Deep Lemniscate

Alexander Madurowicz
Stanford University
CS230 Deep Learning
amaduro@stanford.edu

March 22, 2018

Abstract

The Gemini Planet Imager (GPI) Exoplanet Survey is a mission to directly image planets orbiting around other stars by employing adaptive optics and coronagraphy as well as clever post-processing techniques. Improvements to the sensitivity of the survey could be achieved by creating a model to understand a particular form of noise aberration known as the wind-butterfly. A deep neural network with a cross entropy logistic loss function has been demonstrated as a classifying tool with greater than or equal to human level accuracy in categorizing data from the survey, which will be effective for follow up studies. Future ideas for predictive control are briefly discussed.

Introduction

The Gemini Planet Imager is an instrument installed on the Gemini South Telescope in Cerro Pachon, Chile, designed to search for thermal emission from young hot extrasolar planets at wide angular separation. The GPI Exoplanet Survey is well under way, with many successful detections, but contrast remains limited by atmospheric aberrations and imperfections in the AO system. Notably, the presence of the so-called 'wind butterfly' (so-called because of their figure-8 or lemniscate shape) scatters significant light into the coronagraphic dark hole, causing the residual PSF to break azimuthal symmetry in the image plane, and reducing the final contrast ratio. This PSF pattern is consistent with wavefront errors from servo lag. A two to three microsecond delay in positioning of the deformable mirror in response to the wavefront sensor causes an effective displacement of the atmospheric turbulence relative to the applied correction. This effect produces a phase pattern on the detector whose Fourier transform has preferential direction. This butterfly shaped noise aberration is difficult for planet detecting algorithms to handle, because many of them assume that the background PSF is azimuthally symmetric, although the butterfly is not. Classification of the survey's dataset into categories which separate those images which contain this aberration from those who do not will be useful for doing various follow-up studies which could potentially eliminate the butterfly is post processing with the overall goal of improving survey sensitivity for detecting more planets.

Dataset

The Dataset contains 20,561 Image cubes, each of which is a (37,281,281) datacube stored in a .fits files, which contains thirty seven different spectral images of the target taken simulatenously on an Integral Field Spectrometer (IFS) over the course of a thirty second exposure. Exposure sequences are blocked into hour long chunks on particular targets for use with Angular Differential Imaging (ADI) post processing techniques. The data cubes contain 32-bit floating point numbers, as well as NaNs which arise from post-processing where no measurements were made. These NaNs are converted to 0 to maintain analyticity during matrix multiplication for the purposes of this project. Just over forty percent of the Dataset was labeled by hand over eight hours of playing a simple labeling game, with an estimated self-consistency for human level performance of 83.8% from comparisons of duplicate labels on a particular subset. Then, the labeled dataset was separated in Training, Development, and Testing sets with a ratio of 60/20/20. The figure below shows an example of each class of labeled data, the difference between them is the presence of the lemniscate shaped noise aberration of which we would like to identify.

