

15-780: Deep Learning

J. Zico Kolter

March 1 – April 4, 2016

Outline

Introduction

Machine learning with neural networks

Training neural networks

Convolutional neural networks

Recurrent neural networks

Deep reinforcement learning

Outline

Introduction


Machine learning with neural networks

Training neural networks


Convolutional neural networks

Recurrent neural networks

Deep reinforcement learning





Google scholar counts of papers containing “neural network” divided by count of papers containing “machine learning”



A non-exhaustive list of some of the important events that impacted this trend

“AlexNet” (Krizhevsky et al., 2012), winning entry of ImageNet 2012 competition with a Top-5 error rate of 15.3% (next best system with highly engineered features based upon SIFT got 26.1% error)






Some classification results from AlexNet

Google Deep Dream software: adjust input images (by gradient descent) to strengthen the activations that are present in an image



Question answering network (Vinyals and Le, 2015), using sequence to sequence learning method (Sutskever et al., 2014)



Machine: *hi*

Human: *hello*

Machine: *could you please let me know what are the operating systems you are using ?*

Human: *i am using linux*

Machine: *ok it was in the ssl vpn on which you stated it was broken*

Machine: *please check if you are logged in to corp network*

Machine: *what is the error that you are running please*

Human: *i am seeing an error related to vpn*

Machine: *what is the error message that you are getting when connecting to vpn using network connect ?*

Human: *connection refused or something like that*

Machine: *may i know the version of network connect you connect ?*

Human: *i am not sure i know that*

Machine: *is the network connect prompting for any user certificate*

Machine: *?*

Human: *yes*

Machine: *are you at home or at work at the moment ?*

Machine: *office ?*

Human: *i am at home*

Machine: *try this*

Machine: *goto <URL> page and enter your name and password*

Human: *done*

Machine: *try logging in with <NAME> and check if its working fine now*

Human: *yes , now it works !*

Machine: *great*

Machine: *anything else that i can help*

Human: *no , that was very useful*


Machine: *you are other than <NAME> :)*

Machine: *have a nice night*

AlphaGo (Silver et al., 2016) beats Lee Sedol in 5 game competition

Artificial intelligence: Google's AlphaGo beats Go master Lee Se-dol

12 March 2016 | Technology



A computer program has beaten a master Go player 3-0 in a best-of-five competition, in what is seen as a landmark moment for artificial intelligence.

Google's AlphaGo program was playing against Lee Se-dol in Seoul, in South Korea.

Outline

Introduction

Machine learning with neural networks

Training neural networks

Convolutional neural networks

Recurrent neural networks

Deep reinforcement learning

Supervised learning setup

Input features $x^{(i)} \in \mathbb{R}^n$

Outputs $y^{(i)} \in \mathcal{Y}$ (e.g. \mathbb{R} , $\{-1, +1\}$, $\{1, \dots, p\}$)

Model parameters $\theta \in \mathbb{R}^k$

Hypothesis function $h_\theta : \mathbb{R}^n \rightarrow \mathbb{R}$

Loss function $\ell : \mathbb{R} \times \mathcal{Y} \rightarrow \mathbb{R}_+$

Machine learning optimization problem

$$\underset{\theta}{\text{minimize}} \quad \sum_{i=1}^m \ell(h_\theta(x^{(i)}), y^{(i)})$$

We mainly considered the linear hypothesis class

$$h_{\theta}(x^{(i)}) = \theta^T \phi(x^{(i)})$$

for some set of non-linear features $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^k$

(Note: previously, we just directly included the non-linear features in $x^{(i)}$, but here we separate them for clarity)

Example

$$x^{(i)} = [\text{temperature for day } i]$$

$$\phi(x^{(i)}) = \begin{bmatrix} 1 \\ x^{(i)} \\ x^{(i)2} \\ \vdots \end{bmatrix}$$

Challenges with linear models

Linear models crucially depend on choosing “good” features

Some “standard” choices: polynomial features, radial basis functions, random features (surprisingly effective)

But, many specialized domains required highly engineered special features

- E.g., computer vision tasks used Haar features, SIFT features, every 10 years or so someone would engineer a new set of features

Key question: can we come up with an algorithm that will automatically *learn* the features themselves?

