Object Detection from Radon Transformations using Machine Learning Techniques

Thomas Walker
Lockheed Martin Australia
Boris Repasky, Timothy Payne, Greg Madsen
Lockheed Martin Australia

ABSTRACT

Efficiently examining astronomical images for objects is becoming increasingly important as the space domain becomes more crowded. Our novel approach overcomes several challenges posed by using digital single lens reflex cameras to detect objects from Earth, including low signal-to-noise ratios, small object sizes and interference from other light sources. Leveraging the superior line feature detection of the Radon transformation, we show how a heatmap regression based convolutional neural network applied to the Radon space can drastically improve low and medium Earth object detection when compared to traditional Radon transform based methods when there is significant noise present from artifacts.

1. INTRODUCTION

1.1 Context and Motivation

As the space domain becomes increasingly crowded with man-made objects, the threat posed by space debris on satellite positioning and launches is made more real. To mitigate this risk, constant monitoring of the space domain is becoming more important so that object trajectories can be catalogued, maintained and avoided.

To make a pervasive ground to space observation network viable, low cost sensors are required. This constraint requires the development of image processing techniques that are able to detect streak like signatures that are very close to the noise floor. To further enhance detection capabilities this paper investigates the application of machine learning approaches to extract streaks from pre-processed data using the Radon transform.

1.2 The FireOPAL Network

The Desert Fireball Network is a distributed array of low cost digital single-lens reflex (DSLR) cameras across the Australian Outback designed to monitor the sky for meteoroids entering the atmosphere. With slight lens configuration modifications, this sensor array can be adapted and utilised for space situational awareness (SSA) tasks [1]. In 2018, a small number of modified sensors were deployed to create the FireOPAL Network, designed specifically for the SSA application and managed by Lockheed Martin and Curtin University [2]. An example of the sensors in the FireOPAL network is shown in Fig. 1.

This paper builds on results and processing techniques applied to image data obtained from the FireOPAL Network during 2020.

1.3 Related Work and Contribution

Previous work has shown that streak-like signatures can be extracted in a signal maximising way by applying a Radon Transformation [3].

Many techniques for LEO object detection capitalise on the Radon transformation as a line feature extraction tool and assume objects of interest will generate local maxima within the Radon space for a specified region in the pixel space. To reduce the prevalence of false detections, [4] iteratively increased the Radon integration region and examined the impact on the signal to noise ratio. They showed that if the signal to noise ratio from the signature increased in the radon space, their confidence in its validity increased. Another method proposed by [5] showed that extracting peaks from the Radon space and then examining the corresponding light curve from the pixel domain allowed for the rejection of non-streak-like signatures by checking average brightness compared to a predetermined threshold.

Encoder-decoder networks such as autoencoders and variational autoencoders have been successfully applied to diverse areas such as denoising, compression and image segmentation tasks [6]. Convolutional Neural Networks (CNN) encoders typically evolve an embedding space through successive dimension reducing and feature extracting convolutional layers. However, the dimension reduction of successive layers can impact localisation and accuracy [7]. Ronneberger et al. [7] demonstrate that feature localisation precision and accuracy can be preserved by forming bridging connections between convolution layers.

Zhao et al. [8] show how a similar encoder-decoder implementation combined with a Hough transformation can be used to extract line-like features from more complicated scenes. By incorporating this context aware layer into the analysis, they are able extract line-like environmental boarders such as lake boundaries and fences.

This paper seeks to demonstrate that the application of a encoder-decoder model to a Radon transformed domain can improve the detection and rejection rate of objects of interest and false alarms respectively in astronomical images.

Fig. 1. FireOPAL sensor in the field (*source:* [1])

2. METHOD

The Radon transformation is an effective line feature detector and as such has been widely used for astronomical object detection, particularly where exposure times are long. Approaches using the Radon transformation typically rely on the streak within the image causing the highest peak response within the Radon space, making the detector susceptible to false alarms caused by other line-like artifacts such as partially filtered stars, collinear bright points, or hot-pixels [4,5].

2.1 The Radon Transform

The Radon transformation was introduced by Johann Radon [9] and further developed by Deans [10].

