CS230

Exoplanet Candidate Stellar Companion Detection using a Mask R-CNN

Simon Engler
Stanford School of Engineering
Stanford University
stengler@stanford.edu

Abstract

This project attempts a novel application using a Mask Region Convolutional Neural Networks (Mask R-CNN) to adaptive optics observational data from of exoplanet candidates from the Kepler telescope.. Anomalous light curve behavior observed create inflated exoplanet radii measurements and false positive exoplanet detections. The high volume of stellar observations cannot possibly be processed manually and require intelligent software to conduct the task. Using Mask R-CNN, the code identifies stars with a stellar companion and corrects the Asterosemic data along with exoplanet radii measurements to exclude false positives A single GPU was used processing a single image at a time with the CCD camera in greyscale. The learning momentum is set at 0.9 with a learning rate of 0.001. The mini-mask shape is 56x56 searching for a star. Training is conducted on 903 stellar observations, with 38 development images, and 346 validation images. Results show proficient detection of starts at 98% success level.

1 Introduction

This project will be a novel application of deep learning to adaptive optics observational data of exoplanet candidates from the Kepler telescope. Using Mask Region Convolutional Neural Networks (Mask R-CNN) the developed code will identify stars with a stellar companion and correct the Asteroseismic data along with the exoplanet radii measurements to correct the data, or to exclude false positives.The Kepler space telescope used Asteroseismic observations for over 15K stars. Anomalous light curve behavior was detected in 75% ad 8% of dwarf and giant stars respectively. These anomalies consisted of suppressed or no oscillations in Asteroseismic data, which has the potential to create inflated exoplanet radii measurements and false positive exoplanet detection.

In observations conducted by Zeigler et.al (2015) over 10K high resolution adaptive optics observations of exoplanet candidates were collected. Processing this data manually, 88 stars showed nearby companions stars. Through corrections in light curve oscillations due to the companion star the correct planetary radius estimates are inflated by a factor of 1.54 times. Through this correction, 35 rocky exoplanet candidates proved to be false positives. It is believed that 16% of Earth-sized exoplanets are detected around stars with stellar companions. The high volume of stellar observations cannot possibly be processed manually and require intelligent software to conduct the task.

CS230: Deep Learning, Winter 2020, Stanford University, CA. (LateX template borrowed from NIPS 2017.)

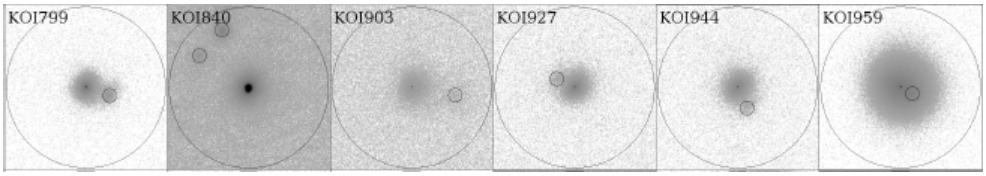

Mask R-CNN combines semantic segmentation which assigns a class to each pixel in the image, and object detection which create object bounding boxes. Then instance segmentation is applied to each pixel that combines the bounding boxes and classes to create a mask to place over the object. The Mask R-CNN is built upon the Fast R-CNN, which consists of a convolutional neural network, a Region Proposed Network, and a Classifier. The Mask R-CNN will use the output to calculate the probability of an object within the bounding box at each pixel. [1]

The observational data is provided by Dr. Barnec of the University of Hawaii Astronomy department. Over 1K observations of stellar images are provided for training with their respective solutions. The unprocessed observations consist of over 10K images. Fortunately, this data is provided already preprocessed into individual images for each observation. Plotting the raw data produces the following images. Figure 1 shows the top view of the starlight captured from the telescope with 100 raw images stacked on top of each other. Figure 2 is a 3D interpolated plot of the intensity of the starlight on the CCD image which displays the noise within the data, plus the Airy disk behavior of the light hitting the CCD surface. Both images clearly show the presence of a secondary stellar object near the main high intensity star.

Detection is done when a companion star is closer than the resolution of the Kepler telescope. Detection is simple given there is enough separation between the host star and the stellar companion. For example, Image K0I799 shows the image of a host star with a fully separated companion. However, given observations as seen in Image KOI944 and KO1959, the companion star is well within the light flux of the host star making its detection difficult. Additional complexities are the presence of multiple stellar companions as seen in image KOI840.

2 Related work

There is a great deal of work in adaptive optics having machine learning automatically adjust optical mirrors using laser light. [2] However, this specific application looking for stellar exoplanet candidates with the intention of creating more accurate light flux profiles is quite specialized and novel.

