps = 0
        self._reward_store = []
    14
        self._max_x_store = []

    16
def run(self):
    18
        state = self._env.reset()
    20
        tot_reward = 0
```

```
max_x = -100
2
while True:
4     if self._render:
        self._env.render()
6
        action = self._choose_action(state)
8     next_state, reward, done, info = self._env.step(action)
        if next_state[0] >= 0.1:
10         reward += 10
        elif next_state[0] >= 0.25:
12         reward += 20
        elif next_state[0] >= 0.5:
14         reward += 100
16
        if next_state[0] > max_x:
            max_x = next_state[0]
18
        # is the game complete? If so, set the next state to
20        # None for storage sake
```

```
2         if done:
3             next_state = None
4
5         self._memory.add_sample((state, action, reward,
6 next_state))
7         self._replay()
8
9         # exponentially decay the eps value
10        self._steps += 1
11        self._eps = MIN_EPSILON + (MAX_EPSILON - MIN_EPSILON) \
12            * math.exp(-LAMBDA * self.
13 _steps)
14
15        # move the agent to the next state and accumulate the
16 reward
17        state = next_state
18        tot_reward += reward
19
20        # if the game is done, break the loop
21        if done:
22            self._reward_store.append(tot_reward)
23            self._max_x_store.append(max_x)
24            break
```

```
2     print("Step {}, Total reward: {}, Eps: {}".format(self._
3         _steps, tot_reward, self._eps))
4
5
6     def _choose_action(self, state):
7
8         if random.random() < self._eps:
9             return random.randint(0, self._model._num_actions - 1)
10        else:
11            return np.argmax(self._model.predict_one(state, self._
12                _sess))
13
14    def _replay(self):
15
16        batch      = self._memory.sample(self._model._batch_size)
17        states     = np.array([val[0] for val in batch])
18        next_states = np.array([(np.zeros(self._model._num_states)
19                                if val[3] is None else val[3]) for
20            val in batch])
21
22        # predict Q(s,a) given the batch of states
```

```
2     q_s_a = self._model.predict_batch(states, self._sess)
4
6     # predict Q(s',a') - so that we can do gamma * max(Q(s'a'))
7     # below
8     q_s_a_d = self._model.predict_batch(next_states, self._sess
9 )
10
12    # setup training arrays
13    x = np.zeros((len(batch), self._model._num_states))
14    y = np.zeros((len(batch), self._model._num_actions))
15    for i, b in enumerate(batch):
16
17        state, action, reward, next_state = b[0], b[1], b[2], b
18
19    [3]
20
21    # get the current q values for all actions in state
22    current_q = q_s_a[i]
23
24    # update the q value for action
25    if next_state is None:
26        # in this case, the game completed after action, so
27        # there is no max Q(s',a')
28        # prediction possible
29        current_q[action] = reward
```



```

    else:
2         current_q[action] = reward + GAMMA * np.amax(
          q_s_a_d[i])
          x[i] = state
4         y[i] = current_q

6         self._model.train_batch(self._sess, x, y)
8
10 if __name__ == "__main__":
12     env_name = 'MountainCar-v0'
13     env_name = 'Acrobot-v1'
14     env = gym.make(env_name)

16     num_states = env.env.observation_space.shape[0]
17     num_actions = env.env.action_space.n
18
19     model = Model(num_states, num_actions, BATCH_SIZE)
20     mem = Memory(50000)
```

