

More on fitting neural networks

This note is a continuation of our high-level introduction to neural networks.

Local optima

The cost function for neural networks is not unimodal, and so is certainly not convex (a stronger property). It's easy to see why by considering a neural network with two hidden units. Assume we've fitted the network to a (local) optimum of a cost function, so that any small change in parameters will make the network worse. Then we can find another parameter vector that will represent exactly the same function, showing that the optimum is only a local one.

To create the second parameter vector, we simply take all of the parameters associated with hidden unit one, and replace them with the corresponding parameters associated with hidden unit two. Then we take all of the parameters associated with hidden unit two and replace them with the parameters that were associated with hidden unit one. The network is really the same as before, with the hidden units labelled differently, so will have the same cost.

It's a common difficulty with "hidden" or "latent" representations of data, that there are usually many equivalent ways to represent the same model. As machine learning is usually concerned about making predictions, it doesn't matter that the parameters aren't well-specified. But it's worth remembering that the values of individual parameters are often completely arbitrary, and can't be interpreted in isolation.

In practice there are many more local optima than just the ones corresponding to permuting the hidden units. Some of these optima will have better cost than others, and some will make predictions that generalize better than others. When I've fitted small neural networks, I've tried optimizing many times and used the network that cross-validates the best. However, researchers pushing up against available computational resources will find it difficult to optimize a network many times.

One advantage of large neural networks is that fitting more parameters tends to work better(!). The intuition I have is that there are many more ways to set the parameters to get low cost, so it's less hard to find one good setting. Although it's difficult to make rigorous statements on this issue. Understanding the difficulties that are faced in really high-dimensional optimization is an open area of research. (For example, <https://arxiv.org/abs/1412.6544>.)

Regularization by early stopping

We have referred to complex models that generalize poorly as "over-fitted". One idea to avoid "over-fitting" is to fit less! That is, terminate the optimization routine before it has found a local optimum of the cost function. This heuristic idea is often called "early stopping".

The most common way to implement early stopping is to periodically monitor performance on a validation set. If the validation score is the best that we have seen so far, we save a copy of the network's parameters. If the validation score fails to improve upon that cost over some number of future checks (say 20), we terminate the optimization and return the weights we've saved.

David MacKay's textbook mentions early stopping (Section 39.4, p479). This book points out that terminating the optimizer stops the weights from growing too large. Adding a regularization term to the cost function to achieve a similar effect seems more appealing: if we have a well-defined cost function, we're not tied to a particular optimizer, and it's probably easier to analyse what we're doing.

However, I've found it hard to argue with early stopping as a pragmatic, sensible procedure. The heuristic directly checks whether continuing to fit is improving predictions for held-out data, which is what we care about. And we might save a lot of computer time by terminating early. Moreover, we can still use a regularized cost function along with early stopping.

Regularization corrupting the data or model

There are a whole family of methods for regularizing models that involve adding noise to the data or model during training. Like early-stopping, I found this idea unappealing, as it's hard to understand what objective we are fitting, and it makes the models we obtain depend on which optimizer we are using. However, these methods are often effective...

Adding Gaussian noise to the weights of a linear model during gradient training has the same average effect as L2 regularization¹. We can add noise to the parameters while fitting neural networks too. The procedure will have a regularization effect, but one that's harder to understand. In practice adding noise may work better than optimizing a cost function we can define simply.

Other regularization methods randomly replace some of the weights with zeros ("drop-out"²) or features with zeros (such as in "denoising auto-encoders"³). These heuristics prevent the model from fitting delicate combinations of parameters, or fits that depend on careful combinations of features. If used aggressively, "masking noise" makes it hard to fit anything! Often *large* models are needed when using these heuristics.

Further Reading

Most textbooks are long out-of-date when it comes to recent practical wisdom on fitting neural networks and regularization strategies. However, <http://www.deeplearningbook.org/> is still recent, and would be a good starting point.

1. Non-examinable: there's a sketch in these slides: http://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec9.pdf
2. <https://www.cs.toronto.edu/~hinton/absps/JMLRdropout.pdf>
3. <http://machinelearning.org/archive/icml2008/papers/592.pdf>