




Contents lists available at ScienceDirect

# Astronomy and Computing

journal homepage: [www.elsevier.com/locate/ascom](http://www.elsevier.com/locate/ascom)

Full length article

## Towards asteroid detection in microlensing surveys with deep learning

P. Cowan<sup>\*</sup>, I.A. Bond, N.H. Reyes

School of Mathematical and Computational Sciences, Massey University, Private Bag 102-904 North Shore Mail Centre, Auckland, 0745, New Zealand

### ARTICLE INFO

#### Article history:

Received 18 October 2022

Accepted 23 January 2023

Available online 30 January 2023

#### Keywords:

Microlensing surveys

Asteroid detection

Deep learning

Convolutional neural networks

YOLOv4

MOA

### ABSTRACT

Asteroids are an indelible part of most astronomical surveys though only a few surveys are dedicated to their detection. Over the years, high cadence microlensing surveys have amassed several terabytes of data while scanning primarily the Galactic Bulge and Magellanic Clouds for microlensing events and thus provide a treasure trove of opportunities for scientific data mining. In particular, numerous asteroids have been observed by visual inspection of selected images. This paper presents novel deep learning-based solutions for the recovery and discovery of asteroids in the microlensing data gathered by the MOA project. Asteroid tracklets can be clearly seen by combining all the observations on a given night and these tracklets inform the structure of the dataset. Known asteroids were identified within these composite images and used for creating the labelled datasets required for supervised learning. Several custom CNN models were developed to identify images with asteroid tracklets. Model ensembling was then employed to reduce the variance in the predictions as well as to improve the generalization error, achieving a recall of 97.67%. Furthermore, the YOLOv4 object detector was trained to localize asteroid tracklets, achieving a mean Average Precision (mAP) of 90.97%. These trained networks will be applied to 16 years of MOA archival data to find both known and unknown asteroids that have been observed by the survey over the years. The methodologies developed can be adapted for use by other surveys for asteroid recovery and discovery.

© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

### 1. Introduction

Asteroids are among the millions of small bodies that inhabit our Solar System and are remnants from its formation. While popular sentiment most commonly associates asteroids with mass extinction events, the vast majority of asteroids pose no threat to us. The Main Asteroid Belt between Mars and Jupiter has the largest concentration of asteroids in our Solar System and this is where the majority of the asteroids seen in this research reside. Observing and tracking these small bodies gives us a better understanding of their complex orbital dynamics. Their composition and structure offer clues about the conditions when the terrestrial planets were formed 4.6 billion years ago.

Asteroids are part of the landscape of our night sky and appear in the imaging data of most astronomical surveys. However, as most surveys have a specialized purpose, their data is rarely mined for asteroids. Microlensing surveys like MOA [Sumi et al., 2003](#), OGLE [Udalski et al., 1993](#), and KMTNet [Kim et al., 2016](#) are particularly good for determining the rotation period and orbital trajectory of asteroids because they survey a given region of space several times each night [\(Gould and Yee, 2013\)](#). This means that asteroids could spend several nights in the field of view of

the telescope, giving us the opportunity to both observe their trajectory and analyse the light gathered from them. [Cordwell et al. \(2022\)](#) demonstrates the efficacy of extracting asteroids light curves from the MOA microlensing data.

Automated detection software has been part of surveys dedicated to discovering asteroids since the early 90s [\(Rabinowitz, 1991\)](#). With improved computing power, other techniques for detecting moving astronomical sources such as shift and stack have also proven popular. In recent years, a leader in the field is the Pan-STARRS Moving Object Processing System or MOPS [\(Denneau et al., 2013\)](#). Initially trained with simulated but realistic asteroid data for the Pan-STARRS telescopes, it takes transient candidates not associated with a known source and uses a complex tree-based spatial linking algorithm [\(Kubica et al., 2007\)](#) to further parse and form associations between these point sources. MOPS does not work with imaging data but rather celestial coordinates, which reduces the computational cost. HeliolinC [\(Holman et al., 2018\)](#) further improves on MOPS' efficiency with an approach that combines working with a heliocentric frame of reference and clustering sources that belong to the same object.

While these and other deterministic approaches have been successfully utilized for asteroid detection, applications of deep learning in the field remain in the early stages, potentially because of the lack of labelled data. Deep learning offers the benefit of being able to learn representations directly from the raw data,

<sup>\*</sup> Corresponding author.

E-mail address: [preeti.cowan@auckland.ac.nz](mailto:preeti.cowan@auckland.ac.nz) (P. Cowan).

making it a potentially valuable tool for asteroid discovery in archival astronomical data. The works that do apply deep learning techniques note the benefits, particularly with greatly reducing the amount of data that must be examined by an astronomer, as we see next.

Zoghbi et al. (2017) successfully applied both convolutional and recurrent architectures to reduce the amount of data to be vetted by astronomers looking for debris from long-period comets in the CAMS data.<sup>1</sup> Lieu et al. (2018) applied neural networks to the task of detecting small solar system objects (SSO) in data simulated for the ESA's Euclid space telescope.<sup>2</sup> They successfully used transfer learning and retrained three architectures from TensorFlow's Keras Applications library<sup>3</sup> to distinguish between postage stamp cut-out images of asteroids and objects commonly mistaken for asteroids like cosmic rays, stars, and galaxies. Duev et al. (2019) introduced DeepStreaks to aid in the ZTF's<sup>4</sup> quest for the discovery of near-Earth asteroids, which resemble streaks in the observations. Their model significantly reduced the number of candidate detections that had to be reviewed without sacrificing the detection sensitivity. Rabeendran and Denneau (2021) applied deep learning to the ATLAS<sup>5</sup> pipeline looking for near-Earth objects. It was successful in catching nearly 90% of the false positive detections, thus greatly speeding up the process of followup observations. Duev et al. (2021) introduced Tails, which involved training an object detector to discover comets based on their distinctive morphology and it now forms a part of the ZTF's detection pipeline. Finally, Kruk et al. (2022) used deep learning to hunt for asteroid trails in archival data from the Hubble space telescope (HST).<sup>6</sup> They used composite HST images to make the asteroids trails longer and thus easier to detect. Their research also demonstrates the merits of citizen science for labelling the data and of mining archival data for asteroids with a deep learning-based toolkit.

We have seen that an important aspect of finding asteroids in survey data is isolating the sources that indicate a candidate object. In the case of microlensing data, this task is all the more challenging because of the extremely dense star fields observed, causing even the reference subtracted images to contain numerous spurious artefacts. Thus, rather than extracting the sources and then establishing connections between them, we propose a technique to instead enable extracting a cluster of sources that could represent part of the orbital arc of asteroids.

We present the methodology to create labelled datasets suitable for supervised learning, followed by convolutional neural network-based architectures for finding asteroids tracklets in microlensing surveys. Our classification models facilitate reducing the amount of data to be vetted and the object detection model localizes the potential tracklets within the classifiers' candidate detections. The bounding boxes predicted by the object detector could then be used to extract the potential sources, which in turn could be passed to orbiting link software, such as HeliolinC, to determine if they are valid solar-bound objects.

