

A Journey to the Edge of the Solar System with an AI navigator

by

Aram Lee

Bachelor of Science, Virginia Polytechnic Institute and State University, 2018

A Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of

MASTER OF SCIENCE

in the Department of Physics and Astronomy

© Aram Lee, 2023

University of Victoria

All rights reserved. This thesis may not be reproduced in whole or in part, by photocopying or other means, without the permission of the author.

A Journey to the Edge of the Solar System with an AI navigator

by

Aram Lee

Bachelor of Science, Virginia Polytechnic Institute and State University, 2018

Supervisory Committee

Dr. JJ Kavelaars, Co-supervisor
(Department of Physics and Astronomy)

Dr. Kim Venn, Co-supervisor
(Department of Physics and Astronomy)

Dr. Hossen Teimoorinia, Committee Member
(Department of Physics and Astronomy)

ABSTRACT

I present a deep learning method of searching for solar system objects (SSOs) in wide-field survey imaging data including trans-Neptunian objects (TNOs).

Artificially generated sources are added to mosaic images taken with the Canada-France-Hawaii telescope (CFHT) MegaCam instrument to create the convolutional neural network (CNN) training set. The CFHT MegaCam data images are a time series of observations, and the location of the artificial SSO changes between images, in a way that is consistent with a heliocentric Keplerian orbit. The imaging characteristics of the artificial sources were found to be highly similar to those of real SSOs, with rates of sky motion consistent with TNOs.

My deep learning approach is based on the detection of moving sources within 64×64 -pixel sub-image pairs extracted from the time series of large-format mosaic astronomical imaging data. Each image pair extracted from the training images has been labelled with the presence or absence of a moving source, along with the source location and brightness measured in magnitudes. The labelled sub-images were fed into ImageNet algorithms to train classification models and regression models separately. The algorithm assigns a model-dependent probability that a particular sub-image contains an SSO. The probability threshold required to assert that an SSO has been detected is set based on the evaluation of retrieval and precision of the model and the requirements of the experiment. This thesis evaluates the capabilities of the range of deep learning models and determines which one is most effective in the detection of artificial SSOs.

The MobileNet model was selected as the most efficient for this problem space. A trained classification model derived from the MobileNet model retrieved 91% of sub-images with a moving source with a 90% precision on test data sets. A separate regression model then predicted the location of the moving source with a mean absolute error of ± 1.5 pixels for sources with $SNR > 17$ ($m_r < 23$ in my data set). Although the retrieval rate is high, due to the scarcity of real SSOs in imaging data, the precision achieved (90% of false positives rejected) results in a substantial number of false positives. Further data processing on the candidate list is required to improve the purity of the result.

To improve sample purity, I investigated two post-processing approaches:

- With the classification-filtered sub-images and their regression-measured locations in sky coordinates, each detected source was grouped with nearby detected sources as SSOs exhibit nearly linear sky motion for the duration of the observed time series. Any group of linear source tracks, detected in at least 1/3rd of the images, was considered a candidate detection. This approach achieves an effective detection limit (more than 50% of artificial sources in the data are detected) at $\text{SNR}=7.2$, and the source purity of the sample was greater than 99% in this case. However, the required combinatorics of this approach ($N \times N$ comparison) make it computationally slow, and the high SNR required for detection resulted in very few ‘real’ candidates being proposed.
- I also investigate a ‘scoring’ approach for candidate selection. My CNN classification model output is a model-dependent probability that a particular sub-image contains a moving source. Each sub-image was given a score derived by scaling the classification model probability assigned to that sub-image. A sub-image was then determined to hold a candidate object if its score exceeded a given threshold (determined by the desired purity of the sample). With this approach, I achieved an effective detection limit (50% of artificial sources in the data are detected) at $\text{SNR}=3.4$ and discovered a number of real SSOs within the test data set. Visual inspection of 1800 scoring-based candidates revealed approximately 200 visibly bright real (not from the artificial source list) SSO candidates.

I tested trained models on test sets from different sky regions and found that our models did not learn from the backgrounds or shapes of TNOs, but rather detected the motion of TNOs. I found that deep-learning object detection algorithms can aid in the discovery of TNOs and SSOs. When combined with a scoring approach, my algorithm provides a capability that is similar to that achieved with more classical approaches without making assumptions of motion rates of the SSOs and without requiring any substantive data engineering. The CNN approach to SSO detection is very promising and should be pursued in the development of future SSO discovery software pipelines.

Contents

Supervisory Committee	ii
Abstract	iii
Table of Contents	v
List of Tables	vii
List of Figures	ix
Acknowledgements	xv
Dedication	xvi
1 Introduction	1
1.1 Definition of TNOs	1
1.2 Diversity in TNO population	2
1.3 A Reason to Find More TNOs	4
1.4 How TNOs are discovered	7
1.5 Challenges in discovering TNOs	9
1.6 Artificial intelligence, machine learning and deep learning	11
1.7 Astronomy and Image recognition deep learning algorithm	12
2 A Trans-Neptunian Object and Solar System Object Discovery Method with Convolutional Neural Networks for Future Surveys	15
2.1 Contributions	15
2.2 Introduction	16
2.3 Image Recognition Deep Learning Techniques	17
2.3.1 Tested Architectures	17

2.3.2	Classification and Regression	20
2.3.3	Model Implementation	21
2.4	Data and Pre-processing	22
2.4.1	Original Images	22
2.4.2	Data Calibration	23
2.4.3	Synthetic Trans-Neptunian Objects	24
2.4.4	Building Image Subsets	25
2.4.5	Labels	28
2.5	Model Selection	29
2.5.1	Architecture comparison	29
2.5.2	Training Data	32
2.6	Implementation	37
2.6.1	Approach 1: Linear Fitting Approach	37
2.6.2	Approach 2: Scoring Approach	40
2.7	Results and Discussion	41
2.7.1	Results with the Linear Fitting Approach	42
2.7.2	High Precision Detection with the Probability Cumulative Approach	45
2.7.3	Peak Detection Efficacy	47
2.7.4	Discovery of real SSOs	48
2.8	Summary and Conclusion	50
3	Additional concepts for a CNN-based discovery system.	52
3.1	Discovery of a TNO with just one image: Why it works and why it should not	52
3.2	Ideas to improve the discovery methods	54
4	Conclusion	56
	Bibliography	59

List of Tables

Table 2.1	Architecture List with performance and the number of parameters on ImageNet validation set. The depth is the number of convolutional layers and fully connected layers.	18
Table 2.2	Examples of CFHT Megacam datasets names used in this research	27
Table 2.3	Classification performance of tested networks, based on the validation set of C051020M25. Larger is better for recall, precision, and F1, and the metrics are explained in section 2.5.1.	30
Table 2.4	Regression performance on C051020M25 validation set. ‘pos.’ and ‘mag.’ in front of MAE and RMSE refer to positional (pixel) and magnitude uncertainties respectively. Lower values indicate better performance.	31
Table 2.5	Classification performance on validation sets (same chip) and test sets (different chips) at p threshold=0.5. The models are trained on M25 data sets. The models were validated/tested on a balanced M23 data set, which is a subset of the M25 validation/test sets, for a better presentation of their effectiveness on brighter objects.	33
Table 2.6	Classification performance of the combined data set models on validation sets (same chip) and test sets (different chips) at p threshold=0.5. The table displays the performance of the models trained on an M25 data set when validated/tested on a balanced M23 data set. The suffixes “R” and “M” after the model names denote ResNet and MobileNet.	37

