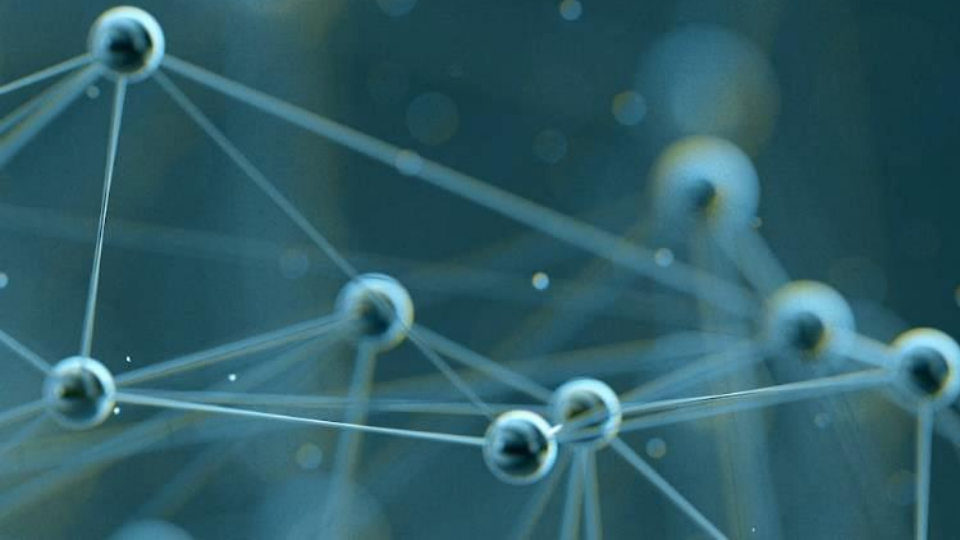


REGULARIZATION TECHNIQUES FOR DEEP LEARNING



REGULARIZING NEURAL NETWORKS

We have several means by which to help “regularize” neural networks – that is, to prevent overfitting

- Regularization penalty in cost function
- Dropout
- Early stopping
- Stochastic / Mini-batch Gradient descent (to some degree)

PENALIZED COST FUNCTION

- One option is to explicitly add a penalty to the loss function for having high weights.
- This is a similar approach to Ridge Regression

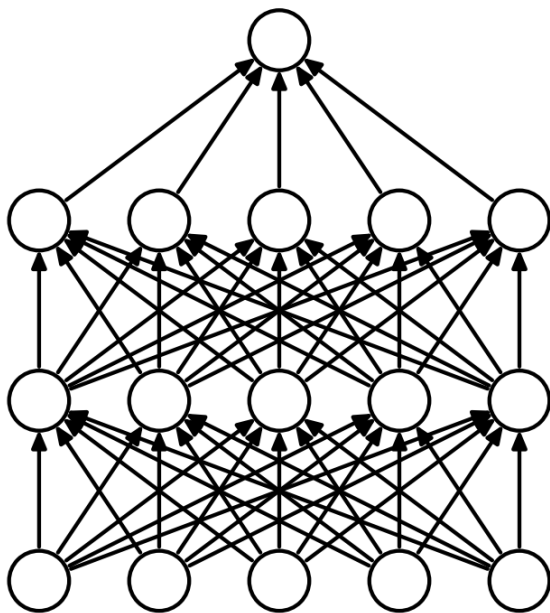
$$J = \frac{1}{2n} \sum_{i=1}^n (\hat{y}_i - y_i)^2 + \lambda \sum_{j=1}^m W_j^2$$

