

## Class notes 11

Model selection: Interesting extensions:

- Dealing with unknown  $\sigma^2$  (e.g. using unbiased estimates assuming linear model)
- Extending beyond squared loss to likelihood-loss (as in logistic regression): AIC
- Extending to cases where  $op$  cannot be calculated but can be estimated in unbiased manner: Stein's Unbiased Risk Estimator (SURE)

We will discuss SURE on the board as time allows.

## Neural networks and deep learning

**1-Hidden Layer Neural net:** A single layer network (from Wikipedia)

Each arc contains a parameter  $w_{ij}$

Each node sums its inputs

Each of  $K$  hidden nodes (here  $K = 4$ ) applies a non-linear function  $\sigma$  :

$$\hat{f}_l(x) = \sum_{k=1}^K w_{kl}^{(2)} \sigma\left(\sum_{j=1}^p w_{jk}^{(1)} x_j\right),$$

where  $l$  is the index of the output nodes (in regression or two-class classification usually a single output node).

Designing a simple Neural net:

- How many hidden layers, how many nodes in each?
- Which non-linearity  $\sigma$  to use?
- One or more output nodes? For regression or 2-class classification, usually only one
- Given training data  $T$ , how do we learn the weights/parameters  $W$ ?

**Gradient descent for parameter learning:** As in other methods, we start from defining a loss function  $L(y, \hat{y})$  for learning, for example

- Squared loss (RSS) for regression, like in OLS or our boosting example
- Bernoulli log-likelihood (AKA cross-entropy) for classification, like in logistic regression

Since  $\hat{y}_i = \hat{f}(x_i)$  is a function of the parameters  $W$ , we can write the loss as function of  $W$ , in our simple example with one hidden layer, one output node, and squared loss:

$$\sum_{i=1}^n L(y_i, \hat{y}_i) = \sum_{i=1}^n \left( y_i - \sum_{k=1}^K w_k^{(2)} \sigma \left( \sum_{j=1}^p w_{jk}^{(1)} x_{ij} \right) \right)^2$$

Given  $W^{(t-1)}$ , we calculate the derivative (gradient) of the loss function relative to every parameter  $\nabla L^{(t)}$ . Dimension of  $\nabla L$  is the number of parameters in our network — can be millions or billions. We then take a *downward* step to decrease the loss

$$W^{(t)} = W^{(t-1)} - \epsilon \nabla L^{(t)}$$

We repeat that for many iterations

## Two important notions: back-propagation and stochastic gradient descent (SGD)

A big part of the Neural Nets knowledge-base deals with ways of updating the weights based on gradient which are:

- Statistically efficient (including regularization, implicitly or explicitly)
- Computationally efficient (how do we get somewhere in such a huge parameter space, with non-convex functions of parameters?)

Back-prop is an efficient approach for taking advantage of the graphical model structure for calculating elements of  $\nabla L^{(t)}$  of the form:

$$\frac{\partial \sum_{i=1}^n L(y_i, \hat{y}_i^{(t-1)})}{\partial w_{jk}^{(l)}}$$

In mathematical terms, back-prop is simply using the chain rule to differentiate for weights that are "deep" in the function, and uses the layer structure to calculate it efficiently. More generally, the important idea is that of a *computation graph* that allows to efficiently calculate the gradient relative to all parameters.

SGD uses the idea that:

$$\frac{\partial \sum_{i=1}^n L(y_i, \hat{y}_i)}{\partial w_{jk}^{(l)}} = \sum_{i=1}^n \frac{\partial L(y_i, \hat{y}_i)}{\partial w_{jk}^{(l)}}$$

SGD calculates the term inside the sum for one observation (or small group of observations) every time, and updates the weights. Because the number of weights can be very large, backprop + SGD allow updating all of them with relatively little computation. In practice, this converges much faster than calculating  $\nabla L$  over all observations before updating.

## From Neural Nets to deep learning

Neural nets have been around for decades, initially proposed as "imitating the brain". They went in and out of fashion, not the leading choice for predictive modeling as of 2005 (Boosting was, including for image processing!) Starting around 2005: the deep learning (DL) revolution.

With the explosion of computing power, could fit *much* bigger neural nets, with millions or billions of parameters. In parallel, architectures got more imaginative, like:

- Convolutional NNets for images (we will discuss)
- Recurrent NNs
- GANS

These new architectures can address a wide range of problems beyond standard predictive modeling, a few examples:

- Reinforcement learning: learning to play games, the "reward" is only observed at the end
- Auto-encoding: find structure in data (non-linear PCA)
- Transfer learning: Learn networks on problems where you have lots of data, adjust to problems with a little data

They also involve more modern statistical elements, such as:

- Dropout: occasionally zero some of the weights in the network to *get out of local minimum*
- Momentum: Instead of just using the gradient (first derivative) remember previous gradients (sort-of second derivative)

Starting around 2005 we observe amazing accuracy gains in many domains from using DL models, fields dominated by DL today: image recognition, speech recognition, natural language processing,...

## Convolutional neural nets: the original DL revolution for images

Motivating example: CIFAR-10

An image is a  $K \times K \times L$  object (for example  $K = 100, L = 3$  layers of color). For simplicity assume here one layer  $L = 1$ . Denote an image by  $A$ , the  $i, j$  pixel is  $A(i, j)$ . A *local area* in the image is a  $q \times q$  region (say  $q = 3$  or  $q = 5$ ).

**Filters as detectors:** A *filter*  $f$  is a  $q \times q$  mask  $f(k, l)$ ,  $k, l = 1, \dots, q$ . Now  $f$  can be applied to every  $q \times q$  area in  $A$ :

$$(f * A)(i, j) = \sum_{k=1}^q \sum_{l=1}^q f(k, l) \times A(i + k, j + l)$$

Gives a new "image" of size  $(K - q + 1) \times (K - q + 1)$ , where each pixel is a *convolution* of a neighborhood with the filter.

**Filter example: finding a corner or a line** Consider the following two  $3 \times 3$  filters:

$$f_1 = \begin{pmatrix} 0 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \quad f_2 = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \end{pmatrix}$$

$(f_1 * A)$  gives high values where  $A$  has a corner, while  $(f_2 * A)$  gives high values where  $A$  has a vertical line Example: Sobel filter finds edges

Filters can find "features" in the image that can help identify it:

- A house has corners
- A snake has a repeating pattern on its back
- A cat has pointy ears

But we don't know what features will allow us to solve a specific problem:

- Corners, lines, edges, patterns?
- What scale? A small cat has small pointy ears, a big cat has LARGE pointy ears — they are both cats

So, we want the model to "learn" filters of different types and different scales that will help it solve the specific problem

### A CNN architecture:

- We don't decide on filters, we "learn" them: a  $q \times q$  filter has  $q^2$  parameters.
- Each filter is applied to the entire image to create a new "image".
- Images are scaled ("pooling"), then more filters are applied, to find features at different scales
- The learned features are then used in a dense feed-forward network to classify images

Example

Some details:

- Number of parameters of convolutional layers is relatively small: assume input is  $K \times K \times L$  and we have  $C$  convolutions of size  $q \times q \times L$ , then the number of parameters is  $C \times q^2 \times L$  regardless of image size  $K$
- Pooling: decreasing image size from  $K \times K$  usually to  $K/2 \times K/2$  by aggregating every  $2 \times 2$  region to a single pixel
- Most common pooling operator: max pooling (are there pointy ears anywhere in this region?)
- By alternating convolutions and pooling we learn features at different scales
- Important: all parameters (filters, dense layers weights,...) are learned \*together\* using SGD

Example: Simpsons vs South Park