```
2   with tf.Session() as sess:
3       sess.run(model._var_init)
4       gr = GameRunner(sess, model, env, mem, MAX_EPSILON,
5           MIN_EPSILON, LAMBDA)
6       num_episodes = 300
7       cnt = 0
8       while cnt < num_episodes:
9           if cnt % 10 == 0:
10              print('Episode {} of {}'.format(cnt+1, num_episodes))
11          ))
12          gr.run()
13          cnt += 1
14
15          plt.plot(gr._reward_store)
16          plt.show()
17          plt.close("all")
18          plt.plot(gr._max_x_store)
19          plt.show()
```

We will rely on the Stochastic Policy Gradient Theorem

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\pi_{\theta}} [\nabla_{\theta} \log \pi_{\theta}(s, a) Q^{\pi_{\theta}}(s, a)] . \quad (146)$$

Hence, the computation of the policy gradient reduces to a simple expectation.

Policy modeling: parameterized by a function θ : $\pi_{\theta}(a|s)$

$$J(\theta) = \sum_{s \in \mathcal{S}} d^{\pi}(s) V^{\pi}(s) = \sum_{s \in \mathcal{S}} d^{\pi}(s) \sum_{a \in \mathcal{A}} \pi_{\theta}(a|s) Q^{\pi}(s, a) \quad (147)$$

where $d^{\pi}(s)$ is the stationary distribution of Markov chain for π_{θ} for which

$$d^{\pi}(s) = \lim_{t \rightarrow \infty} P(s_t = s | s_0, \pi_{\theta}) \quad (148)$$

and this is the probability that $s_t = s$ when starting from s_0 and following policy π_{θ} for t steps.

Problems:

- in generalized policy iteration, the policy improvement step $\arg \max_{a \in \mathcal{A}} Q^\pi(s, a)$ requires a full scan of the action space, suffering from the **curse of dimensionality**

$$\nabla_\theta J(\theta) = \nabla_\theta \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} Q^\pi(s, a) \pi_\theta(a|s) \quad (149)$$

$$\propto \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} Q^\pi(s, a) \nabla_\theta \pi_\theta(a|s) \quad (150)$$

$$\begin{aligned}
 \nabla_{\theta} V^{\pi}(s) &= \phi(s) + \sum_a \pi_{\theta}(a|s) \sum_{s'} P(s'|s, a) \nabla_{\theta} V^{\pi}(s') \\
 &= \phi(s) + \sum_{s'} \sum_a \pi_{\theta}(a|s) P(s'|s, a) \nabla_{\theta} V^{\pi}(s') \\
 &= \phi(s) + \sum_{s'} \rho^{\pi}(s \rightarrow s', 1) \nabla_{\theta} V^{\pi}(s') \\
 &= \phi(s) + \sum_{s'} \rho^{\pi}(s \rightarrow s', 1) \nabla_{\theta} V^{\pi}(s') \\
 &= \phi(s) + \sum_{s'} \rho^{\pi}(s \rightarrow s', 1) [\phi(s') + \sum_{s''} \rho^{\pi}(s' \rightarrow s'', 1) \nabla_{\theta} V^{\pi}(s'')] \\
 &= \phi(s) + \sum_{s'} \rho^{\pi}(s \rightarrow s', 1) \phi(s') + \sum_{s''} \rho^{\pi}(s \rightarrow s'', 2) \nabla_{\theta} V^{\pi}(s'') \\
 &= \phi(s) + \sum_{s'} \rho^{\pi}(s \rightarrow s', 1) \phi(s') + \sum_{s''} \rho^{\pi}(s \rightarrow s'', 2) \phi(s'') + \sum_{s'''} \rho^{\pi}(s \rightarrow s''', 3) \nabla_{\theta} V^{\pi}(s''') \\
 &= \dots; \text{ Repeatedly unrolling the part of } \nabla_{\theta} V^{\pi}(\cdot) \\
 &= \sum_{x \in \mathcal{S}} \sum_{k=0}^{\infty} \rho^{\pi}(s \rightarrow x, k) \phi(x) .
 \end{aligned}$$

The nice rewriting above allows us to exclude the derivative of Q-value function $\nabla_{\theta} Q^{\pi}(s, a)$.

$$\nabla_{\theta} J(\theta) = \nabla_{\theta} V^{\pi}(s_0) \quad (160)$$

$$= \sum_s \sum_{k=0}^{\infty} \rho^k \pi(s_0 \rightarrow s, k) \phi(s) \quad (161)$$

$$= \sum_s \eta(s) \phi(s) \quad (162)$$

$$= \left(\sum_s \eta(s) \right) \sum_s \frac{\eta(s)}{\sum_s \eta(s)} \phi(s) \quad (163)$$

$$\propto \sum_s \frac{\eta(s)}{\sum_s \eta(s)} \phi(s) \quad (164)$$

$$= \sum_s d^{\pi}(s) \sum_a \nabla_{\theta} \pi_{\theta}(a|s) Q^{\pi}(s, a) . \quad (165)$$

In the episodic case, the constant of proportionality ($\sum_s \eta(s)$) is the average length of an episode.

$$\nabla_{\theta} J(\theta) \quad \propto \sum_{s \in \mathcal{S}} d^{\pi}(s) \sum_{a \in \mathcal{A}} Q^{\pi}(s, a) \nabla_{\theta} \pi_{\theta}(a|s) \quad (166)$$

$$= \sum_{s \in \mathcal{S}} d^{\pi}(s) \sum_{a \in \mathcal{A}} \pi_{\theta}(a|s) Q^{\pi}(s, a) \frac{\nabla_{\theta} \pi_{\theta}(a|s)}{\pi_{\theta}(a|s)} \quad (167)$$

$$= \mathbb{E}_{\pi} [Q^{\pi}(s, a) \nabla_{\theta} \ln \pi_{\theta}(a|s)] \quad (168)$$

where \mathbb{E}_{π} refers to $\mathbb{E}_{s \sim d_{\pi}, a \sim \pi_{\theta}}$ when both state and action distributions follow the policy π_{θ} (on policy).

The policy gradient theorem lays the theoretical foundation for various policy gradient algorithms. This vanilla policy gradient update has no bias but high variance. Many following algorithms were proposed to reduce the variance while keeping the bias unchanged.

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\pi} [Q^{\pi}(s, a) \nabla_{\theta} \ln \pi_{\theta}(a|s)] . \quad (169)$$



Genetic Algorithms

Genetic algorithms apply the principles derived from Darwin's principles (natural selection):

- Individuals in population compete for resources.
- Fittest individuals mate to create more offsprings than others.
- Fittest parent propagates genes through generation; parents may produce offsprings better than either parent.
- Generation are coupled to the environment.

The objective is to maintains the population of n individuals along with their fitness scores. Hence, central to genetic algorithms are the notions of population and the fitness function. Each iteration generates from an initial population a new one. There are three operators operating on the individuals from the population:

- **Selection**

Individuals with better fitness scores pass genes on to successive generations.

- **Crossover**

The selection operator is applied to select two individuals, and randomly choose crossover sites to exchange the genes at these sites.

■ Mutation

Insert random genes in offsprings to maintain diversity.

The operation is assumed to be applied to **genes**, often represented by a sequence from an alphabet Σ :

ADEAGEF

Algorithm 12 Generic Genetic Algorithm

- 1: Generate the initial population
 - 2: Compute fitness
 - 3: **repeat**
 - 4: Selection
 - 5: Crossover
 - 6: Mutation
 - 7: Compute fitness
 - 8: **until** population has converged
-

Let us look at the survival probability of individual i with fitness f_i . One possibility is

$$P_i = \frac{f_i}{\sum_i f_i} . \quad (170)$$

Genetic algorithms have the following advantages:

- No gradients are required
- Can be parallelized
- Can optimize continuous as well as discrete functions
- Can be applied to multi-objective problems