- Can have an analogous expression for Categorical Cross Entropy

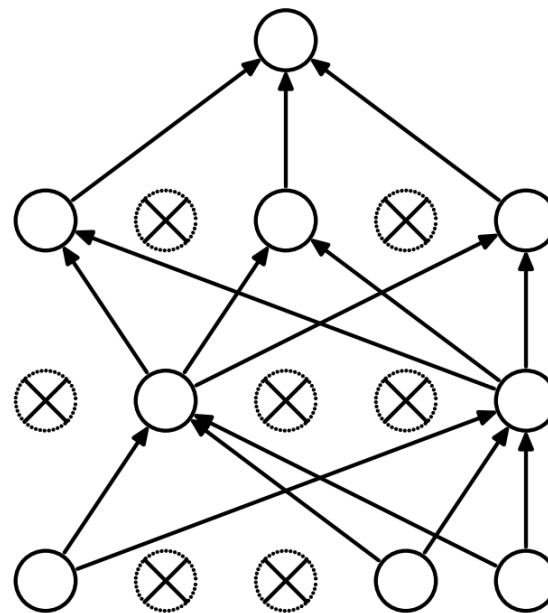
DROPOUT

- Dropout is a mechanism where at each training iteration (batch) we randomly remove a subset of neurons
- This prevents the neural network from relying too much on individual pathways, making it more “robust”
- At test time we “rescale” the weight of the neuron to reflect the percentage of the time it was active

DROPOUT—VISUALIZATION



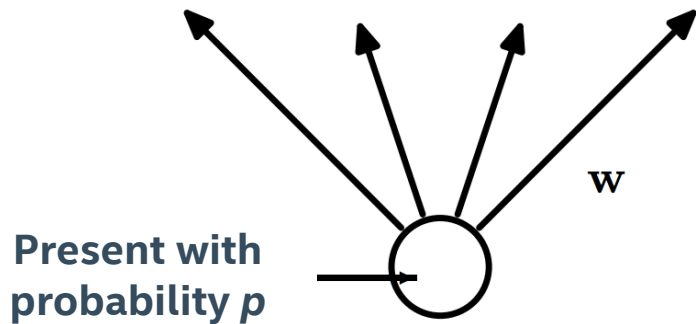
(a) Standard Neural Net



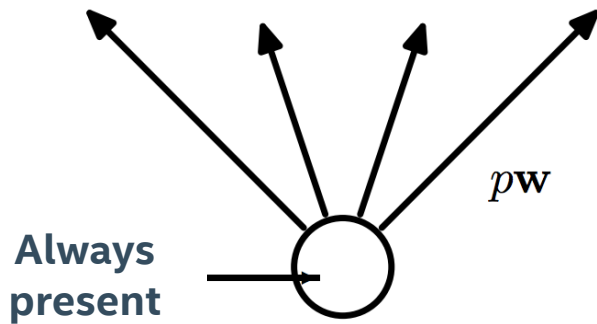
(b) After applying dropout

DROPOUT—VISUALIZATION

If the neuron was present with probability p , at test time we scale the outbound weights by a factor of p .



(a) At training time



(b) At test time

EARLY STOPPING

- Another, more heuristical approach to regularization is early stopping.
- This refers to choosing some rules after which to stop training.
- Example:
 - Check the validation log-loss every 10 epochs.
 - If it is higher than it was last time, stop and use the previous model (i.e. from 10 epochs previous)

OPTIMIZERS

- We have considered approaches to gradient descent which vary the number of data points involved in a step.
- However, they have all used the standard update formula:

$$W := W - \alpha \cdot \nabla J$$

- There are several variants to updating the weights which give better performance in practice.
- These successive “tweaks” each attempt to improve on the previous idea.
- The resulting (often complicated) methods are referred to as “optimizers”.