For the purposes of this paper, the idea of the transform can be distilled to the following concept. Each (r, θ) in the Radon space represents the integration along the line tangential to a circle centred within the image chip of radius r at the angle θ. This means that the intensity of a linear pattern of bright pixels will be summed, causing the magnitude in the Radon space to be significantly higher. This effect is demonstrated in Fig. 2.

The algorithm used to perform the Radon transformation was developed in Python and produces equivalent outputs to implementations available and therefore has not been discussed in detail within this paper.

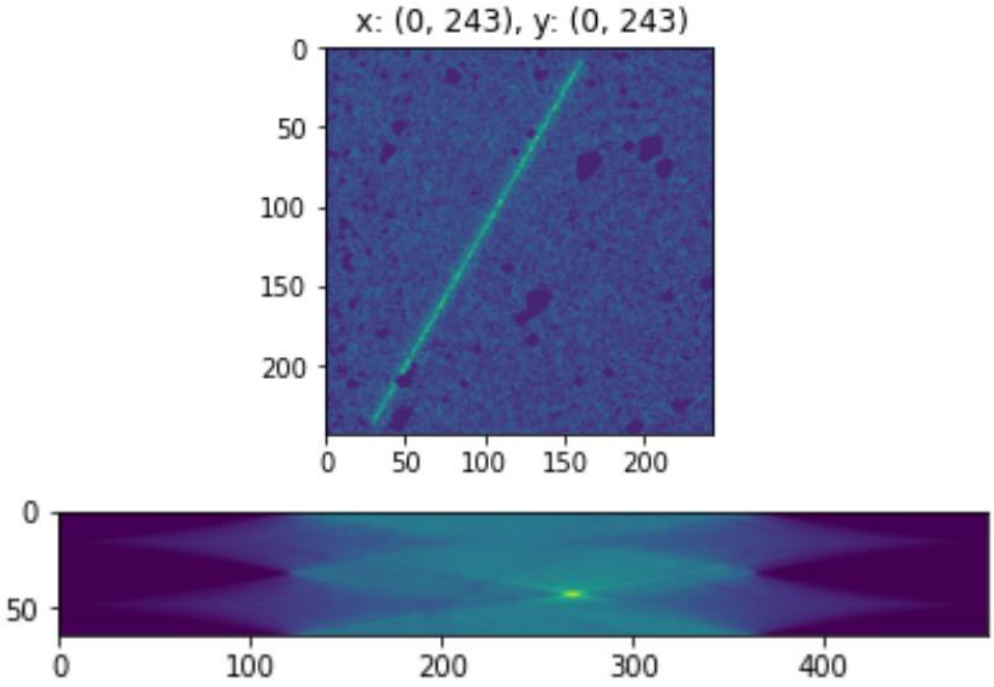

Fig. 2. A synthetic streak caused by an object in low earth orbit (LEO), and the resulting Radon space

2.2 Data Overview and Synthesis

Supervised machine learning approaches depend on a ground truth dataset to allow informed gradient decent based on a cost function. In order to improve upon existing detectors, a synthetic dataset was developed. This allowed for the rapid generation of samples with signals closer to the noise floor than could be observed *en masse* from real images.

Models trained on just synthetic data can often poorly generalise to the real domain when the synthetic data does not perfectly match the characteristics and distribution of the real data. To mitigate this risk, the synthetic data was generated by applying synthetic streaks, derived from an empirical point spread function (PSF) with random noise applied, to a real background noise profile from the real data distribution. The PSF used to simulate the presence of streaks and artefacts within samples was derived from calibration data for the Spitzer Space Telescope shown in Fig. 3 [11]. This source was chosen as it preserved the influence of the diffraction spikes and Airy rings which have an impact on the surrounding image noise profile.