```
1 # Source: https://www.geeksforgeeks.org/genetic-algorithms/  
2 # Python3 program to create target string, starting from  
3 # random string using Genetic Algorithm  
  
5 import random  
  
7 # Number of individuals in each generation  
POPULATION_SIZE = 100  
  
9  
11 # Valid genes  
GENES = '''abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMN  
12 OPQRSTUVWXYZ 1234567890, .-;:_!"#%&/()=?@${[]}'''  
13  
15 # Target string to be generated  
TARGET = "I love GeeksforGeeks"  
  
17 class Individual(object):  
18     '''  
19     Class representing individual in population  
20     '''
```

```
1  def __init__(self, chromosome):
2      self.chromosome = chromosome
3      self.fitness = self.cal_fitness()
4
5
6  @classmethod
7  def mutated_genes(self):
8      '''
9      create random genes for mutation
10     '''
11
12     global GENES
13     gene = random.choice(GENES)
14     return gene
15
16 @classmethod
17 def create_gnome(self):
18     '''
19     create chromosome or string of genes
20     '''
21
22     global TARGET
23     gnome_len = len(TARGET)
```

```
    return [self.mutated_genes() for _ in range(gnome_len)]
2
def mate(self, par2):
4     '''
    Perform mating and produce new offspring
6     '''

8     # chromosome for offspring
    child_chromosome = []
10    for gp1, gp2 in zip(self.chromosome, par2.chromosome):

12        # random probability
        prob = random.random()

14

16        # if prob is less than 0.45, insert gene
        # from parent 1
        if prob < 0.45:
18            child_chromosome.append(gp1)