Table 2.7 Regression performance of the combined data set model on the validation set (same CCDs) and test sets (different CCDs) at threshold = 0.5. The table displays the performance of the models trained on an M25 data set when validated/tested on a balanced M23 data set. The suffixes “R” and “M” after the model names denote ResNet and MobileNet.	37
--	----

List of Figures

- Figure 1.1 Number of TNOs listed to the Minor Planet Center each year (left) and the cumulative number (right), as of October 2023. . . . 9
- Figure 1.2 A diagram to explain the convolutional operation of the CNN. The filter/kernel is a 3D/2D array of values to extract features from images, and values in the array change throughout learning. Training CNN is a process of finding the best weights and bias (omitted in this diagram) to solve a problem. The input is an array, which can be an image or an output layer of the previous layer. The output array is calculated by element-wise multiplication of the input and filter/kernel arrays, applied regionally. For example, only the values highlighted in bold are calculated to generate the output value 6. 13
- Figure 2.1 (Left) An example of the images from CFHT MegaCam before cutting. This is the upper left part of a single CCD. Around 200 artificial TNOs are present in this image; most of them are small and dim, making them hard to find on this scale. (Right) Measured positions by the regression model of all 44 observations at this location. Only positions that passed the classification are shown here. The short, almost horizontal lines are mostly the signature artificial TNOs. The vertical line from the top to the middle is the collection of incorrect measurements of moving objects caused by the noise from very bright stars. Large, empty, circular spaces are locations where a bright star has obscured detections. On the bottom right, a real SSO candidate is detected as a longer inclined streak. . . . 22

Figure 2.2 This figure shows sky coordinates of centres where images are cut into 64×64 pixel sub-images. The left plot has 500 randomized cutting positions, and the positions were different for each FITS file pair, and the right plot has evenly spaced cutting positions with 63-pixel gaps. The left one is to have a randomized training/validation set with a less volume, and the right one is to retrieve the most planted and non-planted TNOs by searching the whole area. 26

Figure 2.3 This figures shows 6 pairs of sub-images. Examples of sub-images used as our inputs. The orange circle marks the location of the synthetic TNOs. From the left upper pair to the lower right pair, 6 pairs of positive-negative sub-images are shown. For the magnitude 22, 23, and 24 cases one can see the bright source in the positive sub-image that is absent in the negative sub-image taken from a different epoch. The magnitude 25, 26 and 27 sources are progressively more difficult to distinguish. The artificial sources range from magnitude 22 to 27. 28

Figure 2.4 Cumulative plots of detection efficiency of the MobileNet model trained using different input data sets. Upper: Trained on data with sources brighter than m=23. centre: m=25. Lower: m=27. On the left side, each panel shows classification probability (p) versus fraction (f) of sources with a moving source correctly classified. The vertical blue/navy/black lines show various thresholds that can be used to select images that may contain a moving source. Thresholds of 0.5, 0.9, and 0.99 are shown, and the middle panels show the model behavior near the 0.9 and 0.99 thresholds. The panels in the right column show the fraction of sources correctly classified as a function of the brightness of the source, also known as the completeness curve. 34

- Figure 2.5 Cumulative plots of the model probability (p) versus the fraction of images classified as containing a TNO (g). All image pairs used for this plot are negatives, i.e. do not contain any artificial moving sources. Thus, g is the number of false positives compared to the total number of negative sub-images. There is very little difference between the M23 and M25 false-positive rates. The M27 model, however, returns false positives at double the rate of the M23/M25 models when $p=0.5$. 35
- Figure 2.6 ROC curves from testing the C051020M25 classifier on image pairs. The x-axis is the false positive rate that shows how frequent false alarms are compared to the total number of negatives (Less is better). The y-axis is the true positive rate that shows how many positives are correctly classified as positives compared to the total number of positives (More is better). These ROC curves show the general performance of the classifier when applied with various p thresholds. The more the curve is skewed to the left-upper area, the better the classifier performs. Area under curve (AUC) indicates the performance of the classification model (More is better). The ROC curves have AUCs of 0.87, 0.82, and 0.74 depending on the maximum magnitude limits of 23, 24, and 25. The diamond markers indicate points at $p=0.99$ classification threshold. . . . 35
- Figure 2.7 Blue dots are regression-predicted positions of a group of sources that the classification model gave $p>0.5$ probability of being a moving source. Yellow crosses are the predicted locations a moving source would be at, in each exposure, based on a maximum likelihood-based linear fit to the regressor-based locations. The linear model provides an excellent match to this group of sources which are then declared as moving object candidates. 39

- Figure 2.8 Analysis of linear fitting approach on the C08132134 test set. (Left) Number of false positive tracks as a function of the minimum length of the track (n) to accept a track as a detection candidate. The vertical lines and end points are at $n=3$, $n=6$, $n=11$, and $n=22$. Setting the minimum track length requirement to these values would result in 14435, 1514, 387, and 98 candidates to vet, respectively. (Right) A completeness graph showing artificial TNO retrieval rates as a function of magnitude for various track length. “Found” objects did not go through the human inspection, therefore the actual retrieval rate will be slightly lower. Figure 2.4 is based on each single image pair classification result while the completeness shown here is for the final candidates after classification, and linear fitting. 42
- Figure 2.9 Analysis of scoring approach on the C08132134 test set (Left) Number of false positive sub-image series as a function of the minimum score of the sub-image series (n) to accept a sub-images series as a detection candidate. The vertical lines are at $s=80$ and $s=1000$. Setting the minimum score requirement to these values resulted in 110 and 1 false-positive candidates. (Right) A completeness graph showing artificial TNO retrieval rates as a function of magnitude for two kinds of score requirements and the manual inspection requirement. Figure 2.4 is based on each single image pair classification result while this completeness graph is to show to the final number of TNOs found after post-processing. 45

Figure 2.10 Sub-images containing artificial sources that were not detected by our process. The orange circles indicate the position at which the artificial source was added to the image. (Left) A bright stationary object (a star) in the background is obscuring the moving object. (Right) Bad columns on the detector prevent the detection of the moving object. For these objects, the classification model assigns a low probability to the presence of a moving source. For those that pass the classification step, the regression model often fails to predict the correct position of the moving source, and the candidate is rejected during the tracking process. 47

Figure 2.11 A sub-image series accepted as being a detected solar system object. The images in the grid are arranged in rows and columns, at the same sky coordinates, depicting the progression of time. The time interval between each image is 248 seconds. The images are arranged in a sequential order from left to right, and from top to bottom. This non-planted candidate returned the highest score from our method. Pixel scale is 0.185 arcsecond per pixel. Based on the rate of motion of ~ 3 arcsecond per hour, this object is likely to be a Centaur or a TNO. 48

Figure 2.12 Sub-image series examples accepted as being detected SSOs. The images are arranged in sequential order from left to right. The first row: A dim oval-shaped light source moves from the center-left to the center-right almost horizontally. The moving object is easier to spot when images are made into an animation. The second row: A bright oval-shaped light source moves from the center-left to the lower right. Based on their rate of motion, these objects are likely to be asteroids. 49

Figure 3.1 The graphs show the relationship between the performance of models and motion rates of TNOs. The left graph represents the case where the model is trained on each image separately, and the right graph represents the case where the model is trained on pairs of images that depict the movement of TNOs. As the rate of motion increases, the recall of the dual-channel model also increases. Unlike single-channel training, dual-channel training enables the model to recognize the movement of the source. . .

ACKNOWLEDGEMENTS

For the land I stayed during the study:

We acknowledge and respect the Lək̓ʷəŋən (Songhees and Esquimalt) Peoples on whose territory the university stands, and the Lək̓ʷəŋən and WSÁNEĆ Peoples whose historical relationships with the land continue to this day.