Feature learning, take one

Instead of a simple linear classifier, let's consider a two-stage hypothesis class where one linear function creates the features, another models the classifier


$$h_{\theta}(x) = W_2\phi(x) + b_2 = W_2(W_1x + b_1) + b_2$$

where

$$\theta = \{ W_1 \in \mathbb{R}^{n \times k}, b_1 \in \mathbb{R}^k, W_2 \in \mathbb{R}^{1 \times k}, b_2 \in \mathbb{R} \}$$

Note that in this notation, we're explicitly separating the parameters on the “constant feature” into the b terms

Graphical depiction of the above function



But there is a problem:

$$h_{\theta}(x) = W_2(W_1x + b_1) + b_2 = \tilde{W}x + \tilde{b} \quad (1)$$

in other words, we are still just using a normal linear classifier (the apparent added complexity is not giving us any additional representational power)

Neural networks




Neural networks are a simple extension of this idea, where we additionally apply a *non-linear* function after each linear transformation

$$h_{\theta}(x) = f_2(W_2 f_1(W_1 x + b_1) + b_2)$$


where $f_1, f_2 : \mathbb{R} \rightarrow \mathbb{R}$ are some non-linear functions (applied elementwise to vectors)

Common choices for f_i are hyperbolic tangent

$\tanh(x) = (e^{2x} - 1)/(e^{2x} + 1)$, sigmoid $\sigma(x) = 1/(1 + e^{-x})$, or
rectified linear unit $f(x) = \max\{0, x\}$



We draw these the same as before (non-linear functions are virtually always implied in the neural network setting)



Middle layer z is referred to as the *hidden layer* or *activations*

These are the learned features, nothing in the data that prescribes what values these should take, left up to the algorithm to decide

Properties of neural networks


It turns out that a neural network with a single hidden layer (and a suitably large number of hidden units) is a *universal function approximator*, can approximate *any* function over the input arguments (but this is actually not very useful in practice, c.f. polynomials fitting any sets of points for high enough degree)

The hypothesis class h_θ is not a convex function of the parameters $\theta = \{W_i, b_i\}$, so we must resort to non-convex optimization methods

Architectural choices (how many layers, how are they connected), become important free parameters, more on this later

Deep learning

“Deep” neural networks refer to networks with multiple hidden layers



Mathematically, a k -layer network has the hypothesis function

$$z_{i+1} = f_i(W_i z_i + b_i), \quad i = 1, \dots, k-1, \quad z_1 = x$$

$$h_\theta(x) = z_k$$

where z_i terms now indicate vectors, not entries into a vector

Why use deep networks?

Motivation from circuits: many functions can be represented more compactly using deep networks than one-hidden layer networks (e.g. parity function would require (2^n) hidden units in 3-layer network, $O(n)$ units in $O(\log n)$ -layer network)

Motivation from neurobiology: brain appears to use multiple levels of interconnected neurons to process information (but careful, neurons in brain are not just non-linear functions)

In practice: works better for many domains

Outline

Introduction

Machine learning with neural networks

Training neural networks

Convolutional neural networks

Recurrent neural networks

Deep reinforcement learning

Optimizing neural network parameters

How do we optimize the parameters for the machine learning loss minimization problem with a neural network

$$\underset{\theta}{\text{minimize}} \quad \sum_{i=1}^m \ell(h_{\theta}(x^{(i)}), y^{(i)})$$

now that this problem is non-convex?

Just do exactly what we did before: initialize with random weights and run stochastic gradient descent

Now have the possibility of local optima, and function can be harder to optimize, but we won't worry about all that because the resulting models still often perform better than linear models

Stochastic gradient descent for neural networks

Recall that stochastic gradient descent computes gradients with respect to loss on each example, updating parameters as it goes

function SGD($\{(x^{(i)}, y^{(i)})\}, h_\theta, \ell, \alpha$)

 Initialize: $W_j, b_j \leftarrow \text{Random}, j = 1, \dots, k$

Repeat until convergence:

For $i = 1, \dots, m$:

 Compute $\nabla_{W_j, b_j} \ell(h_\theta(x^{(i)}), y^{(i)}), j = 1, \dots, k - 1$

 Take gradient steps in all directions:

$W_j \leftarrow W_j - \alpha \nabla_{W_j} \ell(h_\theta(x^{(i)}), y^{(i)}), j = 1, \dots, k$

$b_j \leftarrow b_j - \alpha \nabla_{b_j} \ell(h_\theta(x^{(i)}), y^{(i)}), j = 1, \dots, k$

return $\{W_j, b_j\}$

So how do we compute the gradients $\nabla_{W_j, b_j} \ell(h_\theta(x^{(i)}), y^{(i)})$, this is a complex function of the parameters

Backpropagation

Backpropagation is a method for computing all the necessary gradients using one “forward pass” (just computing all the values at layers), and one “backward pass” (computing gradients backwards in the network)

The equations sometimes look complex, but it’s just an application of the chain rule of calculus

The Jacobian

One (last!) bit of multivariate calculus will help us derive the backpropagation algorithm using purely matrix and vector operations

For a multivariate, vector-valued function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, the *Jacobian* is a $m \times n$ matrix

$$\left(\frac{\partial f(x)}{\partial x} \right) \in \mathbb{R}^{m \times n} = \begin{bmatrix} \frac{\partial f_1(x)}{\partial x_1} & \frac{\partial f_1(x)}{\partial x_2} & \dots & \frac{\partial f_1(x)}{\partial x_n} \\ \frac{\partial f_2(x)}{\partial x_1} & \frac{\partial f_2(x)}{\partial x_2} & \dots & \frac{\partial f_2(x)}{\partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_m(x)}{\partial x_1} & \frac{\partial f_m(x)}{\partial x_2} & \dots & \frac{\partial f_m(x)}{\partial x_n} \end{bmatrix}$$

For a scalar-valued function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the Jacobian is the transpose of the gradient $\frac{\partial f(x)}{\partial x}^T = \nabla_x f(x)$

We'll use a few simple properties of the Jacobian to derive the backpropagation algorithm for neural networks

Chain rule

$$\frac{\partial f(g(x))}{\partial x} = \frac{\partial f(g(x))}{\partial g(x)} \frac{\partial g(x)}{\partial x}$$

Jacobian of a linear transformation, for $A \in \mathbb{R}^{m \times n}$

$$\frac{\partial Ax}{\partial x} = A$$

If f is a function applied elementwise,

$$\frac{\partial f(x)}{\partial x} = \text{diag}(f'(x))$$

Derivation of backpropagation

Using the chain rule to compute derivatives:

$$\begin{aligned}\frac{\partial \ell(h_\theta(x), y)}{\partial b_1} &= \frac{\partial \ell(z_k, y)}{\partial b_1} \\ &= \frac{\partial \ell(z_k, y)}{\partial z_k} \frac{\partial z_k}{\partial b_1} \\ &= \frac{\partial \ell(z_k, y)}{\partial z_k} \frac{\partial z_k}{\partial z_{k-1}} \frac{\partial z_{k-1}}{\partial z_{k-2}} \dots \frac{\partial z_3}{\partial z_2} \frac{\partial z_2}{\partial b_1}\end{aligned}$$

Furthermore, for any i

$$\frac{\partial z_{i+1}}{\partial z_i} = \frac{\partial f_i(W_i z_i + b_i)}{\partial W_i z_i + b_i} \frac{\partial W_i z_i + b_i}{\partial z_i} = \text{diag}(f'_i(W_i z_i + b_i)) W_i$$

and

$$\frac{\partial z_{i+1}}{\partial b_i} = \frac{\partial f_i(W_i z_i + b_i)}{\partial W_i z_i + b_i} \frac{\partial W_i z_i + b_i}{\partial b_i} = \text{diag}(f'_i(W_i z_i + b_i))$$

If we compute derivatives with respect to *all* b_i , and W_i just using this formula, we'd be repeating a lot of work (e.g. all the $\frac{\partial z_{i+1}}{\partial z_i}$ terms that appear in multiple derivatives)