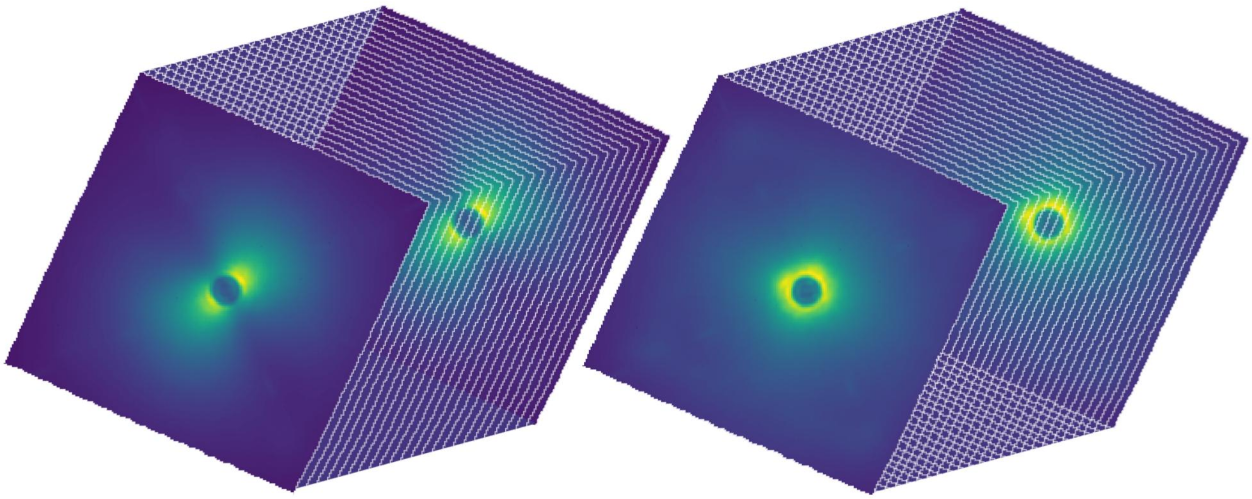

Figure 1: Two image cubes from the dataset, A Butterfly and a Not Butterfly, showing the two categories that we wish to predict with a binary classification algorithm. To display these cube faces we radical summed the frequency axis. Note the presence of the lemniscate shaped aberration in the left PSF.

Network Architecture

The network architecture employed is a rather simple model, consisting of four hidden layers of decreasing size, with Rectified Linear Unit (ReLU) activations and a final sigmoid output for binary classification. The total number of trainable parameters is 73039347. A cross entropy loss function is computed to optimize the parameters during backward propagation.

$$\mathcal{L}(y, \hat{y}) = \frac{1}{m} \sum_{i=1}^{m} y^{(i)} \log(\hat{y}^{(i)}) - (1 - y^{(i)}) \log(1 - \hat{y}^{(i)}) \tag{1}$$

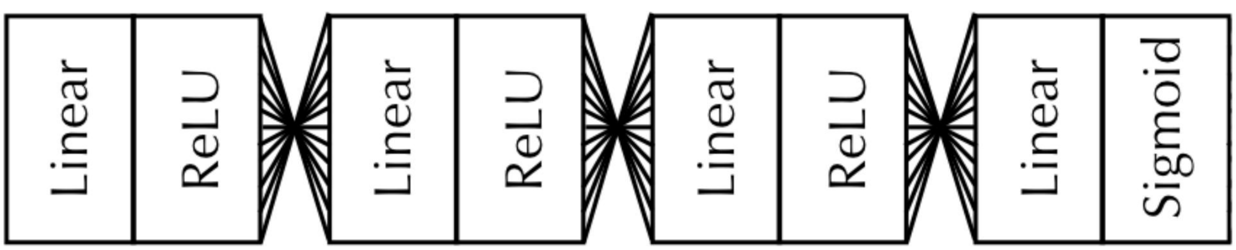

Figure 2: Demonstration of a simple network architecture with four fully connected layers.

Performance

The network was trained over the course of multiple weeks on an Nvidia GTX 980 GPU, for a total of three thousand epochs, with a decaying learning rate of 10^{-3}, 10^{-4}, and 5×10^{-5}, transitioning at epochs 500 and 1000, respectively. The optimization algorithm used was an Adam Optimizer.