3 Dataset and Features

Scientific astronomical images are most commonly taken on CCD cameras. The design of these cameras a CCD is essentially an array of light sensitive elements. These elements are arranged in a grid, and each is referred to as a pixel. Data collected in this project has 2048×2048 pixels — just over 4 million in total. The pixels are typically 15 to 25 microns in size. A photon hitting a pixel knocks loose an electron, and hence deposits a charge on the pixel. The charge on each pixel is thus a measure of the number of photons which struck it. Reducing the noise in CCD images is critical to producing clear, high dynamic range images for good scientific measurement. [1]

Light from a uniformly illuminated circular aperture will exhibit an Airy diffraction pattern far away from the aperture due far-field diffraction. The resulting image is a convolution of the ideal image with the Airy diffraction pattern. The intensity of the diffraction pattern of a circular aperture (the Airy pattern) is given by the squared modulus of the Fourier transform of the circular aperture as seen in Figure 3 [3]

To check the accuracy of the simulated light, a single star is simulated as seen in Figure **??** and Figure **??**. Starlight is simulated to fall on a 800x800 CCD chip with added Poisson and Gaussian Noise. This intensity of the star on the image simulates a bright star of magnitude 5.0.

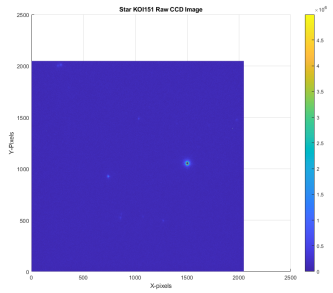

Figure 1: Image of raw observational data showing two stars on CCD camera

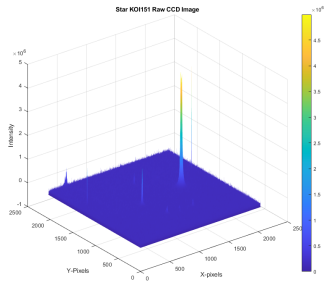

Figure 2: 3D plot of the raw stellar image showing the insensity and noise levels

$$I(\theta) = I_0 \left(\frac{2J_1(ka \sin\theta)}{ka \sin\theta} \right)^2 = I_0 \left(\frac{2J_1(x)}{x} \right)^2$$

Figure 3: The intensity of the diffraction pattern of a circular aperture (the Airy pattern) is given by the squared modulus of the Fourier transform of the circular aperture.

4 Methods

Using open source software VGG Image Annotator provided by [4] and [5] classification annotations of the stars in the images is processed manually. Having to annotate manually greatly reduced the amount of training and validation data due to the time limit of this project. However, using the COCO pre-trained weights allowed for a quick training due to the simple overall shape of the stellar images on the CDD camera. Annotation consisted of classifications primary, secondary, third, and field identifying the main star and companions. Field stars are small background stars that do not appear to be associated with the primary star and are excluded from stellar flux calculations.

The method, called Mask R-CNN, extends Faster R-CNN by adding a branch for predicting an object mask in parallel with the existing branch for bounding box recognition. Mask R-CNN is simple to train and adds only a small overhead. Mask R-CNN is a method that is built upon Faster

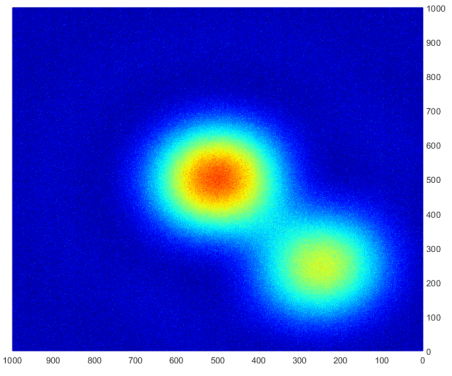

Figure 4: Stellar simulation of a primary star and second secondary star. With added Gaussian and Poisson noise matching the noise behaviour of star KOI799

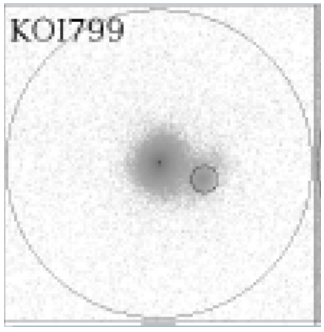

Figure 5: Astronomical observations of KOI799 using adaptive optics. This star has been identified to have an exoplanet with the Kepler telescope, however the secondary star may have introduced serious errors into that calculation.