For the data I used for this thesis:

Based on observations obtained with MegaPrime/MegaCam, a joint project of CFHT and CEA/DAPNIA, at the Canada-France-Hawaii Telescope (CFHT) which is operated by the National Research Council (NRC) of Canada, the Institut National des Science de l'Univers of the Centre National de la Recherche Scientifique (CNRS) of France, and the University of Hawaii. The observations at the Canada-France-Hawaii Telescope were performed with care and respect from the summit of Maunakea which is a significant cultural and historic site.

I acknowledge the use of the Canadian Advanced Network for Astronomy Research (CANFAR) Science Platform. My work used the facilities of the Canadian Astronomy Data Center, operated by the National Research Council of Canada with the support of the Canadian Space Agency, and CANFAR, a consortium that serves the data-intensive storage, access, and processing needs of university groups and centers engaged in astronomy research (Gaudet et al 2010).

DEDICATION

I would like to thank *my dear friends, Soyoun Kim and MK Kim*, for their steadfast support throughout both the happy and tough times of the past 10 years. You are the reason I could keep myself going in challenging moments. I am deeply grateful that we can talk to each other every day, and we will see each other again.

I express my gratitude to *my academic supervisor Professor JJ Kavelaars*, for the mentoring, support, encouragement, and profound patience.

Chapter 1

Introduction

We naturally love unknown objects. The curiosity is coded in our brain with increased dopaminergic modulation [1]. Astronomical Objects among the most unknown lie on the edge of our solar system. By observing them through telescopes and approaching them with space probes, we grow to understand them a little better. As we understand them a little better, our interest in the subject becomes even greater until we are completely satisfied [2]. Of particular interest are trans-Neptunian Objects or TNOs. Until we are completely satisfied with TNOs, we will continue to study them.

In this thesis, I introduce observational studies in TNOs and present a novel discovery technique using deep learning, and I conclude the thesis with some discussion of future work.

1.1 Definition of TNOs

Trans-Neptunian Objects (TNOs) are planetary objects in our solar system whose average distance is beyond the orbit of Neptune. More specifically, they are all planetary objects with a semi-major axis greater than 30.1 au (beyond Neptune) but less than 2000 au (beyond which lies the Oort Cloud). (134340) Pluto was discovered in 1930 and is the first TNO known. Initially, Pluto was thought to be the outermost planet of the Solar System. However, following the discovery of Pluto, astronomers proposed the existence of other planetary bodies in the outermost region. The first two astronomers to claim the existence of such objects quantitatively were Kenneth Essex Edgeworth and Gerard Peter Kuiper [3]. Therefore, TNOs are also called Edgeworth-Kuiper belt objects or just Kuiper belt

objects.

1.2 Diversity in TNO population

TNOs are the history books of our solar system. Some TNOs are in mint condition, having remained almost unchanged since their formation in the early solar system, and provide us with the opportunity to study the conditions of the solar proto-planetary disk. Other TNOs are equally important as they have gone through physical and dynamic evolution and thus provide constraints on the evolution of the solar system. In the following paragraphs, I will describe the diversity of TNOs and what they reveal about the history of the solar system.

TNOs are diverse in their *orbits*. For example, (528219) 2008 KV₄₂ has a highly inclined orbit, $i = 104^\circ$. It is a rare retrograde orbit and goes far beyond the solar system plane [4]. For another example, (241097) 2007 DU₁₁₂ reaches deep into the solar system. Its perihelion is 9.0 au, similar to that of Saturn [5]. On the other side of the spectrum, there are extremely circular and non-inclined TNOs. 2015 GE₅₆ is one of them, having $i = 0.036^\circ$ and $e = 0.049$. Each of these objects has arrived in its orbit via some dynamical process, and the ensemble of these reveals the dynamical history of the giant planets of the solar system.

TNOs are diverse in their *sizes*. The largest ones, such as Pluto, Eris, Haumea, and Makemake, are over 1000 km in diameter. The medium-sized TNOs are tens to hundreds of kilometres in diameter. For example, (15810) Arawn is approximately 133 kilometres in size. The smallest TNO found is around 5 km in diameter (2016 LF₉₀). There are smaller TNOs, but they are too faint to be detected [6]. These sizes can be measured with direct imaging, thermal observation (building a thermal model for both blackbody radiation and reflected sunlight) and, via occultation events (measuring the time that background stars are concealed beneath TNOs as the Earth passes through the shadow.) Determining the distribution of sizes of TNOs is an important measure of the outcome of their formation process. It can reveal the physics of the formation of not just TNOs but of the planetesimals, which are building blocks of the larger planets.

TNOs are diverse in their *albedo*¹. Inner solar system objects, such as the main belt asteroids and Centaurs, often have lower albedos than TNOs as the asteroids

¹Albedo is the ratio of light that impinges onto a surface to that which is reflected or scattered from the surface.

and Centaurs come closer to the Sun, tend to be small, and thus their surfaces do not retain volatile molecules, such as H_2O ice, which are more reflective than C. The TNO with the highest known albedo is Eris: Eris has the visible geometric albedo of $0.96_{-0.04}^{+0.09}$ [7]. While most TNOs are on the darker side of the spectrum, some are particularly dark. 2003 QX₁₁₁ is an example: 2003 QX₁₁₁ has the geometric albedo of $0.018_{-0.009}^{+0.017}$ or $0.033_{-0.025}^{+0.164}$ [8, 9]. This drastic difference between the bright object and the dark object makes the study of TNOs interesting. The surface albedos and the physical materials and processes that lead to them can reveal the chemical nature and evolution of these objects and, thus, the original composition of the solar proto-solar disk. The study of TNO albedos is in its infancy.

Moreover, the range of known albedos is why their sizes are often poorly estimated. When the TNO Sedna was first discovered, the object was assumed to have a diameter of ~ 1800 km [10]. At the time, the albedo of the Kuiper belt objects was thought to be around 0.04. Later, following thermal observation with the Herschel Space Observatory, Sedna's geometric albedo was determined to be 0.32 ± 0.06 , determining the estimated diameter to be 995 ± 80 km. A similar initial overestimation occurred for the large TNOs Pluto and Eris. Hence, constraining the albedo is also important to identify the size of the TNO, which, as noted above, is an important measurable of planetesimal formation processes.

TNOs are diverse in their *colours*² One of the reddest objects is Gonggong: (B-V) = 1.38 ± 0.03 and (V-R) = 0.86 ± 0.02 [11]. This colour is comparable to the (B-V) of the Mars [12]. In contrast, Haumean family members are known to be the bluest objects. As they share their past, the composition and the orbits are similar, and thus, they are almost equally blue. B-V magnitude of the Haumea is 0.64 ± 0.01 and V-R magnitude is 0.33 ± 0.01 [13].

TNOs are diverse in their *surfaces* and *atmospheres*. As we can know from the diversity of albedos and colours, the surfaces of some TNOs are covered with bright (fresh) ices, but others are covered with older, darkened materials. Comparing two objects that have been observed at close approach, Pluto is mainly covered with N_2 and CH_4 ice and has an atmosphere caused by the sublimation of the surface ices [14, 15]. On the other hand, Arrokoth is mainly covered with CH_3OH ice and, due to its smaller size, poses no atmosphere due to its low surface gravity [16]. The surface

²Here colour is a term used to refer to the log-ratio of flux from the object measured in two separate band passes, normally designated by single letters and noted as 'X'-'Y'. For example, the B bandpass is near 400 nm while the V bandpass is near 500 nm, and the colour (B-V) is the log ratio of the fluxes observed through those two band passes $(B-V) \propto \log(\frac{f_V}{f_B})$.

chemistry hints at the physical composition of these objects, providing further insight into the chemistry of the solar proto-planetary disk.