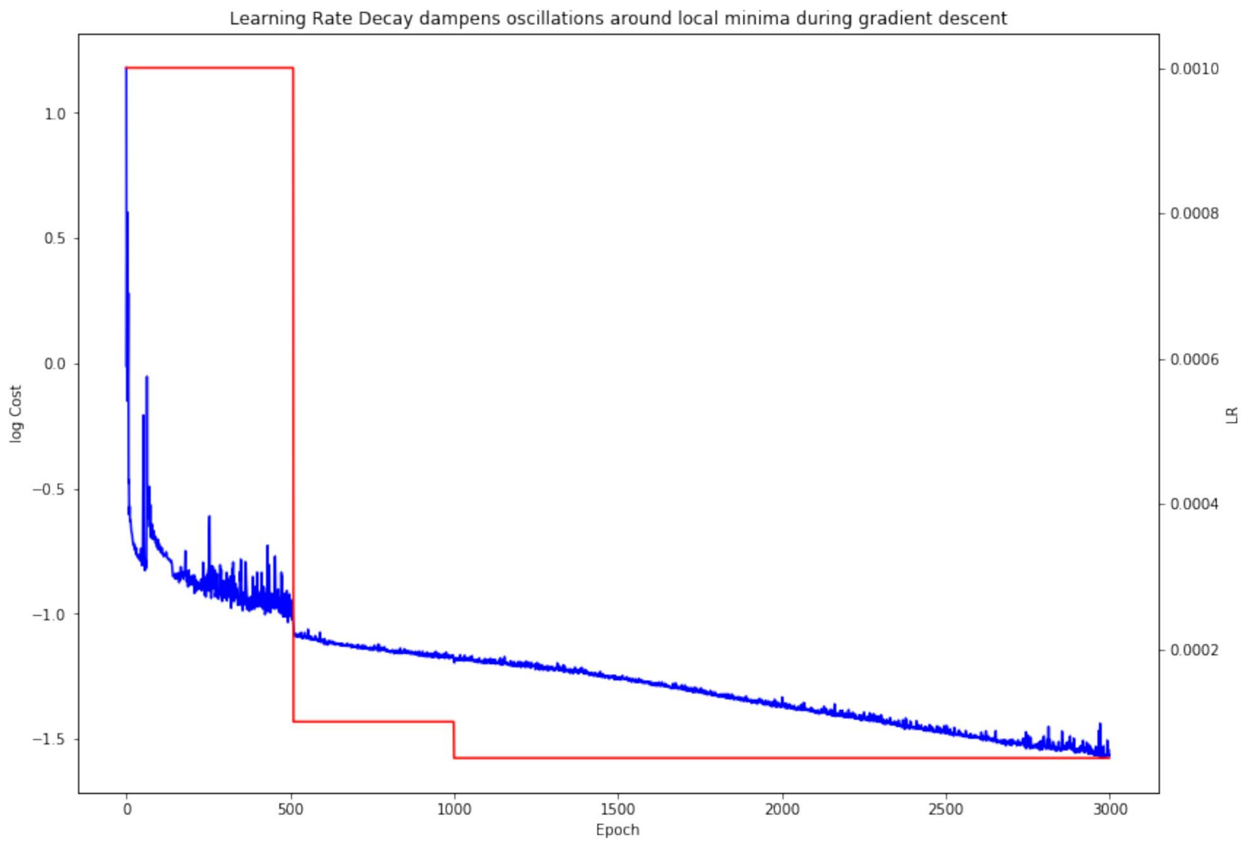

Figure 3: Three thousand epochs of training demonstrating the slow decrease in the cost function. Oscillations around local minima were damped out by diminishing the learning rate parameter in the optimization algorithm.

By the end of the first week of training, the simple network was already performing at the level of our estimate for human-self consistency in the data labeling process. While this in practice should be an upper theoretical bound on the performance of the algorithm, it continued to improve over the following two weeks, albeit gaining only about a percent improvement. This is most likely cause by inaccuracy in the estimation of the human performance, as the success of the algorithm should be bounded above by the noise in the labels it was given.

However, this is still a drastic improvement over the previous optimal technique, which is labeled in Figure 4 as 'Optimal Frac' in black. This metric was the previous classification scheme, where one would take the standard deviation in an annulus centered about the PSF divided by the mean of the annulus in a radial fashion, to find images with the highest deviations from azimuthal symmetry. Maximizing over the cutoff parameters, the best this algorithm could achieve was around 78.7%, around seven percent lower than the performance of the Neural Network implementation.

