
Artificial Neural Networks for Engineering Students

Matthew Burns

mdburns@eng.ucsd.edu

At the request of some ECE students, I have started a neural networks reference page. Instead of overwhelming the reader with the many sources and opinions that exist, I have chosen the few references which I see as important to getting started.

1. Neural net application examples:

- (a) Cat faces: Google trains a billion-parameter neural network on GPUs (\$\$\$) and discovers cats on the internet... so cool.
- (b) Deep Mind: Acquired by Google for \$500M+ because they made an unsupervised net that can learn to play Atari games at superhuman levels. Combines deep learning and reinforcement learning.
- (c) Deep Learning for Decision Making and Control: Advanced application of probabilistic neural networks. Combines deep learning with optimal control.

2. Video lecture:

- (a) Deep Learning for Computer Vision (Rob Fergus)
- (b) Slides

3. Reading:

- (a) Bishop [1]: this is the main theoretical reference for neural networks. It even has a chapter on Bayesian interpretations at the end, tying neural networks to probabilistic graphical models.
- (b) Efficient BackProp [3]: Also referred to as "Tricks of the trade" because it was included in a book with that title.
- (c) Convolutional Neural Networks [2]: Trained a deep, very high dimensional ($\mathbb{R}^{60,000,000}$) convolutional network on ImageNet dataset. This network architecture learns the optimal FIR filters which produce features which are good for separating the different image classes. What is interesting is that the neural network features get better classification performance than human designed features.

4. Tutorial projects:

- (a) UFLDL Tutorial
- (b) These are intense project intended to be completed in teams of three people.
- (c) If you successfully complete these projects, you will know deterministic neural networks well.
- (d) This project sequence starts you off at ECE 174; if you took that class, then you are well prepared to start. You can think of these projects as a template for an independent study, taken with the other reading I mentioned, esp. Bishop.

5. Open source projects:

- (a) Caffe: a cutting edge image classification system. You could hand design your own image features for several years to get the best possible generalized performance. Or you can rip off the features from this project and automatically get state of the art performance on ImageNet.

- (b) NuPIC: This allows for spatiotemporal pattern classification. Some have criticized this system because it has not produced impressive results. On the other hand, there is quite a following, so worth mentioning. video

6. Datasets:

- (a) ImageNet
- (b) MNIST
- (c) kaggle
- (d) Extensive list of computer vision datasets

References

- [1] Christopher M Bishop. *Neural networks for pattern recognition*. Oxford university press, 1995.
- [2] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. “Imagenet classification with deep convolutional neural networks”. In: *Advances in neural information processing systems*. 2012, pp. 1097–1105.
- [3] Yann A LeCun et al. “Efficient backprop”. In: *Neural networks: Tricks of the trade*. Springer, 2012, pp. 9–48.