## 2. MOA: Microlensing observations in astrophysics

This research utilizes archival data from the MOA (Microlensing Observations in Astrophysics) project, which is a Japan/New Zealand collaboration specializing in the search for gravitational




Fig. 1. A single observation of the Galactic Bulge field, GB1, surveyed by MOA-II. North is right to left and east is bottom to top.

microlensing events. They have been operating the 1.8 m MOA-II optical research telescope at the Mt. John Observatory since 2004. It has a wide-field mosaic CCD camera called the MOA-cam3 (Sako et al., 2008), which consists of 10 CCD chips. Each CCD chip is 3 cm x 6 cm and has 2048 x 4096 pixels, with a resolution of 0.01 arc minutes/ 0.6 arcseconds per pixel. Each exposure is 60 s long and all images are in a custom wide MOA-R band (630–1000 nm). The median FWHM seeing is 1.7 arcseconds. The photometric precision is typically 0.01 mag or better for I brighter than 16 and 0.02 mag for I 18. The telescope has a total field of view of 1.7 x 1.4 square degrees. The survey field referred to as GB1 can be seen in its entirety in Fig. 1. The CCD chips are numbered 1 to 10 clockwise, starting from the top-left.

MOA-II surveys the Galactic Bulge (GB) and the Large Magellanic Cloud (LMC), both of which are regions of the sky that are densely packed with stars. It operates at a high sampling rate, with some fields surveyed as often as every 10 min, which makes it particularly good for observing short duration events. The MOA image processing pipeline produces difference images (Tomaney and Crofts, 1996; Alard and Lupton, 1998; Alard, 2000; Bramich, 2008; Bond et al., 2001), where a new observation image is subtracted from a reference image for the same region. The resultant difference image highlights the changes since the reference image was taken. These could either be transient astronomical phenomenon (like microlensing events or asteroids) or noise. The noise could be due to too-bright saturated stars, imperfect subtractions, satellite trails, instrumentation error, atmospheric dust, differential refraction, or proximity to a bright astronomical object (like the moon). As the fields surveyed by MOA-II are very dense, faint asteroids in particular are impossible to identify in the original observation images. Thus, the MOA difference images form the basis of the datasets built for this research.

## 3. Building the dataset

The Galactic Bulge is visible in the Southern Hemisphere between February and October. A combination of long nights and Bulge elevation in the mid-winter months make them the best time for observation. On clear nights with good seeing conditions, several of the MOA-II fields are surveyed frequently, at times as often as every 10 min. This provides ideal conditions for observing asteroids in our solar system. Fig. 2 demonstrates this with six observations of a bright main-belt asteroid (78,153) 2002 NX24,

<sup>1</sup> <http://cams.seti.org/>


<sup>2</sup> <https://sci.esa.int/web/euclid/>

<sup>3</sup> <https://keras.io/api/applications/>


<sup>4</sup> <https://www.ztf.caltech.edu/>

<sup>5</sup> <https://atlas.fallingstar.com/home.php>

<sup>6</sup> <https://hubblesite.org/>




**Fig. 2.** Six consecutive observations of the asteroid (78,153) 2002 NX24 on 23-June-2006, spanning 70 x 108 arcseconds.



**Fig. 3.** Light curve of asteroid (78,153) 2002 NX24 extracted from the MOA-II data.

with a limiting magnitude ( $V$ ) of 17.5, taken at 10–12 min intervals on the 23rd of June, 2006. The light curve for this asteroid, extracted from the MOA-II data can be seen in Fig. 3. The sky-plane velocities are about 0.5 arcseconds per minute, which are representative of the asteroids captured by MOA-II.

The primary data used in this research consists of all observations from 2006–2019 from one chip (chip 5) in the CCD array for the field GB5, amounting to 49,901 difference images (GB5-R5). This field is surveyed very frequently and the cadence of most observations is between 5 and 20 min. The Minor Planet Center (MPC) was queried<sup>7</sup> via a screen scraper to get all known asteroids expected to traverse through the region encompassed by GB5-R5. To determine the limiting magnitude of our survey, we have compared the fluxes measured at the positions of recovered asteroids from GB5-R5 with the  $V$  magnitudes provided by MPC. The results are shown in Fig. 4, where the data has been binned by MPC  $V$  magnitude with a bin size of 0.1 mag with



**Fig. 4.** Measured MOA flux vs the MPC  $V$  magnitude for asteroids recovered from GB5-R5. The MPC value is reported to one decimal place and the results are binned in 0.1 magnitude bins. The red dots represents the median values of the fluxes. This starts to flatten out when  $V$  is around 21, at which point it is effectively the background being measured. The green line represents the working limiting flux commonly used in MOA microlensing work.


the median flux values calculated for each bin. For microlensing measurements, MOA adopts a working value of 2500 for the limiting data number. This corresponds to a MPC  $V$  magnitude 20.5 which we adopt as our limiting magnitude here.

Since asteroids move appreciably with respect to the background stars in each consecutive exposure, asteroid tracklets can be clearly seen by stacking nightly observations. We define a tracklet as part of the orbital arc of the asteroid as it travels through the field of view of the telescope. We started with a stack composed of the brightest pixels. While the tracklets were visible, so was the bright background and the over-bright stars (Fig. 5(a)). To highlight the tracklets better, the stack image was further simplified by subtracting the brightest pixel stack from the median pixel stack. This gave us an image without the bright background and saturated stars (Fig. 5(c)), leaving behind the noise and moving objects. Only nights with 3 or more observations were considered when generating the subtracted stacks. The GB5-R5 dataset resulted in 2252 subtracted stack images within which the search for asteroid tracklets commenced.


The ephemeris data from the MPC, together with the astrometric calibrations specific to MOA-II, were used to extrapolate the minimum and maximum ( $x, y$ ) positions for asteroid tracklets in subtracted stack images. This was used to crop sub-regions expected to contain tracklets from each stack. Asteroids that fell beyond the boundaries of a CCD chip were ignored. Further, by visual inspection of the tracklets, it was established that only asteroids with a limiting magnitude of 20.5 or brighter are visible in the MOA-II exposures. This is also consistent with what we see in Fig. 4, where the median flux around  $V$  21 flattens out and is akin to measuring the background. The objects with magnitude between 19.5 and 20.5 are often at the very edge of visibility, requiring excellent seeing condition and a high signal to noise ratio to be visible in the observations. Careful examination of the images resulted in 2073 tracklets from 1178 distinct asteroids over 1078 distinct nights. Fig. 6 (374 x 387 arcseconds) displays some of the tracklets captured in the observations on the night of 15-May-2008. Some of the tracklets are very faint.

Tracklets come in a variety of sizes, but we need images of a uniform size for training neural networks. The original size of the observations was deemed too big to be used as is because of the computational cost as well as the potential for the tracklets to be lost amongst the other artefacts in the image. Therefore, a decision was made to split each image into 128 x 128 tiles, giving us

<sup>7</sup> <https://www.minorplanetcenter.net/cgi-bin/checkmp.cgi>



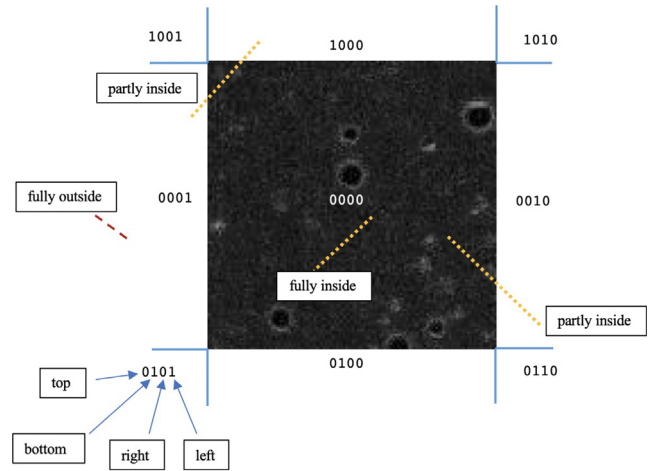
**Fig. 5.** Stacking all of a night's observations (23-June-2006) on one chip in one field gives us a clear tracklet for asteroid (78,153) 2002 NX24. In (a) the observations are stacked by brightest pixel; in (b) they are stacked by the median pixel; and in (c) we see the result of subtracting (b) from (a).