Backpropagation caches these intermediate products, specifically defining

$$g_i^T = \frac{\partial \ell(z_k, y)}{\partial z_k} \frac{\partial z_k}{\partial z_{k-1}} \dots \frac{\partial z_{i+1}}{\partial z_i}$$

which can be computed recursively via the relationship

$$g_k = \frac{\partial \ell(z_k, y)}{\partial z_k}$$

$$g_i = W_i^T (g_{i+1} \circ f'(W_i z_i + b_i))$$

where \circ denotes elementwise multiplication of vectors

Gradients can then be computed via

$$\nabla_{b_i} \ell(h_\theta(x), y) = g_{i+1} \circ f'(W_i z_i + b_i)$$

$$\nabla_{W_i} \ell(h_\theta(x), y) = (g_{i+1} \circ f'(W_i z_i + b_i)) z_i^T$$

As mentioned, algorithmically backpropagation takes the form of one forward pass (to compute z_i terms) and one backward pass (to compute g_i terms) through the network

function Backpropagation($x, y, \{W_i, b_i, f_i\}_{i=1}^{k-1}, \ell$)

 Initialize: $z_1 \leftarrow x$

For $i = 1, \dots, k-1$

$z_{i+1}, z'_{i+1} \leftarrow f_i(W_i z_i + b_i), f'_i(W_i z_i + b_i)$

$L \leftarrow \ell(z_k, y)$

$g_k \leftarrow \frac{\partial \ell(z_k, y)}{\partial z_k}$

For $i = k-1, \dots, 1$:

$g_i = W_i^T (g_{i+1} \circ z'_{i+1})$

$\nabla_{b_i} \leftarrow g_{i+1} \circ z'_{i+1}$

$\nabla_{W_i} \leftarrow (g_{i+1} \circ z'_{i+1}) z_i^T$

return $L, \{\nabla_{b_i}, \nabla_{W_i}\}_{i=1}^{k-1}$

Gradients can still get somewhat tedious to derive by hand, especially for the more complex models that follow

Fortunately, a lot of this work has already been done for you

Tools like Theano

(<http://deeplearning.net/software/theano/>), Torch (<http://torch.ch/>), TensorFlow

(<http://www.tensorflow.org/>) all let you specify the network structure and then automatically compute all gradients (and use GPUs to do so)

Autograd package for Python

(<https://github.com/HIPS/autograd>) lets you compute the derivative of (almost) any arbitrary function using numpy operations using automatic backpropagation

What's changed since the 80s?

All these algorithms (and most of the extensions in later slides), were developed in the 80s or 90s

So why are these just becoming more popular in the last few years?

- More data
- Faster computers
- (Some) better optimization techniques

Unsupervised pre-training (Hinton et al., 2006): “Pre-train” the network have the hidden layers recreate their input, one layer at a time, in an unsupervised fashion

- This paper was partly responsible for re-igniting the interest in deep neural networks, but the general feeling now is that it doesn’t help much

Dropout (Hinton et al., 2012): During training and computation of gradients, randomly set about half the hidden units to zero (a different randomly selected set for each stochastic gradient step)

- Acts like regularization, prevents the parameters for overfitting to particular examples

Different non-linear functions (Nair and Hinton, 2010): Use non-linearity $f(x) = \max\{0, x\}$ instead of $f(x) = \tanh(x)$

Outline

Introduction

Machine learning with neural networks

Training neural networks

Convolutional neural networks

Recurrent neural networks


Deep reinforcement learning

The problem with fully-connected networks

A 256x256 (RGB) image \implies ~200,000 dimensional input


A fully connected deep network would need a large number of parameters, very likely to overfit to training data

A generic deep network also doesn't capture much of the "natural" invariances we expect in images (location, scale)



Convolutional neural networks


Constrain the weights in two ways: require that activations between layers occurs only in a “local” manner, and require that activations share the same weights across all locations



We write the convolutions as $z_{i+1} = z_i * W_i$


Greatly reduces the number of parameters, and provides a model that captures translational invariance

We actually use “3D” convolutions to combine multiple channels, and create multiple features using multiple convolutions at each level



(weights at level i can be represented as four-dimensional tensor)

Also common to include *max-pooling* layers that take maximum over regions, to reduce size of layers



Computing gradients in convolutional networks

How do we compute gradient terms in a convolutional model?

