Deep Lemniscate

Alexander Madurowicz and the GPI Exoplanet Survey Team

Stanford University, Kavli Institute for Particle Astrophysics and Cosmology

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. 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 with the overall goal of improving survey sensitivity.

Dataset

The Dataset contains 20,561 Images, each of which is a (37,281,281) datacube of .fits files, which contains thirty seven different spectral images of the target taken over the course of a thirty second exposure. Exposure sequences are blocked into hour long chunks of particular targets. The data cubes contain 32-bit floating point numbers, as well as NaNs which arise from post-processing where no measurements were made. NaNs are converted to 0 to maintain analyticity during matrix multiplication. Just over forty percent of the Dataset was labeled by hand over eight hours of playing a simple labeling game script, with an estimated self-consistency for human level performance of 83% from comparisons of duplicate labels. 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.

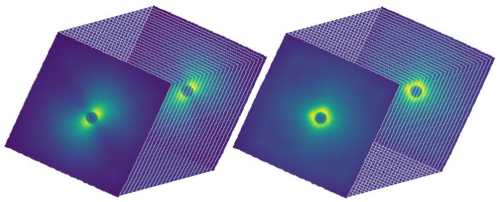

Figure: Two image cubes from the dataset, A Butterfly and a Not Butterfly, showing the two categories that we wish to predict with a binary classification algorithm. To display these cube faces we radical summed the frequency axis. Note the presence of the lemniscate shaped aberration in the left PSF.

Network Architecture

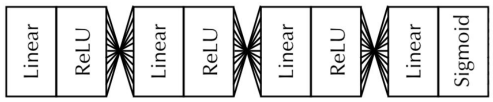

Figure: Demonstration of a simple network architecture with four fully connected layers.

The network architecture is a very 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.

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

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.

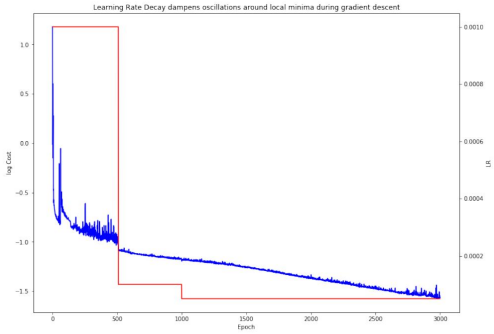

Figure: Three thousand epochs of training demonstrating the slow decrease in the cost function. Oscillations around local minima were damped out by diminishing the learning rate parameter in the optimization algorithm.

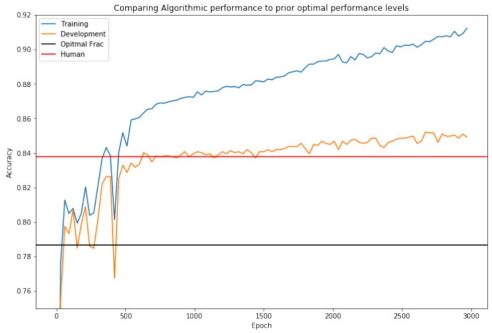

Figure: The increasing performance of the algorithm in classifying data in the training and development sets over three thousand epochs of training is visible. Over fitting to the training set near the end becomes apparent, and the performance is bounded by the human level accuracy in labeling the inputs.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

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.

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.

Exoplanet Survey

Stanford University