**Fig. 6.** A stack of all 51 observations from the night of 15-May-2008 reveals 5 asteroid tracklets clearly visible in a sub-region of the stack image. Numbered in order of appearance, these are: 1 - (103,842) 2000 DQ33 (19.5); 2 - (148,657) 2001 SX124 (19.6); 3 - (152,083) 2004 RH30 (19.9); 4 - (582,743) 2016 AT221 (20.4); 5 - (338,789) 2005 SZ154 (20.4). The line/streak on the top left is from a satellite. This image has been inverted and brightened to improve visibility (original in Fig. A.15).

512 tiles per image. Each  $128 \times 128$  tile spans  $76 \times 76$  arcseconds. For GB5-R5, this gave us a total of 551,936 images from the 1078 nightly stacks that contained visible asteroid tracklets.

The Cohen–Sutherland line clipping algorithm (Newman and Sproull, 1973) was next used to locate the asteroid tracklets within these images. The algorithm involves dividing a rectangular space (in this case, an image) into nine regions – eight “outside” and one “inside”, as seen in Fig. 7 – and determining which of the lines (tracklets) are fully or partially inside the area of interest (the image). Each of the nine regions have associated outcodes (4-bit numbers) that are calculated by performing bitwise operations after comparing the start and end points of the line/tracklet with the coordinates of the image. For example, the outcode 0001 indicates that the end point is centre left and the outcode 1000 indicates that the end point is centre top. A bitwise OR of these two codes returns 1001, which is a non-zero value, and the bitwise AND returns 0000, which is zero, which in turn




**Fig. 7.** The Cohen–Sutherland line clipping algorithm was used to reject tracklets outside the area of interest and determine intersection points of the ones partially inside.

indicates that the tracklet is partially inside the image. Thus, there are three possible solutions for any line:

- If both endpoints of the line are inside the area of interest, the bitwise OR computation returns 0 (trivial accept).
- If both endpoints of the line are outside the area of interest, they will share at least one outside region and the bitwise AND computation returns a non 0 value (trivial reject).
- If both endpoints are in different regions, at least one endpoint will be outside the image tile. In this case, the intersection point of the tracklet's outside point and image tile boundary becomes the new endpoint for the tracklet and the algorithm repeats until the bitwise operation returns a trivial accept or reject.

The application of this algorithm gave us 8341 tiles that potentially had visible tracklets and 543,595 without known asteroids. Once again these were meticulously scanned to ensure that a tracklet could clearly be seen. Any tile with less than 3 points of a tracklet were rejected, along with tiles that did not contain a visible portion of a tracklet. At the conclusion of this process, there were 4153  $128 \times 128$  tiles with visible tracklets from GB5-R5. Some of these can be seen in Fig. 8(a).

For classification, the images where the tracklet was too obscured by noise or where there were less than three clearly visible adjacent point sources related to a tracklet were removed. Following this, there were 4072 images with tracklets, which were split into the training/ validation/ test set with 3322/ 415/ 335 in each, respectively. The images in the training and validation set were further augmented by rotating 180 degrees, flipping horizontally and vertically, brightening, darkening, blurring, and both increasing and decreasing the contrast. There was no overlap between the training, validation, and test sets. A further test set composed of all the observations on a single night from all ten chips in fields GB3, GB4, GB5, GB9, GB10, and GB14 was also constructed (GB-All), and it had 300 tracklet and 2000 no-tracklet images. The purpose of the GB-All test set was to evaluate how the networks performed with data they had never seen; as before, the tracklets were from observations with a maximum cadence of 20 min and there are a minimum of 3 sources per image. The classification networks were also trained to recognize images with no tracklets. Forty images without known tracklets were randomly chosen from each of the 512 sub-regions, giving us 20,480 images. This collection was also visually inspected and



**Fig. 8.** Asteroid tracklets seen in  $128 \times 128$  sub-regions from subtracted stack images. In (b), the tracklets are localized with bounding boxes. Asteroids in the top row, left to right: (38,102) 1999 JM18 (18.4); (152,083) 2004 RH30 (19.9); (103,842) 2000 DQ33 (19.5); (283,261) 2011 FR142 (20.0); and (375,674) 2009 HD5 (19.8). Asteroids in the bottom row, from left to right: (48,617) 1995 HR2 (19.5); (148,657) 2001 SX124 (19.6); (97,948) 2000 QF124 (19.6); and (74,978) 1999 TY234 (18.7). Note that these images have been inverted and brightened to improve visibility (original in Fig. A.16).

any images that could potentially have tracklets were removed. This process resulted in 19,682 images, which were split into the training/validation/test sets with 15,595/2039/2048 images, respectively. The same set of augments was generated for the no-tracklet training and validation data, but since this set was already significantly larger, only a random 35% of the no-tracklet augments were used. Note that this distribution does not reflect the true ratio of tracklet vs no-tracklet images; realistically, on a good night, we might expect as many 10 tracklet images for every 100 no-tracklet images. The roughly 1:5 ratio here serves to better focus the network to learn the tracklet pattern while providing a suitable number of contrasting no-tracklet images.

As the asteroid tracklets are a well-defined pattern of blobs in an image, we further considered training a deep-learning based object detector to localize them in the images. This would facilitate locating the tracklets that were faint and/ or obscured by noise as well as distinguish the area of interest in the images. We utilize YOLOv4 (Bochkovskiy et al., 2020), which requires the objects of interest – in our case, tracklets – to be enclosed in bounding boxes, with the coordinates of the object’s centroid saved with respect to the dimensions of the image. As a result of applying the Cohen–Sutherland line clipping algorithm, we had the probable intersection/end points of tracklets within the  $128 \times 128$  images. These intersection points were used to extrapolate bounding boxes for tracklets, thus highlighting them in the images. Further visual inspection and manual adjustment was undertaken to ensure that bounding boxes encapsulated the tracklets correctly without any superfluous background included (Fig. 8(b)). All 4153 tracklet images were used, with a training/test split of 3737 and 416, respectively.

#### 4. Deep learning framework

The appeal of neural networks and deep learning is in their capacity to make predictions about complex data in real time after they have been trained with a representative dataset. In the last decade, deep learning has matured significantly and has

proven to be highly effective for classification and object detection tasks. Here, we discuss training several convolutional neural network (CNN) based architectures to find asteroids tracklets in the subtracted stack images.

The proposed model consists of an ensemble of five classifiers that produces a probability that a  $128 \times 128$  composite image tile contains a tracklet. Tiles believed to contain tracklets are then passed to YOLOv4, which has been trained to localize tracklets by adding a bounding box around the tracklet within the image.

##### 4.1. Classification

A convolutional neural network or CNN is a network architecture that is designed to take advantage of the 2D structure of an input image. Using a series of convolutional layers that utilize filters, implemented using local connections and shared weights, it can extract meaningful features directly from data, thereby eliminating the need for manual feature extraction.