TNOs are diverse in their *multiplicity*. For example, 2001 QW₃₂₂ is a binary system composed of two approximately same-sized objects[17]. This binary TNO has a large separation ($a \approx 114,000\text{km}$) and period ($P \approx 27$ years). For some objects with $\sim 100\text{km}$ diameter, they orbit each other at a large distance, which might jeopardize, due to possible close encounters with other bodies, the long-term survival of the binary system. There is another extreme case, which is (47171) Lempo, a ternary system. This wonderful system has two larger components at a short orbiting distance ($a \approx 867\text{km}$) and a smaller component orbiting the inner binary from afar ($a \approx 7411\text{km}$). The far component orbits with a period of 50 days. It appears as an observer, slowly orbiting around the two central figure skaters performing on the stage, orbiting each other very quickly with a period of just 1.9 days [18]. From a singular system to a ternary system, the multiplicities of TNOs are various. The complex processes required for the formation of such systems and their fragility against disruption provide further insight into the initial conditions and long-term evolution of the outer solar system.

TNOs are diverse in their *shape*. A good example is Haumea, which is rotating rapidly and made itself to have a flattened shape. On the other hand, Arrokoth is a contact binary. The New Horizons close approach enabled us to study its planetesimals directly. Likewise, smaller TNOs such as Arrokoth are likely to have irregular shapes as the spherical shape would form due to gravitational compression and requires at least a radius of ~ 200 km [19]. The diversity of shape informs of the internal properties of the TNOs and, again, the formation processes of these planetesimals.

The collection book of TNOs has thousands of stamps, which are all unique in a way. Studying TNOs is rewarding because they are all different and tell us stories of the solar system from various viewpoints. The diversity of TNOs and the rich information they provide about the formation of the solar system make the discovery of TNOs a valuable scientific endeavour.

1.3 A Reason to Find More TNOs

In the previous section, I discussed the diversities of TNOs. But aside from their own beauty, some problems can be solved by discovering more TNOs. I introduce

examples of problems that can be solved or better understood with the discoveries of more TNOs.

Where are the extreme TNOs from? The mysterious rabbit hole problem TNO researchers get into is the origin of extreme TNOs or eTNOs. The eTNOs are TNOs with extremely large orbital semi-major axis ($a > 150$ au) and orbital pericentre beyond Neptune, $q > 30$ au. At the time of the first discovery of such an object, namely 2000 CR₁₀₅, astronomers speculated that they would find more objects on such ‘extreme’ orbits. Objects on these orbits cannot be explained via simple scattering from a stable-orbit Neptune [20]. Indeed, more and more eTNOs have been found, starting with Sedna and 2003 SS₄₂₂. Following those, a new eTNO was discovered in 2004, and another was found in 2005 (474640 Alicanto [21], and 2005 RH52 [22]). And after a 5-year time gap, dozens more have been found. By October 2023, 30 eTNOs have been listed in the database of all known TNOs maintained by the Minor Planet Center. This group of should-not-be-there objects intrigued astronomers in the field.

At the point of discovery of 2000 CR₁₀₅, astronomers who observed the object already had a theory about it. That was, a migrating Neptune (Neptune moving from one location of the solar system to another) scattered, via close encounters, the TNOs out, and then Neptune migrated to a location that removed the object from the gravitational domain of Neptune [23]. Additional authors also suspected the origin of 2000 CR₁₀₅ could be because of the primordial Neptune, just like other scattered objects [e.g. 24]. However, the hypothesis of scattering interactions with a migrating Neptune is a possible but unlikely explanation for the orbit of eTNO 2015 BP₅₁₉ [25].

Another possible origin of the eTNOs is via gravitational scattering during a stellar flyby. A stellar flyby is the close approach of one star to another star. In planetary astronomy, the stellar flyby means that another star came close enough to the Sun that it will gravitationally perturb the planetary system. There are scenarios where the scattering of eTNOs could result from the stellar flyby [26, 27]. Stellar flybys that affect the solar system do not frequently occur for a field star like the Sun. The most imminent flyby of a stellar mass to the Sun is expected to be at a distance of $10,563 \pm 566$ au with a 90% probability, 1.29 ± 0.04 Myr from now. The effects on the orbit of classical TNOs from such an encounter will be negligible due to the large distance of the encounter [28]. To generate eTNOs, the flyby must be much closer to the solar system. If the Sun formed with a loose cluster of stars that has subsequently dispersed, a close stellar flyby of a few 100 au would have been more likely early in the

history of the planetary system. Some models show that it is possible to create eTNO orbits via early stellar encounters at a distance of 90-160 au [27]. Both corotating and counterrotating stellar flybys could modify the orbits of classical TNOs to be in the current orbits of eTNOs [27]. Surprisingly, such encounters provide a significant likelihood that they would also result in the capture of eTNOs away from the passing star, and the chance was significant [27]. Even after decades of study, the stellar flyby hypothesis continues to be examined as a possible explanation for eTNOs.

As another interesting hypothesis to explain eTNOs, the solar system could have had an indigenous small mass planet (of the order of an Earth mass) for a few 100 million years early in its history. Some TNOs could have become eTNOs via the gravitational vortex of the small mass planet migrating out of the solar system [29, 30]. Indeed, terrestrial mass rogue planets were observationally discovered in 2023 via microlensing: the free-floating planet MOA-9y-5919L has a very small mass ($0.75_{-0.46}^{+1.23} M_{Earth}$) [31]. If such a small planet existed in the early outer solar system for ~ 150 Myrs and then scattered out, the rogue planet would have raised the perihelia and/or inclination of objects that had previously scattered off Neptune, resulting in orbits like those of the eTNOs [29].

And we have the hypothesis of a massive planet still resident in the distant solar system. The hypothesis is that there exists an orbital-inclined mini-Neptune far from the Sun that has not yet been discovered [32]. This hypothesis excites the public but is controversial in the scientific community. The concept of a massive trans-Neptunian (or trans-Plutonian at that time) planet has been discussed for a century, and there have been previous attempts to explain eTNOs with such an object [33, 34]. More recently, the hypothesis not only tries to explain why eTNOs exist but attempts to explain an apparent clustering in the orbits of eTNOs [32, 35, 36]. The statistical significance of this clustering is a highly controversial topic: some researchers find that the reported clustering is not statistically meaningful and/or due to sampling bias [37, 38].

Another hypothesis regarding a planet-bound to the solar system is called the Kuiper Belt planet hypothesis [39]. It is different from the massive planet hypothesis in the aspect of the mass and orbit of the planet. The hypothetical Kuiper Belt planet is modelled to have $a \sim 250-500au$, $q \sim 200au$, $i \sim 30^\circ$, and $m \sim 1.5-3M_\oplus$. This Earth-mass planet in a highly inclined orbit successfully generates eTNOs in simulation. The Kuiper Belt planet would be detectable with a survey that is sensitive sources as faint as $m_r \sim 24$ with sensitivity to apparent sky motion of $\sim 0.2 - 0.7 \frac{''}{hr}$.

Meanwhile, physicists focusing on Modified Newtonian dynamics (MOND) are attempting to find evidence of MOND on the scale of the solar system[40]. They claim that the relative motion of objects in the solar system is affected by the MOND field caused by Phantom mass, and orbits of eTNOs are due to MOND perturbation. They explain why they await the discovery of more eTNOs from projects like the Classical and Large-a Distant Solar SYstem Survey, the Dark Energy Survey, Transiting Exoplanet Survey Satellite, and the Vera C. Rubin Telescope Legacy Survey of Space and Time [41, 42, 43, 44]. Finding more TNOs might accidentally revolutionize physics if we offer matching evidence.