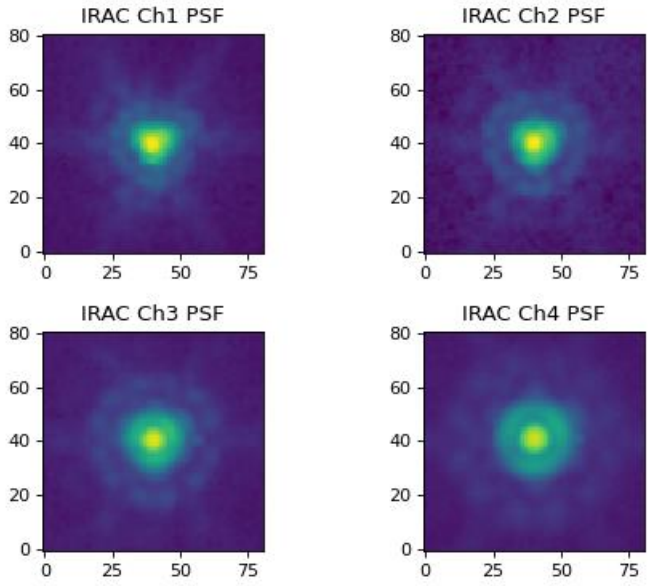

Fig. 3. Empirical point spread function (*Source: [10]*)

Using synthetic data also facilitated the ability to add "confusers" to samples. In this context, we define confusers as non-target objects such as stars or other sources of interference such as lens artifacts that cause significant responses in the Radon space. Confusers were developed by embedding the PSF profile within the sample and applying random warping transformations, noise and mask thresholds.

An example of a real streak and a deliberately similar synthetic streak as well as a confuser can be seen in Fig. 4.

Once the synthetic image had been generated, it was passed through the Radon Transformation, which has an inherent smoothing property through its integration process. This smoothing effect embeds the noise profile within the real background, making the synthetic training image more indistinguishable from real examples.

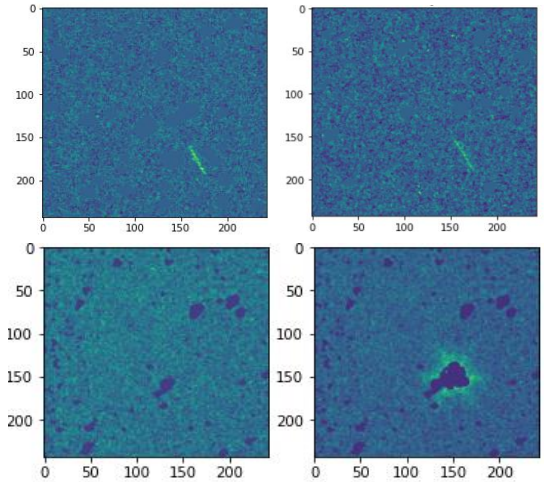

Fig. 4. Top left: real streak found in raw data, top right: deliberately similar synthetic streak, bottom left: real image with just background noise, bottom right: real background noise with a confuser a synthetically added.

2.3 Model Architecture

Given the signature in the Radon space of a streak also captures information from the entire chip in the image domain, an architecture that was aware of broader context was desired. This would enable the model to distinguish between peaks caused by objects of interest and poorly cropped stars. As discussed in [7], the contracting and expanding paths of the U-Net architecture enables context aware localisation that is based on an entire image rather than a sliding window as in region based segmentation models. They also demonstrate that the model is able to learn complex features from relatively small datasets by leveraging data augmentation [7].

3. RESULTS

3.1 Training

The model was trained on a dataset of 5200 samples with the distribution shown in Fig. 5. Training was conducted with a learning rate scheduler, ensuring that the cost function was minimized efficiently. A batch size of 8 was used, as that effectively capitalised on the available memory while allowing a fast learning pace.

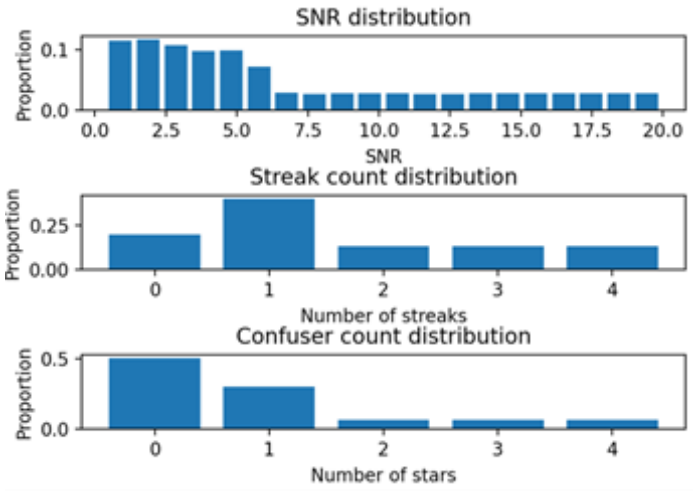

Fig. 5. Training set key parameters

3.2 Detection Results

Testing the model on unseen data from the same distribution shows good performance in both precision and recall.