Traditionally, the number of filters in a layer increases with the depth of the network. The received wisdom is that the more layers a network has, the better it is at extracting more complex features from an image, and thus learning from complex data. The limiting factors here, however, are the size of the input and how localized the features we are interested in are, as well as the amount of available training data. After a given number of layers, a network could start overfitting to the data by focusing on the irregularities in the images. So, while increasing the number of convolutional layers improves the performance of the network, it is not the case that the deeper network is always the best option. Part of the challenge is finding the ideal depth for a network to optimally predict the output. Simonyan et al. introduced VGGNet (Simonyan and Zisserman, 2015) and investigated the effects of increasing the depth of convolutional layers on the model’s classification and localization accuracy in the large-scale image recognition setting (ILSVRC challenge). VGGNet has a compelling simplicity in its architecture; it uses a stack of consecutive convolutional layers to reduce the number of parameters, leading to faster convergence and reducing the overfitting problem.

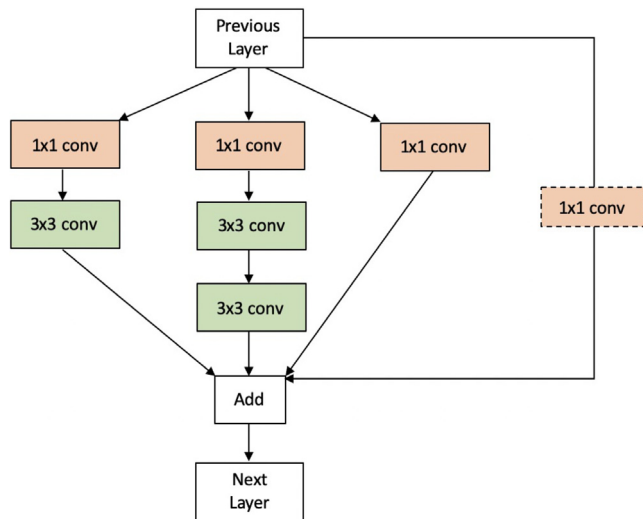


Fig. 9. Hybrid module combining salient features of a ResNet block and an Inception module.

GoogLeNet (Szegedy et al., 2015), popularly known as Inception, introduced the idea of modifying the width of a network as well as employing filters of a variety of sizes to better capture multi-scale data from images. Relative to the VGGNet architecture, it is able to reduce the number of parameters and computational cost further. Residual networks (He et al., 2016), or ResNets, were introduced to address the vanishing gradient problem as networks go deeper by presenting an alternative pathway for the algorithm to follow, called the skip connection. The central element in this architecture is the residual block, which consists of two convolutional layers with  $3 \times 3$  filters. The input of this is added to the output of the second convolution, thus creating a shortcut connection.

In our work, a variety of CNN classification architectures were trialled to determine the best combination of filters and layers to suit our purpose. After considerable experimentation with well-established architectures such as VGG-16 and VGG-19 (Simonyan and Zisserman, 2015), Inception (Szegedy et al., 2015), and ResNet50 (He et al., 2016), various custom models were constructed to determine if better performance could be obtained. Table 1 describes the structure of five of these custom architectures, three of which are VGG-like (MOA-12, MOA-14, MOA-15), and two are composed of hybrid Inception-ResNet modules (Hybrid A, Hybrid B). All of the custom architectures have significantly fewer parameters than their established counterparts. It is our finding that having an over-parameterized deep network architecture causes over-fitting of the training data, and so we reduced the number of layers and filters in the network to reduce the number of neurons and weights involved.

The hybrid module (Fig. 9) consists of four branches composed of convolutional layers, each of which use the ReLU activation (Glorot et al., 2011; Nair and Hinton, 2010) and are combined with the ‘add’ function. The final branch carries the input to the ‘add’ function, with the convolution being used only when the number of filters in the previous layer are not equal to the current layer’s filters. The many  $1 \times 1$  convolutions in the hybrid module may seem superfluous but removing even one negatively affects the performance of the network. This may be because of the beneficial complexity introduced by the associated non-linearity.

All of the classification architectures were trained on a Linux machine running Ubuntu 18.04 with a NVIDIA Quadro M4000 GPU (8 GB, 2.5 TFLOPS). The code was written in Python 3.6 in the Jupyter Notebook environment. TensorFlow GPU 2.4.1 (CUDA

Table 1

Configuration of custom classification architectures MOA-12, MOA-14, MOA-15, Hybrid A, and Hybrid B. ReLU activation is used in each layer and dropout is used after each fully connected layer.

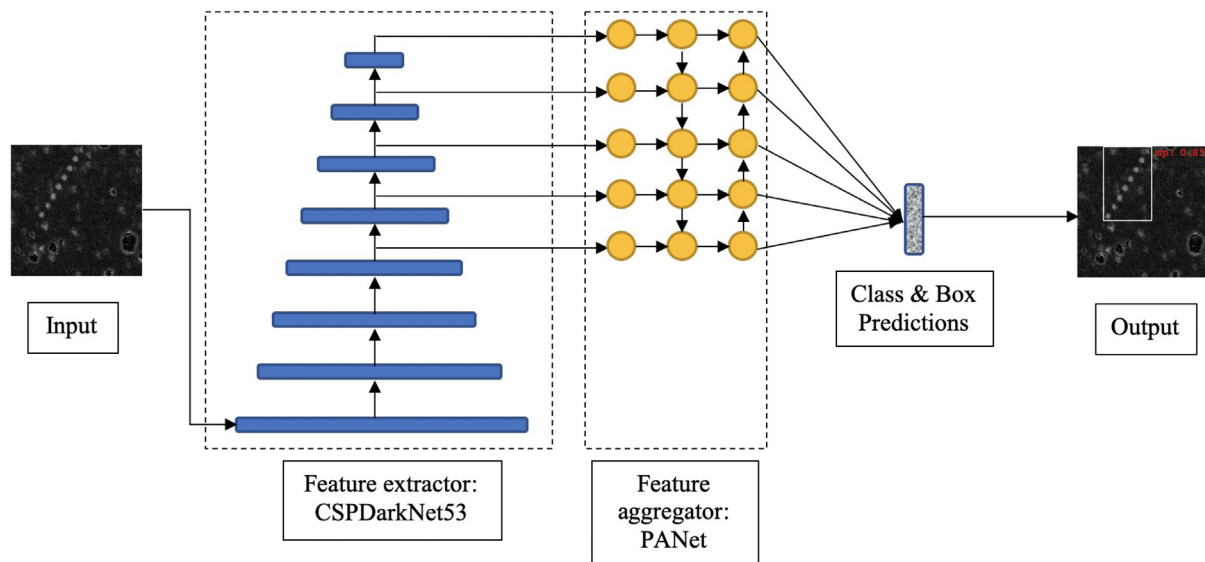
I MOA-12	II MOA-14	III MOA-15	IV Hybrid A	V Hybrid B
128 x 128 subtracted stack image				
32 conv 3x3	32 conv 3x3	32 conv 3x3	32 conv 3x3	32 conv 3x3
32 conv 3x3	32 conv 3x3	32 conv 3x3		32 conv 3x3
		32 conv 3x3		
maxpool 2x2 (stride=1)				
64 conv 3x3	64 conv 3x3	64 conv 3x3	32 hybrid	64 hybrid
64 conv 3x3	64 conv 3x3	64 conv 3x3	64 hybrid	64 hybrid
64 conv 3x3	64 conv 3x3	64 conv 3x3		
maxpool 2x2 (stride=1)				
128 conv 3x3	128 conv 3x3	128 conv 3x3	96 hybrid	128 hybrid
128 conv 3x3	128 conv 3x3	128 conv 3x3	128 hybrid	128 hybrid
128 conv 3x3	128 conv 3x3	128 conv 3x3		
maxpool 2x2 (stride=1)				
256 conv 3x3	256 conv 3x3	256 conv 3x3	160 hybrid	256 hybrid
256 conv 3x3	256 conv 3x3	256 conv 3x3	192 hybrid	256 hybrid
256 conv 3x3	256 conv 3x3	256 conv 3x3		
maxpool 2x2 (stride=1)				
512 conv 3x3	512 conv 3x3	512 conv 3x3	224 hybrid	Global Average Pooling
	512 conv 3x3	512 conv 3x3	256 hybrid	
	512 conv 3x3	512 conv 3x3		
maxpool 2x2 (stride=1)			Global Average Pooling	
dense 256		dense 512	dense 128	
dense 128		dense 256		
sigmoid output				