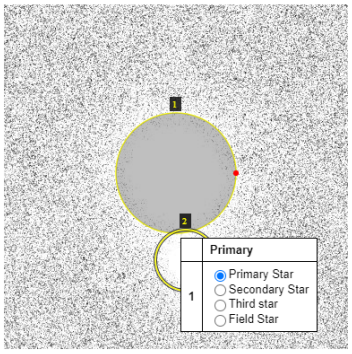

Figure 6: Annotation of the region containing the primary star

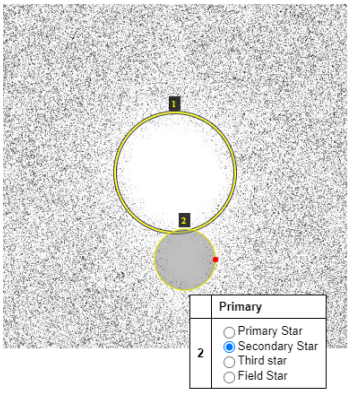

Figure 7: Annotation of the region containing the secondary star

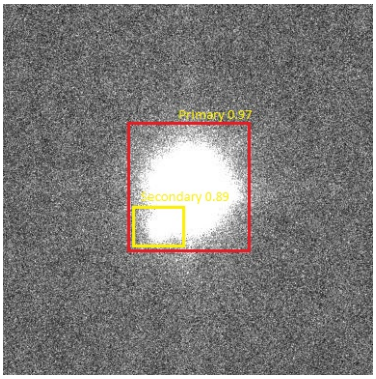

Figure 8: Prediction bounding box for a primary star classification

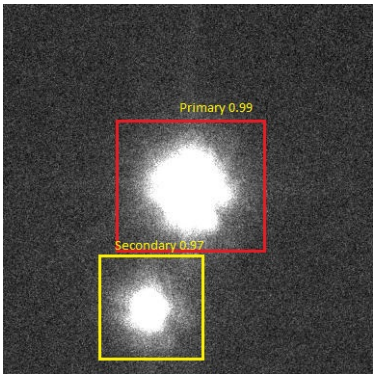

Figure 9: Prediction bounding box for a primary and secondary star classification

R-CNN. It accomplishes this by a parallel addition with bounding box recognition with an object mask prediction. Mask R-CNN consists of a small FCN applied to a Region of Interst (RoI) with bounding box classification regression. The mask branch computational overhead is small. Faster R-CNN wasn't designed for pixel-pixel alignment in the network as seen in (18,12). The solution is to add a layer called RoIAlign that is a quantization free layer which preserves spacial locations.[6]

5 Results

Training a Mask R-CNN is through considering RoI to be positive when it has an Intersection of Union (IoU) of at least 50% otherwise it is set to negative values. The mask consists of positive RoI and the intersection between ROI and the original image. The backbone of the network is a resnet1010 with backbone shapes of 256x256 followed by level of sizes 128x128, 64x64, 32x32, and 16x16. The backbone strides are 4, 8,16,32,64 with a batch size of 1. The bounding box standard deviation is 0.1, 0.1,0.2, 0.2 for each corner. The detection of maximum instances is set at 100 with a minimum confidence of 90% with an initial threshold of 30%. A single GPU was used for this project processing a single image at a time with the image dimensions between 1024x1024 and 800x800 pixels with no RGB channels given the CCD camera is in grey scale, there fore the image shape is 1024x1024x1.

The learning momentum is set at 0.9 with a learning rate of 0.001. [7] The mask pool size is 14 and the shape is 28x28. The mini-mask shape is 56x56 searching for a star. The number of classes is simply star or no star. The pool size is 7 and the RoI positive ratio is set at 33%. The anchor ratios are 0.5, 1, and 2 with scales 32, 64, 128, 256, and 512. The anchor stride is 1 with up to 256 anchors

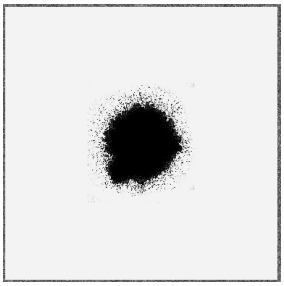

Figure 10: Mask interpolation for primary star. Note the noise in the Mask edges that is not filtered out.