20        # if prob is between 0.45 and 0.90, insert
```

```
2     # gene from parent 2
3     elif probab < 0.90:
4         child_chromosome.append(gp2)
5
6     # otherwise insert random gene(mutate),
7     # for maintaining diversity
8     else:
9         child_chromosome.append(self.mutated_genes())
10
11    # create new Individual(offspring) using
12    # generated chromosome for offspring
13    return Individual(child_chromosome)
14
15    def cal_fitness(self):
16        '''
17        Calculate fitness score, it is the number of
18        characters in string which differ from target
19        string.
20        '''
21
22    global TARGET
```

```
    fitness = 0
2   for gs, gt in zip(self.chromosome, TARGET):
        if gs != gt: fitness+= 1
4   return fitness

6 # Driver code
def main():
8   global POPULATION_SIZE

10  #current generation
    generation = 1

12

    found = False
14  population = []

16  # create initial population
    for _ in range(POPULATION_SIZE):
18      gnome = Individual.create_gnome()
        population.append(Individual(gnome))
```



```
while not found:
2
    # sort the population in increasing order of fitness score
4    population = sorted(population, key = lambda x:x.fitness)

6    # if the individual having lowest fitness score ie.
    # 0 then we know that we have reached to the target
8    # and break the loop
    if population[0].fitness <= 0:
10        found = True
        break

12
    # Otherwise generate new offsprings for new generation
14    new_generation = []

16
    # Perform Elitism, that mean 10% of fittest population
    # goes to the next generation
18    s = int((10*POPULATION_SIZE)/100)
    new_generation.extend(population[:s])
```

```
2 # From 50% of fittest population, Individuals
3 # will mate to produce offspring
4 s = int((90*POPULATION_SIZE)/100)
5 for _ in range(s):
6     parent1 = random.choice(population[:50])
7     parent2 = random.choice(population[:50])
8     child = parent1.mate(parent2)
9     new_generation.append(child)
10
11 population = new_generation
12
13 print("Generation: {} \tString: {} \tFitness: {}".\
14       format(generation,
15             "".join(population[0].chromosome),
16             population[0].fitness))
17
18 generation += 1
19
20 print("Generation: {} \tString: {} \tFitness: {}".\
```