MOMENTUM

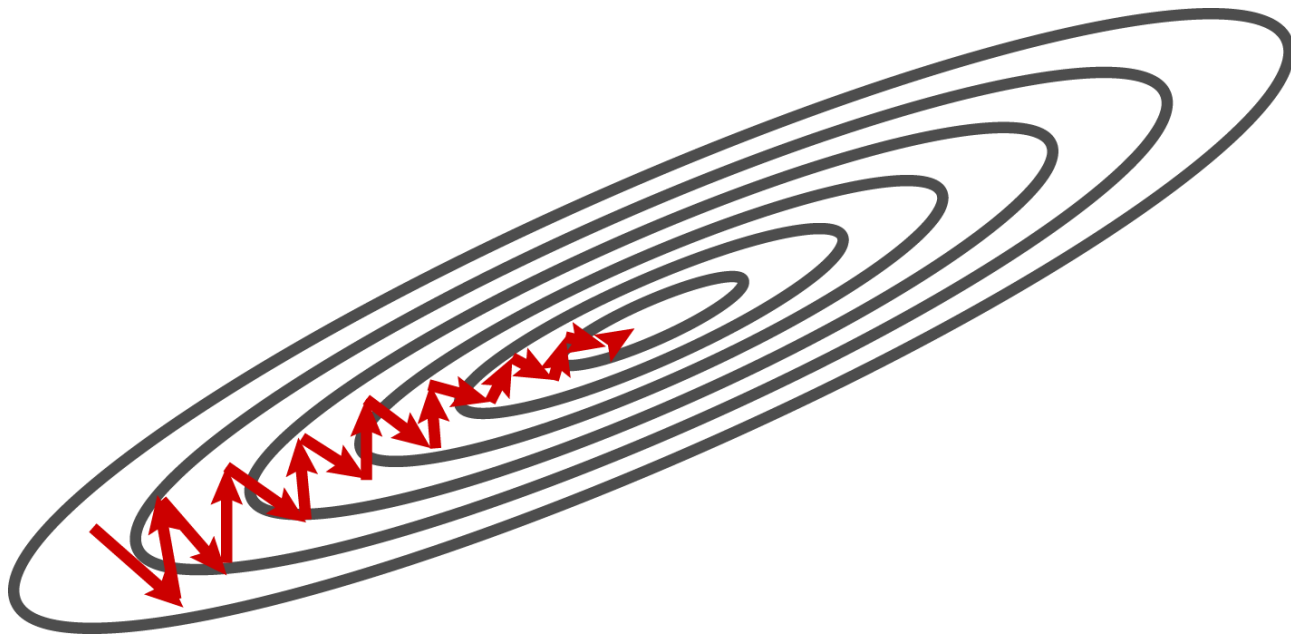
- Idea, only change direction by a little bit each time.
- Keeps a “running average” of the step directions, smoothing out the variation of the individual points.