Almost the same as with standard backpropagation (convolutions are just linear operators)

Consider 1D convolution: $z_i * w$ for $w \in \mathbb{R}^3$, $z_i \in \mathbb{R}^6$

$$z_{i+1} = f_i(z_i * w + b_i) = f_i(W_i z_i + b)$$

where

$$W_i = \underbrace{\begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}}_{\text{Select "valid" entries}} \underbrace{\begin{bmatrix} w_2 & w_3 & 0 & 0 & 0 & w_1 \\ w_1 & w_2 & w_3 & 0 & 0 & 0 \\ 0 & w_1 & w_2 & w_3 & 0 & 0 \\ 0 & 0 & w_1 & w_2 & w_3 & 0 \\ 0 & 0 & 0 & w_1 & w_2 & w_3 \\ w_3 & 0 & 0 & 0 & w_1 & w_2 \end{bmatrix}}_{\text{Circular convolution}}$$

Computing gradients in backprop involves multiplication by W_i (a convolution) and multiplication by W_i^T


We don't want to actually form the W_i matrix described above, but note that

$$W_i^T = \underbrace{\begin{bmatrix} w_2 & w_1 & 0 & 0 & 0 & w_3 \\ w_3 & w_2 & w_1 & 0 & 0 & 0 \\ 0 & w_3 & w_2 & w_1 & 0 & 0 \\ 0 & 0 & w_3 & w_2 & w_1 & 0 \\ 0 & 0 & 0 & w_3 & w_2 & w_1 \\ w_1 & 0 & 0 & 0 & w_3 & w_2 \end{bmatrix}}_{\text{Convolution with } w \text{ flipped}} \underbrace{\begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}}_{\text{Zero padding}}$$


In other words, we can multiply by “ W_i^T ” by zero-padding the vector and applying another convolution

A few more trivial details to compute gradient of weights, subgradients of max-pooling, etc, but it is still all just the chain rule

Some example networks




LeNet-5 (LeCun et al., 1998) architecture, achieves 1% error in MNIST digit classification




“AlexNet” (Krizhevsky et al., 2012), winner of ILSVRC 2012

ConvNet Configuration					
A	A-LRN	B	C	D	E
11 weight layers	11 weight layers	13 weight layers	16 weight layers	16 weight layers	19 weight layers
input (224 × 224 RGB image)					
conv3-64	conv3-64 LRN	conv3-64 conv3-64	conv3-64 conv3-64	conv3-64 conv3-64	conv3-64 conv3-64
maxpool					
conv3-128	conv3-128	conv3-128 conv3-128	conv3-128 conv3-128	conv3-128 conv3-128	conv3-128 conv3-128
maxpool					
conv3-256 conv3-256	conv3-256 conv3-256	conv3-256 conv3-256	conv3-256 conv3-256 conv1-256	conv3-256 conv3-256 conv3-256	conv3-256 conv3-256 conv3-256 conv3-256
maxpool					
conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512 conv1-512	conv3-512 conv3-512 conv3-512	conv3-512 conv3-512 conv3-512 conv3-512
maxpool					
conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512 conv1-512	conv3-512 conv3-512 conv3-512	conv3-512 conv3-512 conv3-512 conv3-512
maxpool					
FC-4096					
FC-4096					
FC-1000					
soft-max					

VGG Network (Simonyan and Zisserman, 2015), 2nd place in ILSVRC
2014



GoogLeNet (Szegedy et al., 2014), winner of ILSVRC 2014



Deep ResNet (He et al., 2015), winner of ILSVRC 2015

Outline

Introduction

Machine learning with neural networks

Training neural networks

Convolutional neural networks


Recurrent neural networks

Deep reinforcement learning

Predicting temporal data

So far, the models we have discussed have been applicable to independent inputs $x^{(1)}, \dots, x^{(m)}$

In practice, we often want to predict a sequence of outputs, given a sequence of inputs




Just predicting each output independently would miss crucial information

Many examples: time series forecasting, sentence labeling, part of speech tagging, etc

Basic recurrent neural network

Recurrent neural network, maintain hidden state over time, hidden state is function of current input *and previous hidden state*



$$z_2^{(t)} = f_1(W_{xz}x^{(t)} + W_{zz}z^{(t-1)} + b_1)$$

$$\hat{y}^{(t)} = f_2(W_{zy}z^{(t)} + b_2)$$

Training recurrent neural networks

Most common for training recurrent neural networks is to “unroll” prediction on some dataset $x^{(1:T)}, y^{(1:T)}$, minimize the loss function

$$\underset{W_{xz}, W_{zz}, W_{zy}}{\text{minimize}} \quad \sum_{t=1}^T \ell(\hat{y}^{(t)}, y^{(t)})$$