The final trained model will be useful for is speed and near-super-human levels of accuracy, to automatically classify new images directly from the telescope, separating the vast datasets into butterfly and non-butterfly categories, assisting in developing atmospheric models to help eliminate the noise.

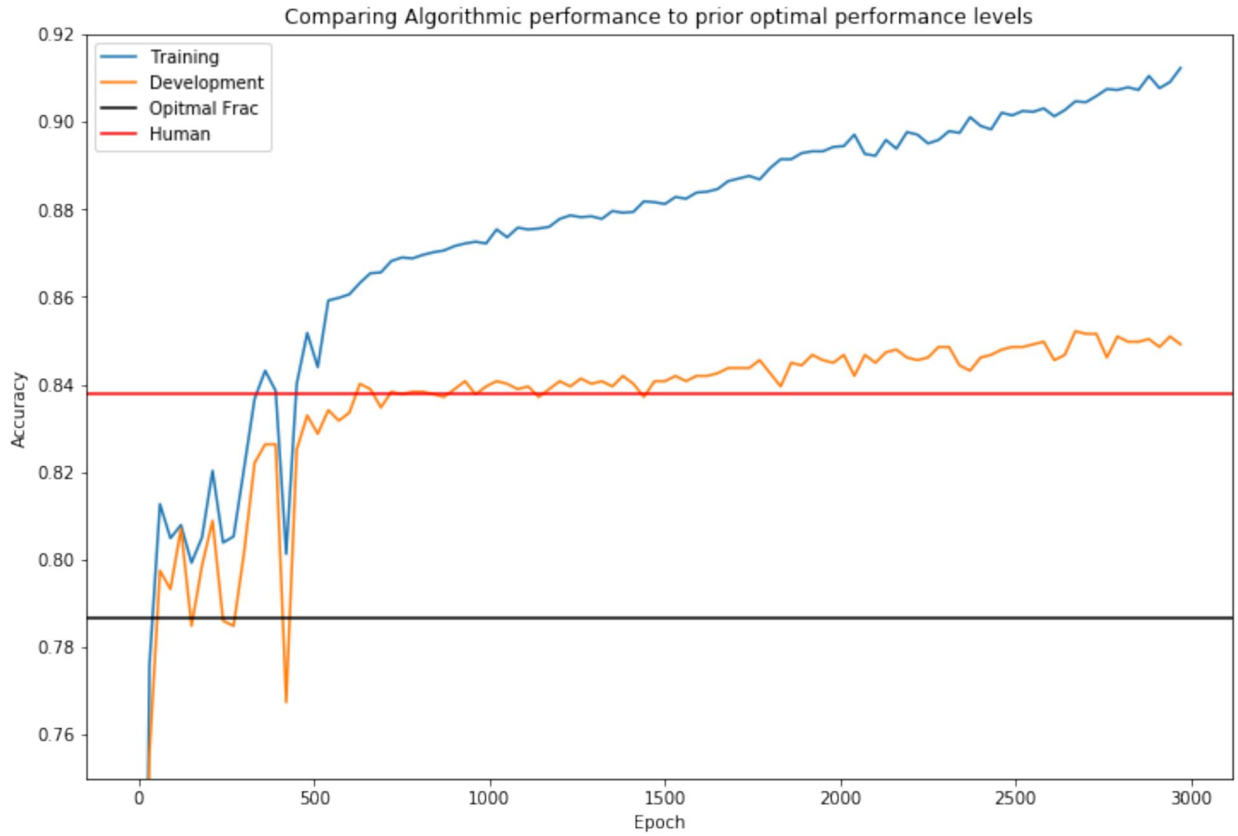

Figure 4: The increasing performance of the algorithm in classifying data in the training and development sets over three thousand epochs of training is visible. Over fitting to the training set near the end becomes apparent, and the performance is bounded by the human level accuracy in labeling the inputs. The network far exceeds the baseline performance expected from previously developed metrics labeled 'Optimal Frac' at 78.7%