```
format(generation,  
2     "".join(population[0].chromosome),  
       population[0].fitness))  
4  
if __name__ == '__main__':  
6     main()
```



Excercises

Exercise 1: Single Layer Perceptron

Consider a simple perceptron (see Figure): what will the output be when the input is $(0, 0)$? What about inputs $(0, 1)$, $(1, 1)$ and $(1, 0)$? What if we change the bias weight to -0.5 ?



Exercise 2: **Basis Functions**

Given a test vector x_i , the output of a neural network is defined as

$$f(x_i) = \sum_{j=0}^M w_j \phi_j(x_i, v_j) . \quad (171)$$

The weights of the neurons can be learned by employing the back-propagation rule with sample-based gradient descent. In the lecture neural networks with sigmoid neurons have been introduced, but it is possible to employ different basis functions:

- Which properties do these basis functions have to fulfill?
- Is the number of parameters for $\phi(x_i, v_j)$ limited? Could several different basis functions be used for the same neural network?

Exercise 3: **Error Convergence**

Given 2-1 network trained with one single pattern by means of back-propagation of error and learning rate $\eta = 0.1$. Let the pattern (p, t) be defined by $p = (p_1, p_2) = (0.3, 0.7)$. Verify whether the error

$$E = \frac{1}{2}(t - y)^2 \quad (172)$$

converges and if so, at what value?



Bibliography

References

- [1] Maxwell W. Libbrecht and William Stafford Noble. Machine learning applications in genetics and genomics. *Nature Reviews Genetics*, 16(6):321–332, 2015.
- [2] Andrey Kan. Machine learning applications in cell image analysis. *Immunology & Cell Biology*, 95(6):525–530, 2020/02/24 2017.
- [3] Andrew W. Senior, Richard Evans, John Jumper, James Kirkpatrick, Laurent Sifre, Tim Green, Chongli Qin, Augustin Židek, Alexander W. R. Nelson, Alex Bridgland, Hugo Penedones, Stig Petersen, Karen Simonyan, Steve Crossan, Pushmeet Kohli, David T. Jones, David Silver, Koray Kavukcuoglu, and Demis Hassabis. Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792):706–710, 2020.
- [4] Ning Qian. On the momentum term in gradient descent learning algorithms. *Neural Networks*, 12(1):145–151, 1999.
- [5] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv e-prints*, page arXiv:1412.6980, 12 2014.

- [6] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mane, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viegas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv e-prints*, page arXiv:1603.04467, 03 2016.
- [7] W. S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5:115–133, 1943.
- [8] Frank Rosenblat. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Spartan Books, 1962.
- [9] Marvin Minsky and Seymour A. Papert. *Perceptrons: An introduction to Computational Geometry*. MIT Press, 1969.
- [10] Proc. Nat. Acad. Sci. 79 2554-2558 J. J. Hopfield. 1982. 1982.
- [11] J. Reinitz E. Mjolsness, D. H. Sharp. *J. Theor. Biol.*, 152:429–453, 1991.
- [12] Mechanisms of Development 49 133-158 J. Reinitz, D. H. Sharp. 1995. 1995.

- [13] M. Kaufman R. Thomas, D. Thieffry. *Bul. Math. Biol.*, 57(2):247–276, 1995.
- [14] Anders Krogh, Michael Brown, I. Saira Mian, Kimmen Sjölander, and David Haussler. Hidden markov models in computational biology: Applications to protein modeling. *Journal of Molecular Biology*, 235(5):1501–1531, 1994.
- [15] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, 1992.
- [16] R. S. Sutton and A. G. Barto. *Reinforcement Learning : An Introduction*. Cambridge, MA, USA: MIT Press, 1998.
- [17] Schulman J. Wolski F. Dhariwal P. Radford A. and Klimov O. Proximal policy optimization algorithms. In *International Conference on Learning Representations*, 2017.
- [18] Razvan V. Florian. Correct equations for the dynamics of the cart-pole system. Technical report, 2020.
- [19] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [20] <https://arxiv.org/abs/1506.02438>, 2018.



Index

- ϵ -greedy strategy, 116
- cumulative discounted return, 69
- action policy, 69
- action-state function, 69, 110
- Actor-Critic policy gradient algorithm, 114
- Adaptive Moment Estimation (Adam), 15
- advantage function, 110
- advantage, reinforcement learning, 70
- ANN , 17
- backpropagation, 24
- Baum-Welch algorithm, 63
- Bellman Equation, 91
- classification, neural network, 32
- continuous state and action space, 131

CpG, 48

crossover, genetic algorithm, 137

curse of dimensionality, 132

decoding problem, 56

deep learning, 13

discount parameter, 69

environment, fully observable, 11

episode, reinforcement learning, 73

evaluation problem, 56

excitatory synapses, 19

expectation maximization (EM) algorithm, 63

fitness function, 137

forward algorithm, hidden Markov model, 57

forward-backward algorithm, hidden Markov model, 63

gradient descent, 13

gradient descent optimization algorithms, 15

Hadamard product, 24

HMM, hidden Markov model, 44

horizon, reinforcement learning, 67

inhibitory synapses, 19

learning problem, 56

learning rate, 80

learning rate schedule, 13

logistic function, 20

Markov chain, 131

Markov chain, reinforcement learning, 67

Markov Decision Process, MDP, 8

Markov reward process, MRP, 9

mini-batch gradient descent, 14

Monte Carlo policy gradient, 81

- Monte-Carlo, Policy Evaluation, 81
- Monte-Carlo, Policy Evaluation with Neural Network, 92
- multiple sequence alignments, 54
- mutation, genetic algorithm, 138
- n-gram Models, 61
- natural selection, 137
- Neural Networks, 17
- neural networks, artificial, 17
- perceptron, 18
- play out, reinforcement learning, 73
- policy gradient, 68
- policy, reinforcement learning, 66
- population, genetic algorithm, 137
- Q-Learning, 112
- REINFORCE, 81

reinforcement learning, 65
return, reinforcement learning, 65
reward function, 9, 70
reward, reinforcement learning, 65
Schur product, 24
selection, genetic algorithm, 137
SGD, with momentum, 15
sigmoid neuron, 19
Smith-Waterman algorithm, 54
softmax, 72
state space, 9
stochastic gradient descent, 13
stochastic gradient policy theorem, 73
tensorflow, 15
value function fitting, 118

value-state function, 69, 110

values policy, 69

Viterbi algorithm, 59

weighted input, 25