In reality, the current status of the eTNOs could emerge from any combination of these hypotheses: scattering effect by Neptune, stellar flyby, a bound/non-bound planet and more. Finding more eTNOs will aid in finding the origin of these eTNOs, eventually establishing a more accurate model of the past solar system. Can we develop systems that could enable the discovery of eTNOs within existing data sets that may have been acquired for purposes other than the discovery of eTNOs?

1.4 How TNOs are discovered

Observations at multiple epochs discover TNOs. When a TNO is observed on just one image, it is not so different from other unresolved point-like objects. As a planet-like object, we detect it as it wanders around the background stars and galaxies. The movement is mainly caused by the parallax induced by observations from different positions as the observer travelling around the Sun in Earth's orbit. When the TNO is observed at two different times, such that the Earth has moved in its orbit, the position on the sky of the distant solar system source appears projected against a different background of very distant stars: the source moves. The rate of motion (a combination of the projection along the line of sight, the orbital speed of the Earth and the orbital speed of the distant object) can be used to estimate the distance to the source. For estimated distances greater than 30 au, one expects that the object is a TNO. But preferably, the object is observed again within a few days or weeks to confirm the detection and also observed over many months and years such that a precise orbit can be determined. After these two or three observing nights with a time gap, we can estimate the 6 Keplarian orbital elements. These observations are then typically reported to the Minor Planet Center of the International Astronomical Union, which requires that an object's location on at least two nights be reported

before it can be classified as detected.

But to measure the positions of previously unknown moving light sources in an image, we must first detect its presence in a background of confusion and noise. Achieving such detections requires significant automation as there are many more irrelevant light sources, such as distant stars and galaxies, that can obscure the presence of a moving source, as well as noise sources, such as cosmic rays and detector defects, that can be present in just a single image or change position relative to background stars, mimicking the behaviour of a moving source. These sources of confusion outnumber, significantly, signals from the sparse population of TNOs.

Traditionally, catalogue-based searches have been key to automatically differentiating TNOs from irrelevant signals [45]. Let's say we have three images taken with a time interval of 1-2 hours on the same sky coordinates. The first step is to find the star-like objects in each image using a convolution kernel based on some assumptions about the shapes of point sources in the image and the expectation that TNOs will present as point sources. Programs like SExtractor or/and wavelet decomposition are commonly used [46]. The collection of these star-like objects is called a catalogue. The second step is to select the potential moving objects out of the catalogue by first removing for consideration any sources that appear at the same location in multiple image catalogues; these are stationary sources. As a result of this step, stars and galaxies are removed from each catalogue, and three catalogues with potentially moving objects are acquired. The third step is to select objects that can be made into linear triplets. To do this, a predicted position of a source can be interpolated from the position in the first and third catalogues in the time series. The predicted position can be compared to the actual positions of sources in the second catalogue to see whether a source exists at the predicted location. Sources that have been detected based on shape convolution and that appear at expected locations, based on some set of orbit expectations, in all three images are declared as TNO candidate detections. After this step, a catalogue of moving object candidates is produced, and further observations of the candidates are sought using predicted locations from orbital assumptions. Catalogue search methods enable the automatic detection of TNOs and are responsible for discovering essentially all the known TNOs (early searches used image blinking instead of catalogue comparison).




Figure 1.1: Number of TNOs listed to the Minor Planet Center each year (left) and the cumulative number (right), as of October 2023.

1.5 Challenges in discovering TNOs

Finding TNOs is challenging because there are a lot more TNOs to find. Although we found a significant sample of large TNOs, we still have not found most. Approximately 3000 TNOs have been discovered and reported (see Figure 1.1). However, the estimated number of TNOs brighter than current detection thresholds is much higher. Over 40000 TNOs are expected to be detectable [47] and we can expect to find tens of thousands of larger TNOs (≥ 100 km) with current instruments. Most TNOs are still out there, not yet discovered.

For smaller TNOs, the number increases even higher. Based on, a nearly completed inventory of the largest ($H_r < 5.5$, where H is the absolute planetary magnitude) cold classical ($42.5 < a < 42.5$; $i < 4.5^\circ$; $q < 37.5$, where a is the orbital semi-major axis, i the orbital inclination, q the pericentre distance) TNOs; the OSSOS++ sample for medium-sized TNOs ($5.5 < H_r < 8.3$); and other estimates ($8.3 < H_r$): the size distribution of col classical TNOs follows an exponentially tapered shape with a small-end slope of $N(< H) \propto 10^{-\alpha H}$ and $\alpha \approx 0.4$ [48]. The hot classical TNOs follow a similar exponential increase with decreasing size but do not exhibit the exponential cutoff at large sizes that cold classical show [49]. For objects with $m_r > 26.5$, the estimated number of TNOs is 70,000-250,500 [48]. There are hundreds of thousands of TNOs to be found.

On the other hand, it is challenging to discover TNOs because they are incredibly

dim. The difficulties of detecting TNOs come from the nature of observation of reflected solar lights. Because the sunlight travels to the TNOs and returns to the Earth-bound observer, the intensity at observation is proportional to only $1/r^4$ where the distance to the TNO is r . It differs from the $1/r^2$ law that self-luminous objects like stars and galaxies outside the solar system follow. In addition, as noted above, the albedos of TNOs can be as low as just a few percent. Hence, the TNOs are very dark, and the aforementioned $H=8$ magnitude objects result in $m \approx 24$ for TNOs at 40 au (the closest distance of the main belt of objects) and $m \approx 25$ at distances of 50 au (the outer edge of the main belt). Developing processes that are effective at detecting faint sources is, as in other areas of astronomy, a key to further exploration of the Kuiper belt.

Some studies use other approaches to overcome this $1/r^4$ problem. An emerging approach is the use of stellar occultation. Occultation is usually used to measure the diameter and shape in TNO studies, but it can also be directly used to discover unknown TNOs. The Taiwanese-American Occultation Survey used this method to study the number density of TNOs with the smallest sizes [50]. A field of stars is monitored for occultation events, and the serendipitous detection of such occultation events provides a measure of the column number density of objects along the line of sight to the stars. Another way to overcome the faintness of TNOs is the shift-and-stack image processing method. Shift-and-stack is a method to find a dim moving object beyond the detection limit of individual exposures by shifting the signal for moving sources to a common location and then co-adding the images. While just co-addition has long been used for stationary sources such as galaxies, the first usage of this shift and stack method was in 1995 [51]. The shift-and-stack approach is computationally bound as one must conduct trial shifts at every rate of sky motion that the target TNOs might have, with offsets between different trials being small enough to ensure that no signal is being lost due to trailing of the source in the shifted stack. With this technique, dimmer objects can be found, but the computational expense is not small. Modern large-area TNO surveys are based on catalogue search (section 1.4), and new GPU-based shift and stack methods are being applied to highly targeted small-area surveys. These techniques work to mitigate the flux loss from the $1/r^4$ loss (occultation) or compensate for it via image processing (shift-and-stack).

This thesis is an effort to find these dim and abundant TNOs using a more high-efficiency processing approach that is not sensitive to assumptions around the orbit

of the TNOs being sought. To study the solar system, it is best not to miss any TNO in the observing field up to the limit of the instruments. Can machine learning help this challenging task without needing a feature-extraction-based catalogue or shift-and-stack data engineering?