To highlight the limitations of a naive maximum detector in the radon space and to replicate conditions found in real data, the test dataset deliberately includes chips in which classifying the peak response as the object of interest would be incorrect. This is true for approximately 50% of the test set.

Given the inherent line feature extraction behaviour of the Radon Transformation, a relevant benchmark is a naïve detector that locates targets based on the peak response value location within the transformed domain. The implementation of this detector outputs exactly one prediction per sample.

So that the U-Net implementation can be meaningfully compared to the naïve detector and have its advantages highlighted, it was tested on three, 1200 sample datasets. The datasets are varied in terms of the prevalence of streaks and confusers and are characterised as follows.

- Dataset 1: Single streak in each sample and a confuser in 50% of the samples.
- Dataset 2: Single streak in each sample and a confuser in 100% of the samples.
- Dataset 3: Multiple streaks in each sample and a confuser in 50% of the samples.

A comparison between the naïve detector and the U-Net implementation is conducted on the datasets with one target per sample. To capture the performance of U-Net on more complex Radon transformation domains, it is further tested on the dataset with multiple targets per sample. The performance of these runs is shown in Table 1.

Table 1: Results from different datasets

Detector	Dataset	Precision	Recall	F1 Score
Naïve Maximum	1: Single streak, 50% confuser rate	0.428	0.494	0.459
Naïve Maximum	2: Single streak, 100% confuser rate	0.032	0.037	0.034
U-Net	1: Single streak, 50% confuser rate	0.892	0.712	0.792
U-Net	2: Single streak, 100% confuser rate	0.731	0.309	0.434
U-Net	3: Multiple streaks, 50% confuser rate	0.912	0.726	0.808

The limitation of the naïve detector is the inability to distinguish between peaks caused by objects of interest and spurious peaks caused by artefacts such as poorly masked stars. Given this characteristic, a precision value (true positives / number of detections) and recall rate (true positives / number of targets) less than the confuser rate is to be expected.

Fig. 6. shows the successful detection of a streak by the U-Net model and the incorrect detection by the naïve algorithm, and Fig. 7. shows the ability of the U-Net implementation to detect multiple signatures within one frame..

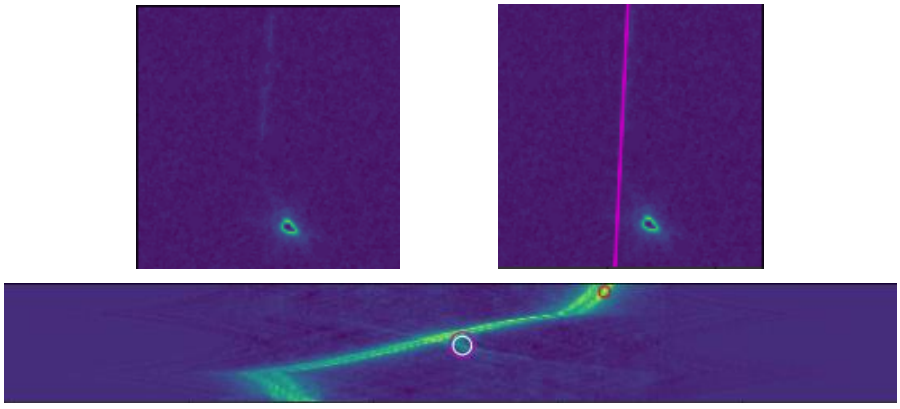

Fig. 6. Example of a confuser causing an incorrect detection from the Naive Maximum detector (shown by the red circle) and the correct detection from the U-Net implementation (shown by the purple circle in the Radon space and the purple trajectory in the middle