11.0), along with the Keras deep learning API, was used for creating the custom CNN models tested. The Keras Applications implementations were used for the established models. A batch size of 32 was used for training each network model. Training was set to run for 50 epochs, with a callback for early stopping if the validation loss failed to minimize after a set number of epochs. The Adam (Kingma and Welling, 2014) optimizer was used and the loss function minimized was binary cross entropy. The learning rate was initialized at 0.0001 for all networks except Hybrid A, which started with a learning rate of 0.001. In each case, the learning rate was reduced after 15 epochs and at scheduled intervals after that point. Callbacks were included to save the best weights for both validation and training accuracy. The established networks performed better when pre-loaded with ImageNet weights before fine-tuning with the MOA-II data. MOA-12, 14, and 15 took on average 3 h to train and the Hybrid models, 6 h.

#### 4.2. Object detection

In recent years, CNNs have been taken beyond classifying images and have been applied to object detection and localization in images. Here, we utilize the YOLOv4 object detection architecture for localizing asteroid tracklets in our tile images.

The aim of the original YOLO (You Only Look Once) object detection architecture (Redmon et al., 2016) was to make object detection both fast and accessible. As a single-shot architecture, it was capable of making class and bounding box predictions with the feature maps produced by a single CNN network. The simplicity of the feature extraction CNN at the heart of YOLO was in direct contrast to its complex loss function, which computed the classification loss, localization loss, as well as the loss quantifying the network’s confidence in the prediction. Later versions of the network introduced default anchor boxes, which are predefined bounding boxes that are adjusted and refined during training to encompass the objects of interest in an image (Redmon and



**Fig. 10.** Architecture of YOLOv4: a feature extractor that is composed of a 53-layer DenseNet with cross-stage partial connections and spacial pyramid pooling along with a feature aggregator that effectively combines the feature maps from the higher and lower resolution layers.

Farhadi, 2017), and more complex feature extractors to make predictions on multiple scales (Redmon and Farhadi, 2018).

This research applied YOLOv4 (Bochkovskiy et al., 2020), which is the latest evolution of the architecture with a major overhaul to include several new techniques, making the model state-of-the-art while still being easy to train. In particular, YOLOv4 ensures that the lower level features are propagated through both the feature extractor as well as the feature aggregator. Fig. 10 illustrates the architecture of YOLOv4. The convolutional backbone feature extractor for the architecture is composed of a 53 layer DenseNet (Huang et al., 2016) with the cross-stage-partial (CSP) connections of CSPNet (Wang et al., 2020). DenseNets extend ResNet’s concept of skip connections by adding connections between all the layers in the network in a feed-forwards fashion. Feature maps from all preceding layers are concatenated and form the input for any given layer, ensuring that low-level features are propagated through the network. CSP connections involve splitting the input feature map into two parts, one of which goes through the dense block and the other goes straight through to the next transitional step. Additionally, the network includes spatial pyramid pooling (SPP) (He et al., 2014) after the last convolutional layer. This has the effect of separating out the most important features and increasing the receptive field. The final feature map is divided into  $m \times m$  bins, following which maxpooling is applied to each bin. The resulting feature maps are concatenated and represent the output of the feature extractor.

CNNs naturally attain a pyramid-like structure with each layer as the image goes from high to low resolution. As we get deeper in a CNN, we lose the fine-grained details of the input, which usually makes it harder to detect small objects. As the resolution lowers, however, the filters learn ever more complex abstractions about the image, making the feature maps more semantically rich. Therefore, is it desirable to combine the feature maps from the higher resolution layer with the more semantically rich ones to facilitate detecting objects at multiple scales. This task falls to a feature aggregator and YOLOv4 uses the approach suggested by PANet (Liu et al., 2018). The feature maps are concatenated from both the top-down and the bottom-up path, ensuring the propagation of semantically rich localization information through to the final part of the network where the class probability and bounding box predictions are made. Each predicted bounding box consists of five elements: centre-x, centre-y, width, height, and

confidence. The  $(centre - x, centre - y)$  coordinates are relative to the dimensions of the predicted box; the width and height are relative to the whole image. The confidence score represents the likelihood that the cell contains the object as well as how confident the model is about its predictions.

YOLOv4 also updates the loss function to include Complete Intersection over Union (CloU) loss (Zheng et al., 2020) to train the network to effectively determine the direction in which to shift the weights to better match the labelled bounding boxes. Finally, the model also includes a raft of new additions: a new activation function, MISH (Misra, 2020) that provides smoother gradients; updates to the spatial attention module (Woo et al., 2018) and multi-input weighted residual connections (Tan et al., 2020) to better suit the architecture; and new data augmentation techniques, Mosaic and self-adversarial training.

The YOLO family of models are supported by Darknet (Redmon, 2013), a custom framework written in C and CUDA and designed for fast object detection. The Darknet implementation of YOLOv4 was trained with the MOA-II dataset via the Google Colaboratory with a hosted GPU runtime environment. The model was first trained with the default anchor boxes, before k-means clustering was used to discover anchor boxes that might prove better suited to finding tracklets in astronomical data. After several variations were tested, the best combination of anchor boxes was discovered by hand-engineering the various clusters. The learning rate of 0.001 was used along with a batch size of 32, with 8 subdivisions, 6000 max\_batches, and with steps set to 48,005,400.

## 5. Results

To evaluate the classification networks, we rely on the metrics derived from a confusion matrix. A confusion matrix breaks down the predictions made by a classifier into 4 outcomes: True Positive (TP), True Negative (TN), False Positive (FP), and False Negative (FN). The “positive” cases are where images were classified as containing tracklets and the “negative” cases are where no tracklets were detected. In this case, the ideal scenarios would be to have as few false negatives as possible, together with a manageable number of false positives. Here, the networks are evaluated based on their recall, F2 score (the weighted harmonic mean of the precision and recall), and the PR AUC (area under the precision–recall curve).