1.6 Artificial intelligence, machine learning and deep learning

Artificial intelligence, machine learning and deep learning are all trendy buzzwords, but they are different in meaning.

$$\text{Artificial intelligence} \supset \text{Machine learning} \supset \text{deep learning} \quad (1.1)$$

Artificial intelligence is an umbrella term of technology that requires computations to solve problems in our world. Anything from a simple system that filters emails by keywords to the DeepMind's AlphaGo that beat the world's Go masters can be examples of artificial intelligence. Machine learning is a subset of artificial intelligence and can be defined as a data-driven automation paradigm. Algorithms trained in the paradigm mainly learn from large data sets rather than from experts who already know how to solve the problem. A good example is a support vector machine, which is a method of classifying data points into two or more classes by finding the optimal hyperplane that separates the data points. Deep learning is a subset of machine learning and can be defined as machine learning with many layers of artificial neurons. Deep learning not only learns from large data sets but also learns how to extract features from the input layer. For example, ChatGPT predicts word-by-word responses to prompting sentences via a deep convolutional neural network with many hidden layers. ChatGPT extracts features such as words, sentences, and contexts from the input and predicts the following word in the sequence based on the extracted features. The category of artificial intelligence provides excellent promise to aid our understanding of the world.

Artificial intelligence, machine learning and deep learning came into practice chronologically. This is because each step to the next paradigm requires a technological advance in data capacity and computing power. Without GPUs with the computing power of more than a teraFLOPS and cheap storage and networks

enabling easy availability of petabytes of data, these advancements in artificial intelligence would not have occurred.

1.7 Astronomy and Image recognition deep learning algorithm

Astronomy intrinsically involves a large amount of data. Just 25 years ago, in 1997, the first astronomical digital image was achieved using a 256×256 digital pixel array [52]. Now, the NIRCam on the James Webb Space Telescope has ten 2048×2048 pixel detectors and the communication system is built to transmit data at 28 Mbps or 270 Gigabits per day [53]. Generating and researching big data is inevitable, as our universe is rendered in high-quality 3D!

Typically, deep learning is established on big data. In a review paper on deep learning methods, Alzubaidi et al. suggested that deep learning can be a better choice when the problem size is extremely large and cannot be handled with human reasoning abilities [54]. An example of large-volume data complexity in astronomy is deblending, detecting, and classifying sources for a wide-field survey. A Region-based Convolutional Neural Network was developed to classify sources in the DECam³ Legacy Survey images [55]. Previously, existing codes for detection and classification were highly sensitive to the density of sources and were found to be inefficient compared to their deep-learning approach. The trained CNN model can process a DECam image in 0.1 seconds and can adapt to new data easily with transfer learning. It is an excellent example of implementing a deep learning method for observational astronomy. As astronomy deals with extremely large data volumes whose complexity can be challenging, deep learning methods will excel in astronomy analysis.

³DECam is a 520-megapixel optical imaging camera operating on the Cerro Tololo Inter-American Observatory in Chile.

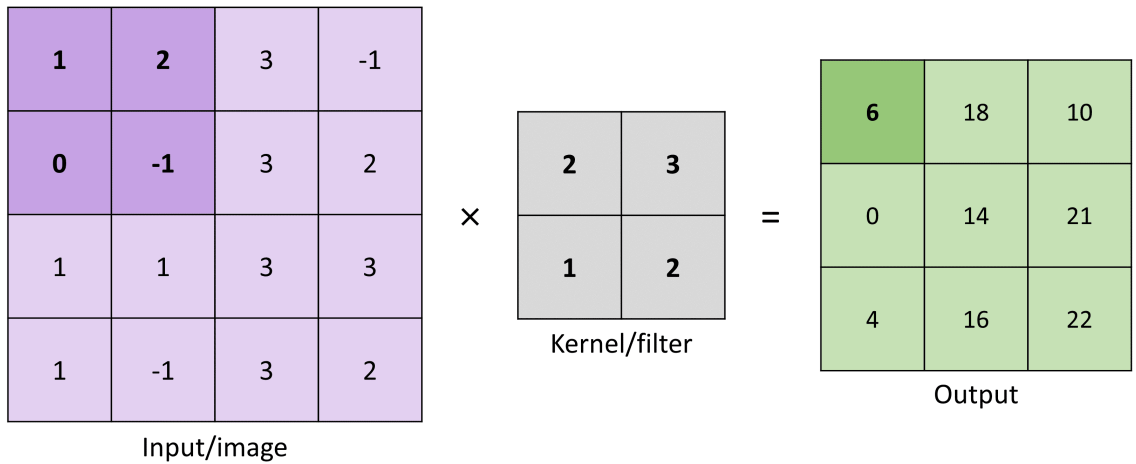


Figure 1.2: A diagram to explain the convolutional operation of the CNN. The filter/kernel is a 3D/2D array of values to extract features from images, and values in the array change throughout learning. Training CNN is a process of finding the best weights and bias (omitted in this diagram) to solve a problem. The input is an array, which can be an image or an output layer of the previous layer. The output array is calculated by element-wise multiplication of the input and filter/kernel arrays, applied regionally. For example, only the values highlighted in bold are calculated to generate the output value 6.

Out of many deep learning methods, we examined the Convolutional Neural Network (CNN) approach for our TNO detection because of CNN's powerful performance in image classification and object detection. CNN is a type of artificial neural network that utilizes convolutional layers. Convolutional layers are hidden layers that contain shared "weights" that are multiplied by the input of the layer. An example of image input, weights and their multiplied outcome is shown on Figure 1.2. In this way, features of the input can be extracted and passed to the next layer. The difference between fully connected neural networks and conventional convolutional layers is that weights are shared. While fully connected layers weight each connection, convolutional layers have weights that are smaller in weight dimension size and regionally process the input. This ability to extract features regionally is suitable for processing image data, and it is currently used to find tumours in our body, detect cars and pedestrians, recognize faces, structurally discover drugs, and more [56, 57, 58, 59]. Due to the widespread usage in object detection on 2D images, CNN was chosen for this thesis research. In this thesis, I present my development of a CNN-based TNO detection system.

Chapter 2 presents a manuscript submitted to the American Astronomical Society

journals; this chapter forms the core of my thesis research. In Chapter 3 I offer some additional analysis that explores the nature of our training data set and describes some concepts for future work. Chapter 4 provides the conclusion of my thesis.

Chapter 2

A Trans-Neptunian Object and Solar System Object Discovery Method with Convolutional Neural Networks for Future Surveys

Authors: Aram Lee, JJ Kavelaars, Hossen Teimoorinia, Wesley Fraser, Edward Ashton. This is a paper to be submitted to AAS journals.

2.1 Contributions

Edward Ashton provided the observational data, which were acquired for a different research project. JJ Kavelaars created the list of artificial sources to add to the images, established the research project and provided guidance and assistance in the analysis. Wesley Fraser conducted the artificial source generation and inserted those sources into the images. Hossen Teimoorinia provided consultation and assistance on the development of the Machine Learning approach to the problem. Aram Lee conducted the model training, data management, analysis and evaluation and wrote the bulk of the text.

2.2 Introduction

Trans-Neptunian Objects (TNOs) are the population of small solar system objects with a semi-major axis of more than 30 au (Neptune) and less than about 2000 au (Oort Cloud), and solar system objects (SSOs) are all natural bodies in the solar system that are neither the Sun nor planets. TNOs exhibit a range of physical sizes, albedos, shapes, surface colors and appear in a variety of different orbits. Some TNOs are as large as thousands of kilometers in diameter, while the smallest known objects are of order 10 km in diameter. Likewise, some have nearly circular and low-inclination orbits (e.g. 79360 Sila-Nunam) [60], while others have highly eccentric (e.g. 541132 Leleākūhonua) [61] or highly inclined orbits (e.g. 136472 Makemake) [62]. By searching for and discovering a large sample of diverse TNOs we can better understand the processes of planet formation and migration at work in the solar system. To achieve this goal requires developing efficient algorithms and processes for the discovery of these faint, slow-moving bodies.

