Advancing the Search for Dark Matter with Deep Convolutional Neural Networks

Sebastian Wagner-Carena, Ben Sorscher
{swagnerc, bsorscher}@stanford.edu
Deptartments of Physics and Applied Physics, Stanford University

Overview

We use deep Bayesian convolutional neural networks to achieve state-of-the-art predictions on the salient features of gravitational lenses, along with a full covariance matrix of uncertainties --providing a powerful tool for next-generation dark matter surveys.

Background

One of the profound predictions of Einstein's general theory of relativity, gravitational lensing –the bending of light's path by sufficiently massive objects– has become a powerful tool in modern observational cosmology. Images of gravitational lenses have distinctive signatures, including a prominent ring-like structure (see Fig. 1). These signatures enable us to infer the dominant features of distant galaxies, probe the presence of dark matter, and obtain precise predictions of the Hubble constant.

Unfortunately, the calculations required to infer these parameters are computationally costly. Before next-generation sky surveys such as LSST, Euclid, and WFIRST come online in the next 2-3 years, we need efficient, accurate tools to predict the parameters of gravitational lenses.

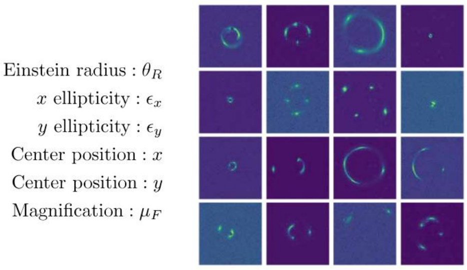

Figure 1: Images of gravitational lenses and their associated parameters

Data

We have gathered a training set of 100,000 high-quality, labeled images of gravitational lenses (Fig.1), generated using background galaxy images from GalaxyZoo [1] and GREAT3 [7], We augment this data by (1) Adding Gaussian noise, (2) Adding masked pixels, (3) Applying a point spread function characterizing the beam used to capture the image, (4) Adding Poisson shot noise associated with the intrinsic statistics of measuring a finite number of photons, (5) Translating the center of the lens. The resulting images are 192x192 grayscale normalized images.

Methods

We train a deep convolutional neural network to predict the six parameters specifying the singular ellipsoid density profile of a gravitational lens. The network architecture is as follows:

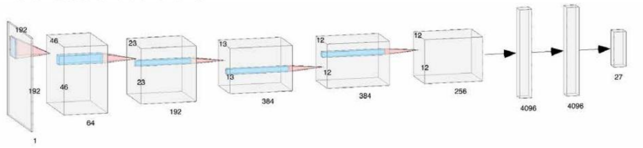

Figure 2: Model architecture

In physics, obaining uncertainties on our predictions is essential. As such, we utilize recent techniques in Bayesian neural networks [1] to estimate both aleatoric and epistemic uncertainties by introducing a variational distribution over our network weights

$$q(w_i) = w_i \text{Bernoulli}(p)$$

$$\mathcal{L}(y_n, \hat{y}_n(x_n, w), \Sigma_n(x_n, w)) = \frac{1}{2}(y_n - \hat{y}_n(x_n, w))^T \Sigma_n^{-1}(x_n, w)(y_n - \hat{y}_n(x_n, w)) + \frac{1}{2}\log|2\pi\Sigma_n(x_n, w)| + \lambda||w||^2$$

We then approximate this variational distribution by implementing the above loss function, and tuning the dropout rate. This enables us to predict the full covariance matrix of parameter values.

Results and Discussion

After training our model for 40 epochs, we obtain the following errors on these parameters, evaluated on a held-out test set of 15,000 images:

	θ_R	ϵ_x	ϵ_y	x	y	μ_F
State of the Art	0.03	0.04	0.05	0.06	0.06	-
Ours	0.004	0.005	0.01	0.004	0.005	4.7

We hope that with further training and hyperparameter tuning we will approach or improve upon the state of the art [4].

We first assess the validity of the diagonal elements of our covariance matrix, corresponding to the standard deviations on each parameter estimate. Confidence intervals are shown in figure 3.

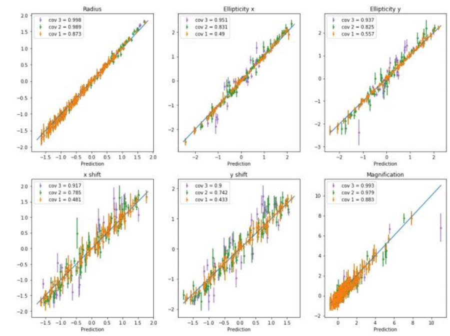

Figure 3: Predicted parameter values and corresponding predicted uncertaintines. Orange represents points within one standard deviation of true value, green within two, and purple within three.

To assess the validity of the full covariance matrix, we generalize our confidence intervals to the full six-dimensional gaussian and compute a chi-squared statistic (see Figure 4)

$$(x-\mu)^T \Sigma^{-1}(x-\mu) \sim \chi_p^2$$

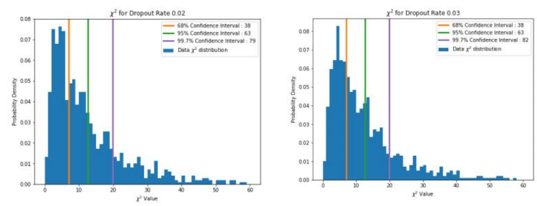

Figure 4: Histogram of chi-squared values of our model's prediction along with bars defining where the 68%, 95%, and 99.7% confidence intervals should lie.

While our model does seem to be capturing the uncertainties, it is generally overconfident in its predictions leading to large chi-squared values.

We examined a few examples of lenses were our model has either a very high or very low level of certainty. The images it is uncertain on are those where the lensed image is small and there are two or fewer copies of the original image - exactly the types of images where it is difficult to define an ellipse. The images it has high certainty have multiple extended copies of the source that trace out the boundary of the ellipse nicely.

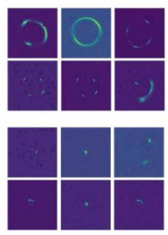

Figure 6: Representative lenses with low (top) and high (bottom) predicted uncertainty.

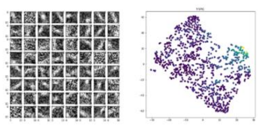

Figure 7: (Left) First convolutional layer weights. (Right) t-SNE evaluated on layer 7 weights (color given by μ_F)

Finally, we visualized the weights of the first convolutional layer, and performed t-SNE on the representation of the seventh network layer. The network seems to focus on elliptical shapes, which is what we would expect.

Future directions

- Explore concrete dropout [5] as a more flexible alternative to standard dropout for EM approximation of the distribution on the weights
- Alter the base architecture to see if that allows for better model performance

1. Gal, Y., & Ghahramani, Z. (2015). Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning, 48. https://doi.org/10.1109/TKDE.2015.2507132
2. Krizhevsky, A., Sutskever, I., & Hinton, G. E. ImageNet Classification with Deep Convolutional Neural Networks. Retrieved from http://code.google.com/p/cuda-convnet/
3. Radford, A., Metz, L., & Chintala, S. (n.d.). Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. https://arxiv.org/pdf/1511.06434.pdf
4. Hezaveh, Y. D., Levasseur, L. P., & Marshall, P. J. (2017). Fast automated analysis of strong gravitational lenses with convolutional neural networks. Nature, 548(7669), 555–557.
5. Gal, Y., Hron, J., & Kendall, A. (2017). Concrete Dropout.