**Table 2**  
Evaluation metrics for the GB5-R5 test set taken at probability threshold 0.5.

	Recall	Precision	F2 Score	PR AUC	FN/FP
VGG16	83.88%	94.61%	85.83%	92.98%	56/16
VGG19	90.15%	89.09%	89.93%	95.09%	33/37
Inception	66.27%	67.27%	66.47%	75.80%	113/108
ResNet50	19.10%	54.24%	21.95%	38.96%	271/54
MOA-12	89.85%	86.49%	89.16%	93.28%	34/47
MOA-14	90.15%	83.20%	88.67%	93.26%	33/61
MOA-15	90.15%	90.15%	90.15%	94.35%	33/33
Hybrid A	90.15%	80.11%	87.94%	90.62%	33/75
Hybrid B	89.55%	80.00%	87.46%	90.54%	35/75

**Table 3**  
Evaluation metrics for the GB-ALL test set taken at probability threshold 0.5.

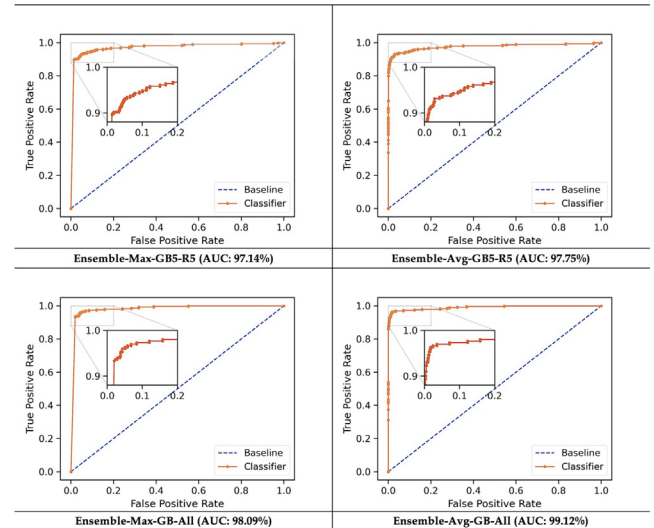
	Recall	Precision	F2 Score	PR AUC	FN/FP
VGG16	82.33%	93.21%	84.30%	90.69%	53/18
VGG19	90.33%	83.90%	88.97%	93.90%	29/52
Inception	67.33%	42.80%	60.41%	59.39%	98/270
ResNet50	37.67%	26.40%	34.71%	28.03%	187/315
MOA-12	94.00%	79.89%	90.79%	95.77%	18/71
MOA-14	94.33%	74.28%	89.50%	94.95%	17/98
MOA-15	94.00%	80.11%	90.85%	95.74%	18/70
Hybrid A	94.67%	72.63%	89.25%	92.04%	16/107
Hybrid B	94.67%	76.96%	90.50%	93.79%	16/85

The evaluation metrics for each classifier for the GB5-R5 test set are in Table 2 and for the GB-All (28-06-2013) test set are in Table 3. All of the metrics were taken at the 0.5 confidence threshold, with values over 0.5 indicating the presence of an asteroid tracklet in the image. Reviewing a combination of the PR AUC, F2 Score, and Recall, we can see the custom networks generalized well when making prediction about data from fields and chips they had never seen. Rather than selecting a single network from among these, all five custom classifiers were configured as an ensemble. Each network makes predictions about an input image and two approaches were trialed for selecting the winning prediction: averaging the predictions from all five classifiers (avg) or selecting the highest predicted value (max).

We see that while the max-ensemble results in a greater number of false positives, it improves the recall by four points (Table 4) to 94.33% and 97.67% for the GB5-R5 and the GB-All test set, respectively. The avg-ensemble has the advantage of having far fewer false positives and a judiciously selected prediction threshold could see false negatives further minimized for this

**Table 4**  
Evaluation metrics after configuring the five custom networks as an ensemble taken at probability threshold 0.5.

	Recall	Precision	F2 Score	PR AUC	FN/FP
Max-GB5-R5	94.33%	63.71%	86.06%	91.64%	19/180
Avg-GB5-R5	90.75%	91.29%	90.85%	95.46%	31/29
Max-GB-All	97.67%	50.78%	82.44%	91.52%	7/284
Avg-GB-All	96.33%	86.27%	94.14%	97.67%	14/27



**Fig. 11.** ROC curves for the CNN ensemble.

configuration. The ROC curves for the ensemble (Fig. 11) further illustrate this trade-off.

The performance of object detection models is typically quantified by the mean Average Precision or mAP. The average precision (AP) measures the trade-off between precision and recall and is calculated by integrating the area that falls under the precision–recall curve for each unique values of recall where the precision value decreases. The mAP is the mean of the AP across all object classes that the model can detect, which is identical to the AP for this dataset. Object detection models aim to have a high overlap between the predicted and the ground truth bounding boxes (Intersection of Union or IoU) and predictions are grouped as follows:


- True Positive: IoU > 0.5
- False Positive: IoU < 0.5 (or a duplicate)
- False Negative: box not detected or the IoU > 0.5 but the object is classified wrong

The average precision is calculated as: Calculated as:


$$AP = \sum (r_{(n+1)} - r_n) \tilde{p}(r_{(n+1)})$$

$$\tilde{p}(r_{(n+1)}) = \max_{\tilde{r} \geq r_{(n+1)}} (p(\tilde{r})),$$


where  $r$  is the recall value,  $n$  represents the locations where the precision decreases, and  $\tilde{p}(r_{(n+1)})$  is the maximum precision where the  $r$  value changes.




**Fig. 12.** Tracklets localization by YOLOv4. From left to right, these are: (538,147)2016 BT90 (18.7); (67,001) 1999 XN117 (19.2); (80,667) 2000 BA15 (18.8); (206,629) 2003 WT154 (19.0); and (281,890) 2010 OA74 (20.1). Note that these images have been inverted and brightened to improve visibility (original in Fig. A.17).



**Fig. 13.** The false positives reported by the classification networks fall into four broad categories: (a) ghosts (short for ghosts of noise past), are the most common false positives and are a result of the additive noise from creating the composite images; (b) streaks are either satellites or other near-Earth objects or cosmic rays; (c) chimera are optical artefacts potentially caused by over-saturated stars; and (d) loki objects are artefacts that move around erratically from one observation to the next. Note that these images have been inverted and brightened to improve visibility.



**Fig. 14.** The false negatives reported by the classification networks are mostly either faint objects or too much like streaks. It is likely that these sorts of tracklets were not well represented in the training set. Note that these images have been inverted and brightened to improve visibility.



**Fig. A.15.** A stack of all 51 observations from the night of 15-May-2008 reveals 5 asteroid tracklets clearly visible in a sub-region of the stack image. Numbered in order of appearance, these are: 1 - (103,842) 2000 DQ33 (19.5); 2 - (148,657) 2001 SX124 (19.6); 3 - (152,083) 2004 RH30 (19.9); 4 - (582,743) 2016 AT221 (20.4); 5 - (338,789) 2005 SZ154 (20.4). The line/streak on the top left is from a satellite.

YOLOv4 achieved an mAP of 90.96% with default anchor boxes and 90.95% with custom anchor boxes on the GB5-R5 data. Neither network could detect tracklets in the GB-All data, indicating that YOLOv4 will need to be trained with labelled data from other fields and chips before it can be used to make detections in them. Some of the detections made by YOLOv4 on the GB5-R5 dataset can be seen in Fig. 12.

The proposed model was applied towards discovering tracklets in the 543,595 images tiles from GB5-R5 that had no known asteroid tracklets. The ensemble could analyse 10,000 images in an hour and took three days to analyse all the images and proposed 50,227 candidate detections. Both version of YOLOv4 were then used for localizing tracklets in these candidate detections, with each version taking 1.5 days to make bounding box predictions. We are still analysing these candidate detections to determine if any new objects are present.


## 6. Discussion

Our classifier ensemble, trained with a relatively small amount of labelled data from just one chip in one survey field, was able to classify a set of images from unseen survey fields and chips, highlighting that it generalizes well. The success of the custom


CNN architectures with fewer training parameters is perhaps down to two aspects. First, the networks might be at the optimal depth for learning to identify the pattern of fuzzy blobs that represent a tracklet. Second, the small size of the input images perhaps leads to more overfitting in the larger models. Further, rather than adding more layers, the hybrid architectures lean into adding complexity with several  $1 \times 1$  convolutions. It is possible this caused the networks to learn representations that ultimately boosted the performance of the ensemble.