Wide-field image TNO surveys, such as the Dark Energy Survey [42] and the Outer Solar System Origins Survey [63], used dedicated TNO discovery pipelines. Developing such pipelines is a time-consuming task. These surveys, and others, rely on detecting “features”, or sources. Determining which of those features remain stationary and then determining if any of the non-static features can be grouped in a way that is consistent with the expected linear motion of a solar system body [see 45, for a complete description of such a detection system]. A key component of the feature detection process is that one first engineers what features of interest might look like (e.g. a two-dimensional Gaussian with a particular profile width) and then uses convolutional processes to search for similar features within the data. This feature detection process relies on knowledge of the image quality and often needs to be tuned for the data taken on specific nights. Our goal was to conduct moving source detection directly from the pixel data. Working directly from the pixel data removes the need for feature-engineering based source catalogues [64, section 3] and may provide a more universally usable detection process. We refer to this as feature-agnostic moving source detection.

We demonstrate, using data acquired for a TNO specific search, that our technique detects TNOs as faint as those found in the original search while providing a low false-positive rate. Using this detection pipeline will aid in the detection of TNOs in imaging not specifically tuned to TNO dedication, such as the wide layer of Canada-

France Imaging Survey [65], which acquired imaging with 10 or 20 minutes between visits pairs and a total exposure time of 150 or 200 seconds of exposure time per visit.

Section 2.3 briefly summarizes the deep neural network (DNN) techniques we used. Section 2.4 describes the data preparation. Section 2.5 compares trained DNN models, and the best classification/regression model for TNO detection was selected and implemented on our data set. Section 2.6 describes how we used trained models for precise and time-efficient image analysis. Section 2.4, 2.5, and 2.6 make the method part of this paper all together. If you are familiar with machine learning, reading section 2.5 and 2.6 would be helpful to find SSOs with a similar approach. In section 2.7, we search data for TNOs, and conclude that the developed method can be used for other surveys as well.

2.3 Image Recognition Deep Learning Techniques

Modern and effective image recognition neural networks are deep and convolutional. Deep neural networks (DNNs) have hidden layers between the input and output layers. Those hidden layers are connected via weights that are initially random or uniform numbers and then optimized during the model training. (A “model” is a program that has been trained on a set of data to recognize certain types of patterns.) Back-propagation and stochastic gradient descent are used to change the initial weights to match the characteristics of the output layer against training labels [66, 67]. The trained weights act as circuits that identify features of the image input, thus, the algorithm learns from the image what the image features are rather than having those imposed at a source detection step. For example, some weights could be switches that only turn on if the input image has a bright galaxy filling the background. The training process of DNN is the process of extracting features from the training set, and training the weights to predict the class of the image as the output layer [68].

2.3.1 Tested Architectures

We explored a number of network architectures to determine which might be most effective at moving source detection. Here we briefly describe each of these architectures and then report on their accuracy and determine the most effective architecture for our problem.

Fully-connected Neural Networks

A fully-connected neural network (FCNN) is a basic type of neural networks. In an FCNN, every single node is connected with the all the nodes in the previous layer. In other words, if there are 32 nodes in both the previous layer and the current layer, there would be 1024 weights to be trained between the layers. We utilized the `keras` package for modelling and construct a basic FCNN architecture by stacking multiple **dense** layers. Our FCNN was built with an input layer (2, 64, 64), one flattened layer (8192), two dense layers (256), and an output layer (2).

Table 2.1: Architecture List with performance and the number of parameters on ImageNet validation set. The depth is the number of convolutional layers and fully connected layers.

Arch.	Year	Accu.	Param.	Depth	Characteristics
FCNN	-	-	2M	3	A simple fully-connected NN
AlexNet	2011	57.2%	60M	8	ReLU, overlapping pooling
VGG-16	2014	71.5%	138M	16	simple 3×3 filters
ResNet50	2015	79.3%	26M	50	residual units
MobileNet	2017	70.6%	4.2M	28	low computational cost

Convolutional Neural Networks

A Convolutional neural network (CNN) connects each layer with shared weights. A block of shared weights is called a filter, and this filter moves along the entire layer and abstracts the previous layer. If the filter is set to size 3×3 , the number of weights will be 9 compared to 1024 in a fully connected network. CNNs are not only inexpensive in terms of processing time, but also express spatial information from the image, making them essential for image recognition using deep neural networks.

ImageNet is a large image data set with hundreds of categories defining images [69]. AlexNet and the following DNNs were developed for ImageNet Large Scale Visual Recognition Challenge, or ILSVRC which is a competition for developing neural networks of classification and detection of ImageNet images [70]. We used architectures designed to perform well on the ImageNet, and the table 2.1 summarizes characteristics of FCNN and CNN-based architectures used in the competition. Note that the accuracy values of AlexNet, VGG-16, and MobileNet are from Howard et al, and the value of ResNet50 is from He et al. [71, 72], and also the

accuracy varies slightly from year to year for the same architecture because the ILSVRC data set used differs from year to year.

AlexNet-like: 2011 AlexNet was one of the early architectures to make use of multi-GPUs for training [73]. Even though AlexNet is a relatively shallow architecture with only 8 neural network layers, AlexNet outperformed other architectures at the time by introducing the following concepts. First, hidden layers are activated with ReLU non-linear functions instead of former tanh functions [74]. Using ReLUs is much faster than tanh in training. Second, dropout was introduced to mitigate over-fitting [75]. Third, max-pooling layers were overlapped. Fourth, local response normalization was introduced. The third and fourth techniques reduce error rates slightly [76].

VGG-like: 2014 VGG is another classical CNN architecture designed to be trained on the 224×224 RGB images of ImageNet. VGG builds 3×3 convolutional layers with the “same” padding.¹ The architecture mixes the max-pooling of 2×2 with strides 2 between convolutional layers followed by 3 fully connected layers [77].

AlexNet and VGG are built around (3, 224, 224) input shape where each number means number of channels, the width of each channel, the height of each channel, respectively. We modified AlexNet and VGG to match the input shape of (2, 64, 64), but kept the overall architecture.

Advanced Deep Neural Networks

Out of the currently popular high-performance architectures, we chose ResNet and MobileNet. They are both based on CNNs but have their own tricks to achieve their advantages, such as having deeper hidden layers with less complexity or being mobile-friendly (Low latency and model size).

ResNet: 2015 The performance of deep neural networks does not increase consistently with increasing depth. The error rate tends to increase after the model reaches a certain depth that is optimized for the data set. Residual units can overcome this problem and allow the use of deeper networks [72]. The Residual unit is a function to build the ResNet. The function is defined as $y = F(x) + x$, where the $F(x)$ is the combination of weight layers and ReLU with the input of x , and the

¹In this context, “same” padding means that padding is applied around the input layer with zeros in order to maintain the same width and height as the input layer after the convolution.

second x is the input information itself. This addition of the identity function, $+x$, is what makes ResNet different from previous CNNs, and it gives an information shortcut to the previous layers. ResNet uses those residual units for every two convolutional layers, solving the problem that extremely deep neural networks perform worse than relatively shallow neural networks, and creating better networks.