Future Work

There are a few suggestions that I made previously about ideas to possibly improve this network, but as I progressed through the project I determined that they would not be worthwhile to pursue, notably, a convolutional implementation is likely better at classifying images. However, since it was determined that the noise given as the labels was a theoretical upper bound on the performance, it does not make sense to try and build a more complicated network structure, when the simplest possible one already can achieve that level of optimal performance. The main takeaway lesson that I learned from this project is that a supervised learning project can only do as well as the data given to it, and that is why for future work I am suggesting a project which is much different from the task completed in this project.

For the same telescope and instrument, there exists a dataset of Adaptive Optics Telemetry measurements, roughly 4000 data cubes of (9,9, 20000) shaped measurements which can be thought of as a twenty second video of the sky, taken at millisecond resolution, with a 9x9 wavefront sensor. The wavefront sensor measures the relative phase of the incoming radiation, which the adaptive optics system uses to adjust the deformable mirrors to correct for turbulence on the sky. The butterfly is thought to arise from a few millisecond delay in this system, so I propose a neural network based predictive control system, which could estimate the phase a few milliseconds ahead of what is currently measured, using an RNN based framework where the output of each node is the actual value but just a few milliseconds ahead. Such a network would not be confined by noise in the human labeled inputs, because the ground truth value is known as it is the measurement, and so such a network could be unbounded in its performance, and potentially solve an large challenge in adaptive optics.

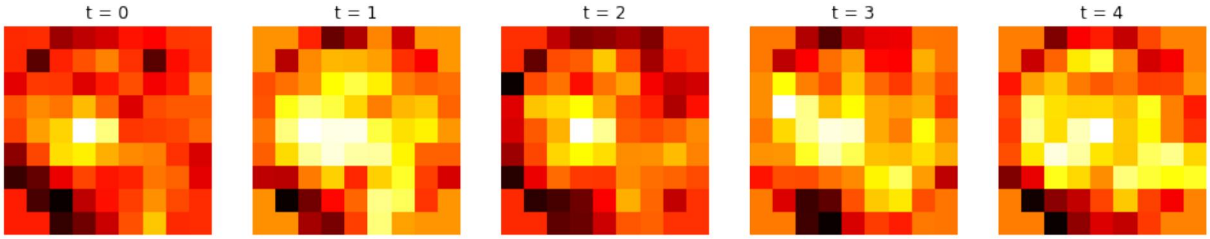

Figure 5: Five sample images from the wavefront sensor measurement dataset, showing a 9x9 slice each space one thousand milliseconds apart. This dataset could be used for a predictive RNN based approach to eliminating errors from the delay of the AO system.

Conclusion

In Summary, we were able to demonstrate that a simple deep learning model was able to match (and perhaps barely exceed) human level performance in the categorization of astronomical data in the GPI Exoplanet Survey Dataset, with an unbiased estimate of performance on the Test Set of 84.9%, which is an improvement over our estimate of human accuracy of 83.8%. Previous sorting algorithms designed with application-specific knowledge were even less superior, with performance accuracies of 73.7%. Use of this categorization algorithm will be useful for automatically classifying new image cubes from the telescope, which will in turn assist in developing reduction models and improve planet detection sensitivity. A proposed extension to this project is briefly discussed, with an idea to generate ground truth labels in an RNN framework that would be unbounded by human error in labeling.

References

[1] F. J. Rigaut, J.P. Veran, and O. Lai. Analytical model for Shack Hartmann-based adaptive optics systems. The Proceedings of SPIE 3353. 1998.

[2] O. Guyon. Limits of Adaptive Optics for High-Contrast Imaging. The Astrophysical Journal 629:592-614. August 2005.