The false negatives (Fig. 14) reported by the ensemble include instances where the tracklets resemble satellite streaks or where the point sources are slightly further apart and obscured by noise. In some cases, parts of the tracklet are on other images that have been correctly identified as a candidate detection by the ensemble. As to the false positives, they are largely cases where noise and other spurious artefacts have been mistaken for tracklets, as indicated in Fig. 13. These also lead to false detections by the object detector. We believe that fine tuning the network models as more labelled data becomes available will aid in minimizing misclassifications such as these.

The YOLOv4 object detector proved to be up to the task of localizing tracklets in the GB5-R5 dataset but does not generalize to unseen data. The results from the GB5-R5 set are promising



**Fig. A.16.** Asteroid tracklets seen in  $128 \times 128$  sub-regions from subtracted stack images. In (b), the tracklets are localized with bounding boxes. Asteroids in the top row, left to right: (38,102) 1999 JM18 (18.4); (152,083) 2004 RH30 (19.9); (103,842) 2000 DQ33 (19.5); (283,261) 2011 FR142 (20.0); and (375,674) 2009 HD5 (19.8). Asteroids in the bottom row, from left to right: (48,617) 1995 HR2 (19.5); (148,657) 2001 SX124 (19.6); (97,948) 2000 QF124 (19.6); and (74,978) 1999 TY234 (18.7).



**Fig. A.17.** Tracklets localization by YOLOv4. From left to right, these are: (538,147) 2016 BT90 (18.7); (67,001) 1999 XN117 (19.2); (80,667) 2000 BA15 (18.8); (206,629) 2003 WT154 (19.0); and (281,890) 2010 OA74 (20.1).

and thus indicate that further training with data from the other survey fields and chips will be beneficial.

While both types of networks are currently limited by the high cadence observations they have been trained with, we are confident that, as the networks are retrained with data from other fields, they will learn to identify tracklets of slow objects or those with longer intervals between observations. Further, other surveys could use our networks, together with their pre-trained weights, as a starting point for discovering asteroids in their data. In this case, we recommend that the other surveys re-train/ fine-tune the architectures with their own data for optimal results.

## 7. Conclusion

We have shown that it is possible to train both CNN-based classifiers as well as the YOLOv4 object detector to find asteroid tracklets in the MOA difference images. The classifier ensemble proved resilient to discovering tracklets in unseen data and will be invaluable for extending the search for asteroids to the rest of the MOA-II archival data. The networks will be fine-tuned or re-trained as more labelled data is available and we will investigate automating the bounding boxes required for training YOLOv4. We will also investigate working with orbit-linking software such

as HelioLinC to determine the validity of the source clusters localized by YOLOv4.

While the classifiers performed well, there is potential for further improvement and development. Investigating effective denoising techniques for the stacked images or the difference images would lead to an immediate performance boost for both the classification and object detection networks. The classification network may benefit from having two inputs – perhaps the subtracted stack image along with the median stack image. This would provide the model with additional information it could use to distinguish between images with or without tracklets. Training the YOLO backbone feature extractor with the classification data first might also lead to better results. Additionally, since much of the salient information is contained in the first layer, a small CNN based on DenseNet could also potentially be successfully trained as a classifier.

Overall, we have presented an effective toolkit for finding asteroids tracklets in the archival data of ground-based telescopes. The code for our neural network models as well as the trained weights are available at <https://github.com/pcowan-astro/MOA-Asteroids>. Our methodology and network architectures can be used to discover and recover asteroids in other archival survey data as well as to strengthen the analysis pipeline for current and future surveys.

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Ian A. Bond reports financial support was provided by Royal Society of New Zealand Marsden Fund.

## Data availability

The authors do not have permission to share data.

## Acknowledgements

We thank the MOA collaboration for use of the MOA-II archival difference images. IAB acknowledges support from grant MAU1901 from the Royal Society of New Zealand - Marsden. This research has made use of data and/or services provided by the International Astronomical Union's Minor Planet Center.

## Appendix. Original images

See Figs. A.15–A.17.

