
Shallow Dreams

The Unseen World of the Interior of a Neural Net

Overview

In the world of multi-layered neural nets, it is clear that machines can be “taught” a wide range of things and by extension act on what they’ve learned. But, what is it that they are actually “learning”?

Google’s [Deep Dream](#) project offered a very striking visualization of this in asking a neural net to see things when they weren’t actually there. But even this is just an exterior look at the machine’s world. But what if we could look inside the net? Could we then use that information to create even more powerful neural nets? More nimble or creative ones?

Goals

I propose to create a visualization of the internal workings of a neural net, and then provide a framework to easily adjust the hyper-parameters of the net and see through the visualization how this alters the machine’s success at learning tasks. The goal of this would be two-fold:

- To more quickly adapt the hyper-parameters of the net via human intuition
- To give insight into how neural nets are constructed to guide future projects

Project Outline

- Build an architecture for creating easily extensible neural nets
- Develop at least 3 neural nets (linear, recurrent, and convolutional)
- Hosted on an AWS instance
- Develop a web interface for live user input and net feedback
- Explore visualization options for the interior elements of the net

Presentation

Results of this project would be developed into a talk for data science minded groups or alternatively a research paper to be posted to [arXiv](#).

Contact

Cole Howard
uglyboxer@gmail.com
(503) 799-9706