$$v_t := \eta \cdot v_{t-1} - \alpha \cdot \nabla J$$

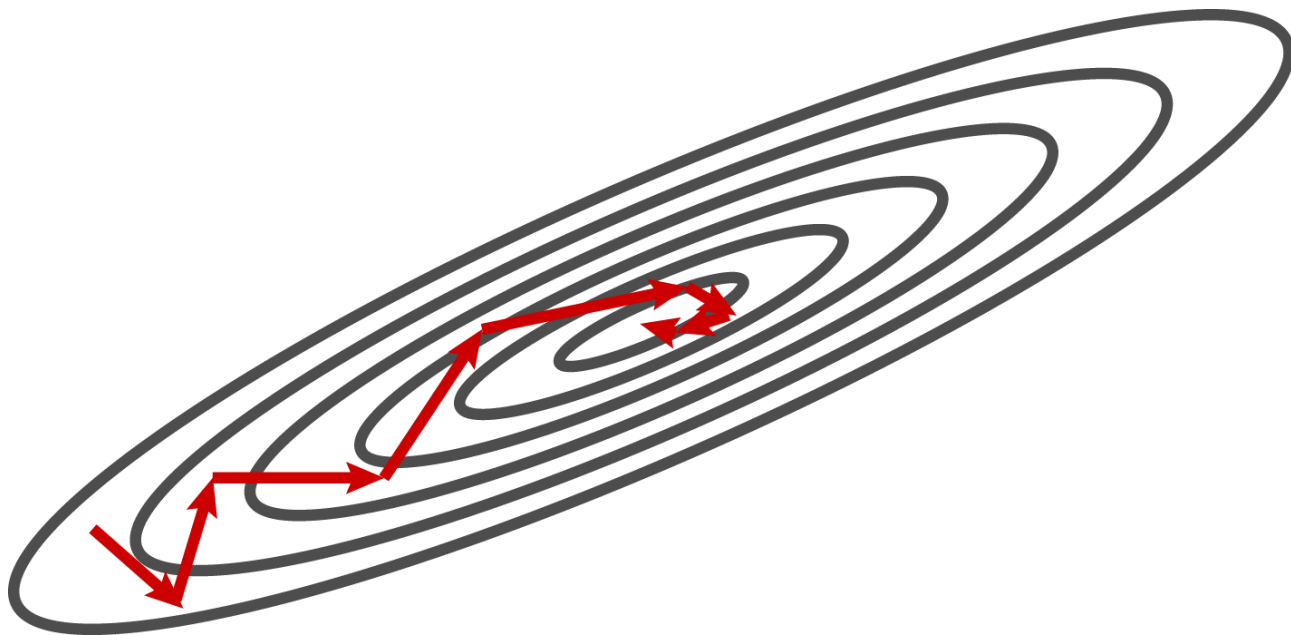
$$W := W - v_t$$

- Here, η is referred to as the “momentum”. It is generally given a value < 1

GRADIENT DESCENT VS MOMENTUM



GRADIENT DESCENT VS **MOMENTUM**



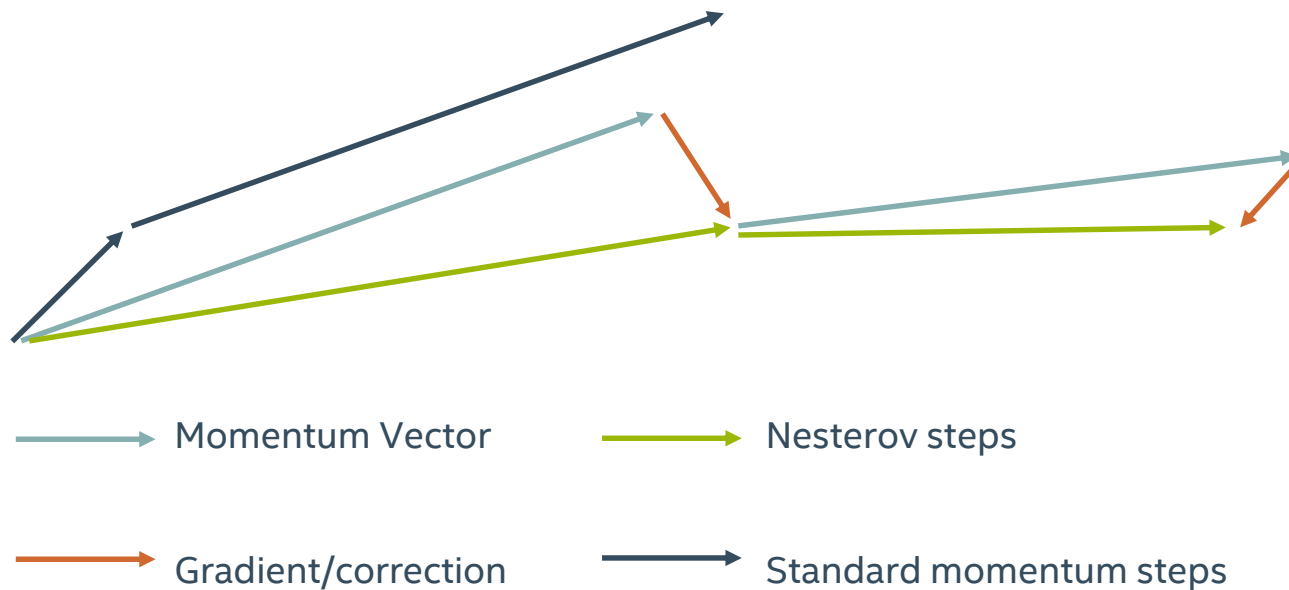
NESTEROV MOMENTUM

- Idea: Control “overshooting” by looking ahead.
- Apply gradient only to the “non-momentum” component.

$$v_t = \eta \cdot v_{t-1} - \alpha \cdot \nabla(J - \eta \cdot v_{t-1})$$

$$W := W - v_t$$

NESTEROV MOMENTUM



ADAGRAD

- Idea: scale the update for each weight separately.
- Update frequently-updated weights less
- Keep running sum of previous updates
- Divide new updates by factor of previous sum

$$W := W - \frac{\eta}{\sqrt{G_t} + \epsilon} \odot \nabla J$$


RMSPROP

- Quite similar to AdaGrad.
- Rather than using the sum of previous gradients, decay older gradients more than more recent ones.
- More adaptive to recent updates


ADAM



Idea: use both first-order and second-order change information and decay both over time.

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) \nabla J$$


$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t}$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) \nabla^2 J$$


$$\hat{v}_t = \frac{v_t}{1 - \beta_1^t}$$


$$W := W - \frac{\eta}{\sqrt{\hat{v}_t} + \epsilon} \odot \hat{m}_t$$

WHICH ONE SHOULD I USE?!

- RMSProp and Adam seem to be quite popular now.
- Difficult to predict in advance which will be best for a particular problem.
- Still an active area of inquiry.

