THE DIRE CONSEQUENCES OF NEURAL NETWORKS IN SCIENCE AND ENGINEERING

Matthew Burns

mdburns@eng.ucsd.edu

Instead of overwhelming the reader with the many sources and opinions that exist, I have chosen to demonstrate in narrative how obscenely dangerous these technologies could be. If you have the opportunity, THINK ABOUT HOW THESE COULD BE USED FOR EVIL, AS WELL AS GOOD. PLEASE.

1. **Neural net applications**:

    (a) *Cat faces*: Google trains a billion-parameter neural network on GPUs ($$$) and discovers cats on youtube

    (b) *Deep Mind*: Acquired by Google for $500M+ because they made a deep 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*: A PhD thesis out of Berkeley: Combines deep learning with optimal control

    (d) *Inceptionism - Going Deeper into Neural Networks*: beautiful images in this article; no comment on the science part of it

2. **Video lecture**:

    (a) *Deep Learning for Computer Vision*

    (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 republished in a book which bears the 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:** *UFLDL*

    (a) These are programming-intensive projects, intended to be completed in teams of 2-3.

    (b) If you successfully finish, you will know deterministic neural networks well.

    (c) Sequence starts you off at ECE 174; if you took that class, then you are well prepared. You can think of these projects as a template for an independent study, taken with the other reading I mentioned, especially Bishop.

5. **Open source projects:**

    (a) *Caffe*: cutting edge, pre-trained convolutional neural net for image classification

    (b) *matconvnet*: Matlab convolutional neural net toolbox

6. **Datasets:**

    (a) *ImageNet*

    (b) *MNIST*

    (c) *kaggle*

    (d) *Extensive list of computer vision datasets*

7. **Current Events:**

    (a) *Blowback*

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.
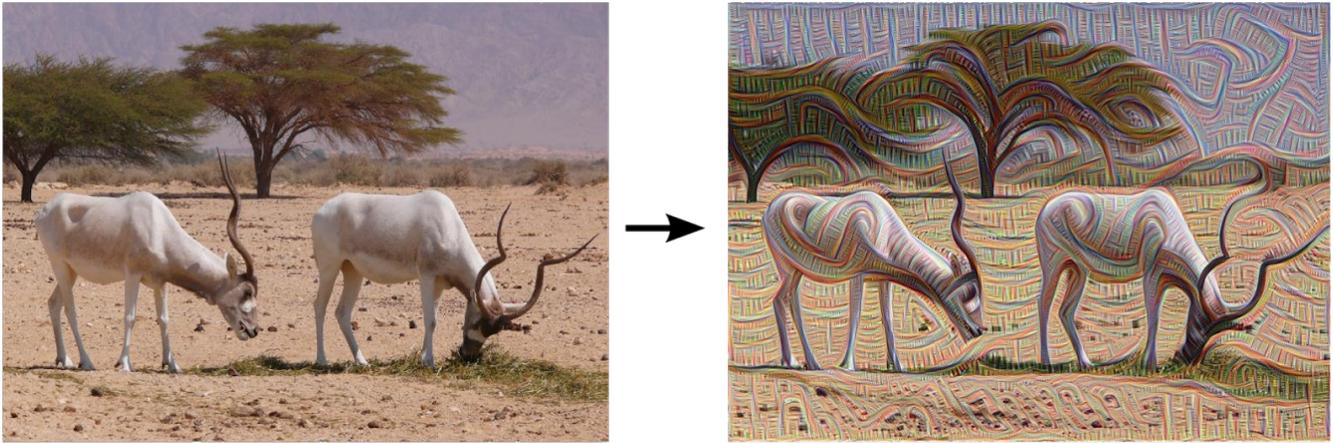
Figure 1: From *Google research blog*: the left image is the original. To the right, we see how the network represents the image. The specific "brush strokes" in the reconstructed image are selected from a subset of the non-linear eigenvectors which the network learned from training data.

X

[3]   Yann A LeCun et al. "Efficient backprop". In: *Neural networks: Tricks of the trade*. Springer, 2012, pp. 9–48.