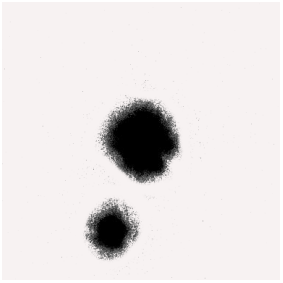

Figure 11: Mask interpolation for primary and secondary star. Note the noise in the mask from the star is not quite filtered out.

per image. The training session takes 100 steps per epoch with a weight decay of 0.0001. Rather than using bounding box coordinates from the datasets, the bounding boxes are computed from the masks. This allows more consistent handling of the bounding box despite the dataset source. It also allows for easier image transformations as rather than having to computer the bounding box again after a transformation we can use the masks to resize, rotate, and crop the images. Images are resized to 1024x1024 from their existing size of 800x800, and maintaining the aspect ratio by adding zero padding around the image where required. [8]

High resolution images can produce quite large binary masks if there are a lot of objects to be masked on an image. Each masking could be around 1MB, but if there are over 100 instances of different masking that adds up to 100 MB for the masks alone. To improve the speed the masks can by optimized by storing the mask pixels that are I inside the bounding box, rather than the mask of the entire image. This saves a great deal of space not storing the entire image for each object. Additionally, the mask can be resized to a size that is typical of an object image. This comes with the risk that the object may be larger than the mini mask. [9]

Currently, the the object is detected or not, and the mask layer is used to calculate the center of each stellar object, along with the circumference on the CCD camera. Using this the relative flux of the star is calculated and subtracted from both stars giving an accurate estimate of the magnitude and distribution of the light sources. Following the PNR, targets for PNR are become the training values. An anchor grid is created to cover the full image at different scales, which are modified by calculating the intersect and union (IoU) of the anchors. With a range of $0.3 >= \text{IoU} < 0.7$ are excluded from training. The PNR is trained through shifting and scaling to cover the image completely. Following this refinement, the PNR graph can be predicted. Anchors within the image are selected for object detection probability of classification. Mask generation generates segmentation masks for each instance and detected for class. [6]

6 Conclusion

Mask R-CNN has shown to be an excellent choice for this application with an ability to find multiple objects intertwined in starlight to an accuracy of 0.98 in our verification tests. The masks produced for each star give a highly accurate estimate of the stellar flux. It would be of interest to compare the accuracy of Mask R-CNN to astronomical software such as IRIS. Further work is needed to build a larger training annotation set which unfortunately will have to be produced manually.Future work will see the integration of this program into the astronomical data pipeline, allowing the system to automatically detect secondary companions to the primary star.

7 Contributions

This project was done by Simon Engler. Dr. Barnec an the astronomer who provided observational data and pointed to research papers in support of this research.

References

[1] W Abdulla. Mask r-cnn for object detection and instance segmentation on keras and tensorflow: matterport/mask_rcnn, 2018.

[2] P.H. Barchi, R.R. de Carvalho, R.R. Rosa, R.A. Sautter, M. Soares-Santos, B.A.D. Marques, E. Clua, T.S. Gonçalves, C. de Sá-Freitas, and T.C. Moura. Machine and deep learning applied to galaxy morphology - a comparative study. *Astronomy and Computing*, 30:100334, 2020.

[3] Gabriel J Brostow, Julien Fauqueur, and Roberto Cipolla. Semantic object classes in video: A high-definition ground truth database. *Pattern Recognition Letters*, 30(2):88–97, 2009.

[4] A. Dutta, A. Gupta, and A. Zissermann. VGG image annotator (VIA). http://www.robots.ox.ac.uk/ vgg/software/via/, 2016.

[5] Abhishek Dutta and Andrew Zisserman. The VGG image annotator (VIA). *arXiv preprint arXiv:1904.10699*, 2019.

[6] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.

[7] Isabel Lily Colman. *Pixels, photometry, and population studies: variable stars across four years of Kepler data*. Doctor of philosophy ph.d., 2019-12-20.

[8] Jessica Schonhut-Stasik, Daniel Huber, Christoph Baranec, Claire Lamman, Maïssa Salama, Rebecca Jensen-Clem, Dmitry A. Duev, Reed Riddle, S. R. Kulkarni, and Nicholas M. Law. Robo-ao kepler asteroseismic survey. ii. do stellar companions inhibit stellar oscillations? *The Astrophysical Journal*, 888(1):34, Jan 2020.

[9] Carl Ziegler, Nicholas M. Law, Christoph Baranec, Ward Howard, Tim Morton, Reed Riddle, Dmitry A. Duev, Maïssa Salama, Rebecca Jensen-Clem, and S. R. Kulkarni. Robo-ao kepler survey. v. the effect of physically associated stellar companions on planetary systems. *The Astronomical Journal*, 156(2):83, Aug 2018.