Equivalent to backpropagation in a “deep” network where each layer is constrained to have the same weights

Some issues: initializing first hidden layer (just have a hidden “previous” layer of all zeros), how long of sequence (pick something big, say 100)

Long short term memory (LSTM) networks

Trouble with plain RNNs is that it is difficult to capture long-term dependencies (e.g., if we see a "(" character, we expect a ")" to follow at some point)

One solution, long short term memory (LSTM) network (Hochreiter and Schmidhuber, 1997), has more complex structure that specifically encodes memory and pass-through features, able to capture these longer dependences


$$\begin{aligned} i_t &= \tanh(W_{xi}x_t + W_{hi}h_{t-1} + b_i) \\ j_t &= \text{sigm}(W_{xj}x_t + W_{hj}h_{t-1} + b_j) \\ f_t &= \text{sigm}(W_{xf}x_t + W_{hf}h_{t-1} + b_f) \\ o_t &= \tanh(W_{xo}x_t + W_{ho}h_{t-1} + b_o) \\ c_t &= c_{t-1} \odot f_t + i_t \odot j_t \\ h_t &= \tanh(c_t) \odot o_t \end{aligned}$$

LSTM architecture, figure/equations from (Jozefowicz et al., 2015)

Some example settings

“Char-RNN” network predicts text one character at a time




Figure and subsequent example from Karpathy, 2015,
<http://karpathy.github.io/2015/05/21/rnn-effectiveness/>


Char-RNN trained character-by-character level on the Linux source code

Some random source code sampled from the resulting model

```
/*
 * Increment the size file of the new incorrect UI_FILTER group information
 * of the size generatively.
 */
static int indicate_policy(void)
{
    int error;
    if (fd == MARN_EPT) {
        /*
         * The kernel blank coeld it to userspace.
         */
        if (ss->segment < mem_total)
            unblock_graph_and_set_blocked();
        else
            ret = 1;
        goto bail;
    }
    segaddr = in_SB(in.addr);
    selector = seg / 16;
    setup_works = true;
    ...
}
```

Combining convolution and recurrent networks

Multimodal RNN (Karpathy and Fei-Fei, 2015, though many other similar ideas), uses the output of a convolution neural network as the first input to the hidden units of an RNN



Some resulting image captions



"man in black shirt is playing guitar."



"construction worker in orange safety vest is working on road."



"two young girls are playing with lego toy."



"boy is doing backflip on wakeboard."

Outline

Introduction

Machine learning with neural networks

Training neural networks

Convolutional neural networks

Recurrent neural networks

Deep reinforcement learning

Recall Q-learning setup

Q^* function encodes discounted sum of future rewards

$Q^*(s, a)$ gives value of being in state s , taking action a , and acting optimally thereafter

Q-learning update: in state s , take action a , obtain rewards r and end up in state s'

$$\begin{aligned}\hat{Q}^*(s, a) &\leftarrow (1 - \alpha) \hat{Q}^*(s, a) + \alpha \left(r + \gamma \max_{a'} \hat{Q}^*(s', a') \right) \\ &= \hat{Q}^*(s, a) + \alpha(r + \gamma \max_{a'} \hat{Q}^*(s', a') - \hat{Q}^*(s, a))\end{aligned}$$

Q-learning with function approximation

For large state and action spaces, we clearly can't represent the value function explicitly for each state/action pair

Leads to an equivalent update for Q-learning (and TD methods in general) with function approximation

Given a Q function approximated by parameters θ , $\hat{Q}_\theta^\star(s, a)$, Q-learning update becomes

$$\theta = \theta + \alpha \left(r + \gamma \max_{a'} \hat{Q}_\theta^\star(s', a') - \hat{Q}_\theta^\star(s, a) \right) \nabla_\theta \hat{Q}_\theta^\star(s, a)$$

Deep Q-learning

Deep Q-learning: simply use a neural network (with continuous output) to represent $\hat{Q}_\theta^*(s, a)$, backprop as before to compute gradients

Example: Google Deepmind DQN network used to learn to play Atari games (Mnih et al., 2015)