## References

- Alard, C., 2000. Image subtraction using a space-varying kernel. *Astron. Astrophys. Suppl. Ser.* 144, 363–370.
- Alard, C., Lupton, R.H., 1998. A method for optimal image subtraction. *Astrophys. J.* 503, 325–331.
- Bochkovskiy, A., Wang, C.-Y., Liao, H.-Y.M., 2020. YOLOv4: Optimal speed and accuracy of object detection. doi:10.48550/arXiv.2004.10934, ArXiv arXiv:2004.10934.
- Bond, I., Abe, F., Dodd, R., Hearnshaw, J., Honda, M., Jugaku, J., Kilmartin, P., Marles, A., Masuda, K., Matsubara, Y., Muraki, Y., Nakamura, T., Nankivell, G., Noda, S., Noguchi, C., Ohnishi, K., Rattenbury, N., Reid, M., Saito, T., Sato, H., Sekiguchi, M., Skuljan, J., Sullivan, D., Sumi, T., Takeuti, M., Watase, Y., Wilkinson, S., Yamada, R., Yanagisawa, T., Yock, P., 2001. Real-time difference imaging analysis of MOA Galactic bulge observations during 2000. *Mon. Not. R. Astron. Soc.* 327, 868–880.
- Bramich, D., 2008. A new algorithm for difference image analysis. *Mon. Not. R. Astron. Soc.* 386, L77–L81.
- Cordwell, A.J., Rattenbury, N.J., Bannister, M.T., Cowan, P., Collaboration:, T.M., Abe, F., Barry, R., Bennett, D.P., Bhattacharya, A., Bond, I.A., Fujii, H., Fukui, A., Itow, Y., Silva, S.I., Hirao, Y., Kirikawa, R., Kondo, I., Koshimoto, N., Matsubara, Y., Matsumoto, S., Muraki, Y., Miyazaki, S., Okamura, A., Ranc, C., Satoh, Y., Sumi, T., Suzuki, D., Tristram, P.J., Toda, T., Yama, H., Yonehara, A., 2022. Asteroid Lightcurves from the MOA-II Survey: a pilot study. *Mon. Not. R. Astron. Soc.* 514 (2), 3098–3112. doi:10.1093/mnras/stac674.
- Denneau, L., Jedicke, R., Grav, T., Granvik, M., Kubica, J., Milani, A., Vereš, P., Wainscoat, R., Chang, D., Pierfederici, F., Kaiser, N., Chambers, K.C., Heasley, J.N., Magnier, E.A., Price, P.A., Myers, J., Kleyna, J., Hsieh, H., Farnocchia, D., Waters, C., Sweeney, W.H., Green, D., Bolin, B., Burgett, W.S., Morgan, J.S., Tonry, J.L., Hodapp, K.W., Chastel, S., Chesley, S., Fitzsimmons, A., Holman, M., Spahr, T., Tholen, D., Williams, G.V., Abe, S., Armstrong, J.D., Bressi, T.H., Holmes, R., Lister, T., McMillan, R.S., Micheli, M., Ryan, E.V., Ryan, W.H., Scotti, J.V., 2013. The pan-STARRS moving object processing system. *Publ. Astron. Soc. Pac.* 125, 357–395.
- Duev, D.A., Bolin, B.T., Graham, M.J., Kelley, M.S.P., Mahabal, A., Bellm, E.C., Coughlin, M.W., Dekany, R., Helou, G., Kulkarni, S.R., Masci, F.J., Prince, T.A., Riddle, R., Soumagnac, M.T., van der Walt, S.J., 2021. Tails: Chasing comets with the Zwicky transient facility and deep learning. *Astron. J.* 161, 218.
- Duev, D.A., Mahabal, A., Ye, Q., Tirumala, K., Belicki, J., Dekany, R., Frederick, S., Graham, M.J., Laher, R.R., Masci, F.J., Prince, T.A., Riddle, R., Rosnet, P., Soumagnac, M.T., 2019. DeepStreaks: Identifying fast-moving objects in the Zwicky Transient Facility data with deep learning. *Mon. Not. R. Astron. Soc.* 4158–4165.
- Glorot, X., Bordes, A., Bengio, Y., 2011. Deep sparse rectifier neural networks. *AISTATS*.
- Gould, A., Yee, J.C., 2013. Microlens surveys are a powerful probe of asteroids. *Agron. J.* 767.
- He, K., Zhang, X., Ren, S., Sun, J., 2014. Spatial Pyramid Pooling in Deep Convolutional Networks for Visual Recognition. *Lecture Notes in Comput. Sci.* 8691 LNCS, 346–361.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. *IEEE Conf. Comput. Vis. Pattern Recognit.* 770–778.
- Holman, M., Payne, M., Blankley, P., Janssen, R., Kuindersma, S., 2018. HeliLinC: A novel approach to the minor planet linking problem. *Astron. J.* 156, 135. doi:10.3847/1538-3881/aad69a.
- Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q., 2016. Densely connected convolutional networks. In: *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*. 2017-January, Institute of Electrical and Electronics Engineers Inc., pp. 2261–2269.
- Kim, S.-L., Lee, C.-U., Park, B.-G., Kim, D.-J., Cha, S.-M., Lee, Y., Han, C., Chun, M.-Y., Yuk, I., 2016. KMTNET: A network of 1.6 m wide-field optical telescopes installed at three southern observatories. *J. Korean Astron. Soc.* 49 (1), 37–44. doi:10.5303/JKAS.2016.49.1.037.
- Kingma, D.P., Welling, M., 2014. Auto-encoding variational bayes. In: *2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings*. International Conference on Learning Representations, ICLR.
- Kruk, S., García Martín, P., Popescu, M., Merín, B., Mahlke, M., Carry, B., Thomson, R., Karadağ, S., Durán, J., Racero, E., Giordano, F., Baines, D., de Marchi, G., Laureijs, R., 2022. Hubble asteroid hunter. *Astron. Astrophys.* 661, A85.
- Kubica, J., Denneau Jr, L., Moore, A., Jedicke, Robert, Connolly, A., 2007. Efficient algorithms for large-scale asteroid discovery. In: *Astronomical Data Analysis Software and Systems XVI ASP Conference Series*. vol. 376, pp. 395–404.
- Lieu, M., Conversi, L., Altieri, B., Carry, B., 2018. Detecting solar system objects with convolutional neural networks. *Mon. Not. R. Astron. Soc.* 485, 5831–5842.
- Liu, S., Qi, L., Qin, H., Shi, J., Jia, J., 2018. Path aggregation network for instance segmentation. In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. pp. 8759–8768.
- Misra, D., 2020. Mish: A Self Regularized Non-Monotonic Activation Function. doi:10.48550/arXiv.1908.08681, ArXiv arXiv:1908.08681.
- Nair, V., Hinton, G.E., 2010. Rectified linear units improve restricted Boltzmann machines. *ICML*.
- Newman, W.M., Sproull, R.F., 1973. Cohen-sutherland algorithm computer graphics, Internatio ed. McGraw-Hill Education, p. 124, 252.
- Rabendran, A.C., Denneau, L., 2021. A two-stage deep learning detection classifier for the ATLAS asteroid survey. *Publ. Astron. Soc. Pac.* 133.
- Rabinowitz, D.L., 1991. Detection of earth-approaching asteroids in near real time. *Astron. J.* 101, 1518.
- Redmon, J., 2013. Darknet: Open source neural networks in C. URL <https://pjreddie.com/darknet/>.
- Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You only look once: Unified, real-time object detection. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pp. 779–788.
- Redmon, J., Farhadi, A., 2017. YOLO9000: Better, faster, stronger. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pp. 6517–6525.
- Redmon, J., Farhadi, A., 2018. YOLOv3: An Incremental Improvement. doi:10.48550/arXiv.1804.02767, ArXiv arXiv:1804.02767.
- Sako, T., Sekiguchi, T., Sasaki, M., Okajima, K., Abe, F., Bond, I.A., Hearnshaw, J.B., Itow, Y., Kamiya, K., Kilmartin, P.M., Masuda, K., Matsubara, Y., Muraki, Y., Rattenbury, N.J., Sullivan, D.J., Sumi, T., Tristram, P., Yanagisawa, T., Yock, P.C., 2008. MOA-cam3: A wide-field mosaic CCD camera for a gravitational microlensing survey in New Zealand. *Exp. Astron.* 22, 51–66.
- Simonyan, K., Zisserman, A., 2015. Very deep convolutional networks for large-scale image recognition. In: *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*.
- Sumi, T., Abe, F., Bond, I., Dodd, R., Hearnshaw, J., Honda, M., Honma, M., Kanya, Y., Kilmartin, P., Masuda, K., Matsubara, Y., Muraki, Y., Nakamura, T., Nishi, R., Noda, S., Ohnishi, K., Petterson, O., Rattenbury, N., Reid, M., Saito, T., Saito, Y., Sato, H., Sekiguchi, M., Skuljan, J., Sullivan, D., Takeuti, M., Tristram, P., Wilkinson, S., Yanagisawa, T., Yock, P., 2003. Microlensing optical depth toward the galactic bulge from microlensing observations in astrophysics group observations during 2000 with difference image analysis. *Astrophys. J.* 591, 204–227.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z., 2015. Rethinking the inception architecture for computer vision. In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. vol. 2016-December, pp. 2818–2826.
- Tan, M., Pang, R., Le, Q.V., 2020. EfficientDet: Scalable and efficient object detection. In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. pp. 10778–10787.
- Tomaney, A.B., Crotts, A.P., 1996. Expanding the realm of microlensing surveys with difference image photometry. *Astron. J.* 112, 2872.
- Udalski, A., Szymanski, M., Kaluzny, J., Kubiak, M., Krzemiński, W., Mateo, M., Preston, G.W., Paczynski, B., 1993. The optical gravitational lensing experiment. Discovery of the first candidate microlensing event in the direction of the galactic bulge. *Acta Astron.* 43, 289–294.

- Wang, C.Y., Mark Liao, H.Y., Wu, Y.H., Chen, P.Y., Hsieh, J.W., Yeh, I.H., 2020. CSPNet: A new backbone that can enhance learning capability of CNN. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops. vol. 2020-June, pp. 1571–1580.
- Woo, S., Park, J., Lee, J.Y., Kweon, I.S., 2018. CBAM: Convolutional Block Attention Module. Lecture Notes in Comput. Sci. 11211 LNCS, 3–19.
- Zheng, Z., Wang, P., Ren, D., Liu, W., Ye, R., Hu, Q., Zuo, W., 2020. Enhancing geometric factors in model learning and inference for object detection and instance segmentation. IEEE Trans. Cybern..
- Zoghbi, S., Cicco, M.D., Ordonez, A.J., Stapper, A.P., Collison, J., Gural, P.S., Ganju, S., Galache, J.-I., Jenniskens, P., 2017. Searching for Long-Period Comets with Deep Learning Tools. In: Workshop on Deep Learning for Physical Sciences.