

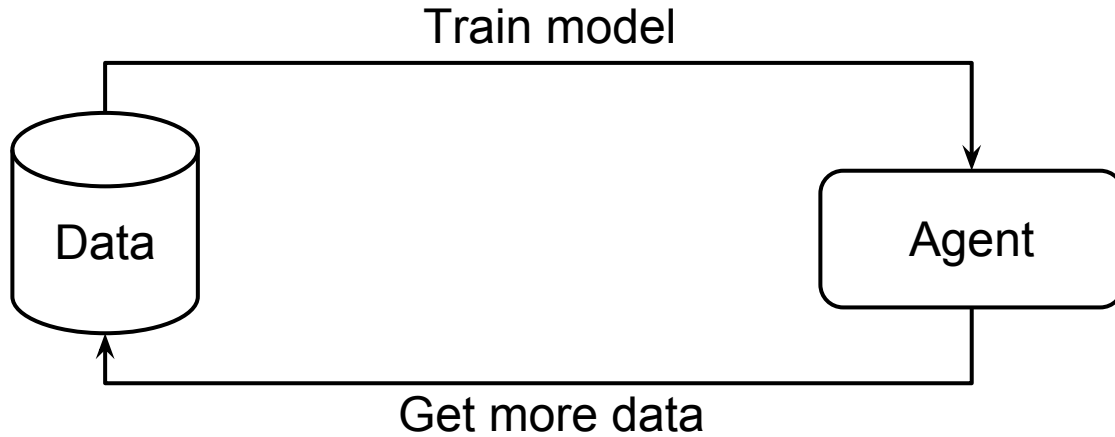
Tensorflow Overview

CS294-112: Deep Reinforcement Learning


Goal of this course

Train an agent to perform useful tasks

Goal of this course



Goal of this course



Machine learning

How to do this?

$$\theta^* = \arg \min_{\theta} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \| \underline{f_{\theta}(\mathbf{x})} - \mathbf{y} \|$$


gradient descent

neural networks


What is Tensorflow?

- Library for
 - Defining computation graphs
 - Calculating gradients

Tensorflow: defining computation graphs



Tensorflow: calculating gradients



input layer

hidden layer 1

hidden layer 2

output layer

$$\mathbf{h}_1 = \sigma(\mathbf{W}_1 \mathbf{x})$$

$$\mathbf{h}_2 = \sigma(\mathbf{W}_2 \mathbf{h}_1)$$


$$\mathbf{y} = \sigma(\mathbf{W}_3 \mathbf{h}_2)$$

$$\frac{\partial \mathbf{y}}{\partial \mathbf{W}_1} = \frac{\partial \mathbf{y}}{\partial \mathbf{h}_2} \frac{\partial \mathbf{h}_2}{\partial \mathbf{h}_1} \frac{\partial \mathbf{h}_1}{\partial \mathbf{W}_1}$$

$$\frac{\partial \mathbf{y}}{\partial \mathbf{W}_2} = \frac{\partial \mathbf{y}}{\partial \mathbf{h}_2} \frac{\partial \mathbf{h}_2}{\partial \mathbf{W}_2}$$

$$\frac{\partial \mathbf{y}}{\partial \mathbf{W}_3}$$

Tensorflow: calculating gradients



input layer

hidden layer 1

hidden layer 2

output layer

You define

$$\mathbf{h}_1 = \sigma(\mathbf{W}_1 \mathbf{x})$$

$$\mathbf{h}_2 = \sigma(\mathbf{W}_2 \mathbf{h}_1)$$

$$\mathbf{y} = \sigma(\mathbf{W}_3 \mathbf{h}_2)$$

TF calculates

$$\frac{\partial \mathbf{y}}{\partial \mathbf{W}_1} = \frac{\partial \mathbf{y}}{\partial \mathbf{h}_2} \frac{\partial \mathbf{h}_2}{\partial \mathbf{h}_1} \frac{\partial \mathbf{h}_1}{\partial \mathbf{W}_1}$$

$$\frac{\partial \mathbf{y}}{\partial \mathbf{W}_2} = \frac{\partial \mathbf{y}}{\partial \mathbf{h}_2} \frac{\partial \mathbf{h}_2}{\partial \mathbf{W}_2}$$

$$\frac{\partial \mathbf{y}}{\partial \mathbf{W}_3}$$

Alternatives to Tensorflow

PyTorch, Caffe, Theano, MXNet, Chainer,

fundamentally the same
implementation differences

Hands-on with Tensorflow

<switch to Jupyter notebook>