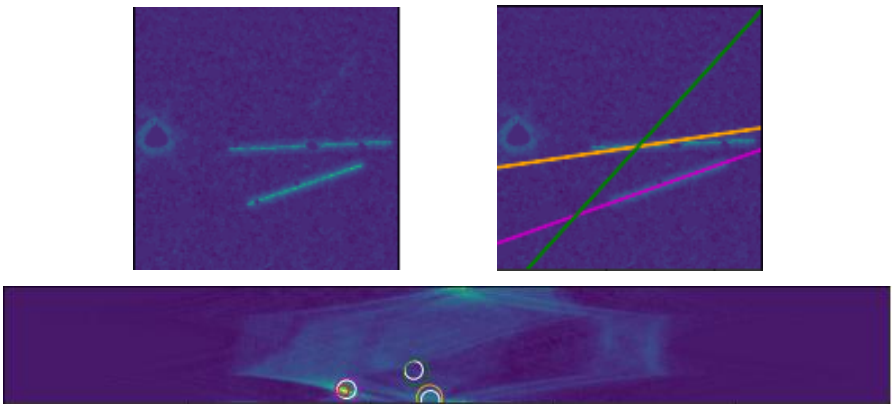

Fig. 7. Demonstration of the ability to detect multiple streaks within an input frame with the U-Net implementation.

4. CONCLUSION

This paper explores the use of a machine learning approach applied to a classical transformation space for object of interest detection withing astronomical images. In the Radon transformation representation of an image, line-like features are made prominent. We show that by using a heatmap formulation based on the U-Net model, one is capable of distinguishing responses from desired targets and spurious peaks resulting from bright confusers such as poorly masked stars or hot-pixels. In a test set where a naive maximum classifier is able to identify a correct response in 3% of the samples, the U-Net implementation is able to correctly identify approximately 30%, highlighting the difference in detection sophistication.

4.1 Future Work

One challenge associated with this classifier is that it is significantly more resource intensive compared to a naïve maximum algorithm. Depending on the application, performing the U-Net inferencing on all chips of a frame may not be possible. As such, the methodology outlined in this paper may be suited best for use as a false alarm discriminator where a naive maximum or similarly lightweight process identifies candidate chips and the U-Net model seeks to localize the target or reject if necessary.

5. REFERENCES

1. Jansen-Sturgeon, T. et al., FireOPAL: Continental-scale Coordinated Observations of the OSIRIS-REx Flyby. *Proceedings of AMOS*, 2018.
2. Madsen, G.J. et al. FireOPAL: Technical Performance and First Results. *Proceedings of AMOS* , 2018.
3. Toft, P. A. et al., The Radon Transform - Theory and Implementation. *Kgs. Lyngby, Denmark: Technical University of Denmark (DTU)*, 1996.
4. Nir, G. et al., Optimal and Efficient Streak Detection in Astronomical Images, *The Astronomical Journal*, 156:229, 2018.
5. Ciurte, A., Danescu, R., Automatic Detection of MEO Satellite Streaks from Single Long Exposure Astronomic Images, *Proceedings of the 9th International Conference on Computer Vision Theory and Applications*, 2014.
6. Shervin, M. et al., Image Segmentation Using Deep Learning: A Survey, *Computing Research Repository (CoRR)*, 2020.
7. Ronneberger, O. et al., U-Net: Convolutional Networks for Biomedical Image Segmentation, *University of Freiburg*, 2018.
8. Zhao, K. et al., Deep Hough Transform for Semantic Line Detection, *IEEE Transactions On Pattern Analysis And Machine Intelligence*, 2020.
9. Radon, J., Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten, *Berichte über die Verhandlungen der Königlich-Sächsischen Akademie der Wissenschaften zu Leipzig, Mathematisch-Physische Klasse*, Leipzig: Teubner. (69): 262–277, 1917.

Translation: Radon, J., Parks, P.C. (translator), On the determination of functions from their integral values along certain manifolds, *IEEE Transactions on Medical Imaging*, 5(4): 170-176, 1986.
10. Deans, S.R., The Radon Transform and Some of Its Applications, *John Wiley & Sons*, 1983.
11. Hora, J.L. et al., In-flight performance and calibration of the Infrared Array Camera (IRAC) for the Spitzer Space Telescope, *Optical, Infrared, and Millimeter Space Telescopes*, 2004