MobileNet: 2017 MobileNet has a relatively low number of parameters compared to ResNet and a short training time, but with a comparable accuracy to modern DNNs [71]. In MobileNet, standard convolution filters are separated into depth-wise convolutional filters and point-wise convolutions. In this way, the computational cost of MobileNet is 8 to 9 times less than the standard convolutional filters. Sometimes computationally expensive architectures are unnecessary; therefore, we experimented with this kind of lightweight architecture to determine if the savings in computing cost comes without a loss in effectiveness.

ImageNet architectures such as ResNet and MobileNet are available as ready-made architectures of Keras and are automatically modified to the required input shape (e.g. 2, 64, 64). Out of various versions of those architectures, we used ResNet50 and ResNet50V2 for ResNet, and MobileNetV1 and MobileNetV2 for MobileNet.

2.3.2 Classification and Regression

We created separate binary classification and regression models for all the architectures mentioned in section 2.3.1. The only difference in architecture between the classification model and regression model is the final few layers. The classification models are trained on an equal number of positive (has a moving source) and negative (does not have a moving source) cases, while regression models are trained only on positive cases.

For classification, the output layer is 2 valued (p_1, p_2), the probability that the sub-image (one value for each channel) is in the ‘has a moving source’ class. To make p_1 and p_2 a value between 0 and 1, the final layer has a sigmoid activation function. For regression, the output layers are the model’s estimate of the position of the source in each of the two images (x_1, y_1, x_2, y_2) and the magnitude of the source (m). Here the activation function is not a sigmoid, but it was a linear function for the regression model. The predicted source locations from the linear output values are then used in

our linear fitting and candidate vetting steps.

In detail, the Keras applications like ResNet and MobileNet have arguments such as `input_shape`, `classes`, and `classifier_activation`. Using these arguments, a user can modify the architecture easily for their desired input and output shapes. As default final layers were sometimes not suitable for the user's goal, `include_top` argument can be set to `False` and custom final layers such as multi-output `dense` layers for separate position and magnitude measurement and additional dropout layers can be inserted. In this way, the same regression model can be used for the measurements of positions and a magnitude.

2.3.3 Model Implementation

Our analysis was carried out using Keras module within the TensorFlow package based on TensorFlow 2.4. The Keras module in TensorFlow offers various models and layers, and ready-made architectures described in 2.3.1 were imported from the `keras.applications`. For the FCNN and some CNN architectures, the model was built from a series of layers such as `dense`, `Conv2D`, and `MaxPooling2D`. These models were compiled with optimizer Adam and loss function of binary cross-entropy for classification and MAE for regression [78]. With the model initialized, the model is optimized with callbacks to `EarlyStopping` [79]. The maximum number of epochs was set to 30, and the early stopping function usually halts the training before reaching 30 epochs when it was set to have `patience=7`. `EarlyStopping` stops the training at the lowest validation loss and aids in avoiding over-fitting.

The hardware we used was a single NVIDIA A100 GPU. With the A100, each epoch took up to 10 minutes, taking about up to 2 hours for the model to finish training when it was trained on a data set taken from a single CCD of the CFHT MegaCam mosaic. We also trained using a data set that combined data from 3 CCDs (C051020M27), which required up to 6 hours to train a model. While any graphics card with over 8GB RAM is capable of this job, employing an NVIDIA GPU supporting fp16 mixed precision feature would decrease the training time greatly.

2.4 Data and Pre-processing

2.4.1 Original Images




Figure 2.1: (Left) An example of the images from CFHT MegaCam before cutting. This is the upper left part of a single CCD. Around 200 artificial TNOs are present in this image; most of them are small and dim, making them hard to find on this scale. (Right) Measured positions by the regression model of all 44 observations at this location. Only positions that passed the classification are shown here. The short, almost horizontal lines are mostly the signature artificial TNOs. The vertical line from the top to the middle is the collection of incorrect measurements of moving objects caused by the noise from very bright stars. Large, empty, circular spaces are locations where a bright star has obscured detections. On the bottom right, a real SSO candidate is detected as a longer inclined streak.

Our experimental data set consists of a series of *gri.MP9605* filter images obtained with the Canada-France-Hawaii Telescope (CFHT) MegaPrime camera [80] (See

Fig. 2.1). The images were obtained as part of proposal ID 19AC24 (“The Size Distribution of Saturn’s Irregular Moon”). From this program, we selected 44 images of the target field “SatEastDis” obtained between 2019-07-01T09:07:02Z and 2019-07-01T12:07:47Z.

Each gri.MP9605 filter image in the sequence was approximately 205 seconds in exposure time, with a time interval of 248 seconds between each image (205 seconds of exposure time and 33 seconds for camera readout). The MegaPrime camera consists of a mosaic of 40 CCDs, with each camera image stored in a multi-extension FITS file. For this research, only images from 36 out of the 40 CCDs were used, excluding the images from 4 CCDs on the sides. We treated each CCD in the mosaic and each exposure independently, resulting in $44 \times 36 = 1584$ independent images. We selected 3 of the 36 CCDs (5, 10, and 20) to train and validate our model. In all, the training data set consists of $3 \times 44 \times 2048 \times 4612$ pixels = 1.247 Gigapixel. We then tested the trained models on data from other CCDs in the mosaic.

2.4.2 Data Calibration

Before passing the telescopic images to the machine learning process, we must remove from the data characteristics of the instrumental system (sensitivity variations). In CFHT parlance, this is referred to as “detrending” the observations. The raw images were detrended using the CFHT Elixir pipeline, which is a collection of programs and tools for processing and evaluating telescope data [81]. This detrending step was performed by the CFHT and distributed to the users via the Canadian Astronomy Data Centre (CADDC).

The detrended data must then be calibrated to place the pixel values onto a spatial and flux reference system, in this way the values in pixels in different images can be inter-compared. We used version 19.1 of the Legacy Survey of Space and Time (LSST) Science Pipeline to perform this calibration step [82]. In particular, the `processCcd.py` command was used to photometrically and astrometrically calibrate these data against the Pan-STARRS DR1 catalog. At the time of use, the *MegacamMapper*, which enabled ingestion of the Elixir images into the *LSST* pipeline, did not include a term converting known colours of PS1 sources to the gri filter used for the observations. Rather, the images were treated as having been acquired in the Megacam r-band filter, and so the final photometric calibration will

suffer a small offset (± 0.05 magnitude depending on the color of the source of interest). Such an offset, however, does not impact the astrometric calibration and does not impact the goal of the current study. The output of this pipeline was the *calexp* files used for the rest of this project.

2.4.3 Synthetic Trans-Neptunian Objects

To generate a list of known moving bodies from which our CNNs could be trained, artificial moving sources were injected into the imagery. Artificial sources with orbits consistent with TNOs were randomly generated on heliocentric orbits with semi-major axes $30 < a < 50$ au, eccentricities $e < 0.5$ and inclinations, $i < 45^\circ$, with uniform sampling for each parameter. The remaining 3 orbital angles were sampled uniformly. Each source was assigned an apparent magnitude drawn uniformly between $22 < r < 27$. Roughly 350,000 sources were generated in this way. We then computed ephemeris locations for each source at the mid-point of each of the exposures in our training set. Those with ephemerides placing them outside the field of view of the imagery in question were discarded while those on the field of view were ‘added’ to the imaging.

Injection of artificial sources was done with the software package *TRailed Image Photometry in Python* TRIPPY [64]. For each image, Point-Spread Functions (PSFs) were generated automatically. Candidate reference PSF sources were first selected as those sources containing no saturated pixels, and with a signal-to-noise ratio (SNR) greater than some threshold. Each source was checked for a nearby background source that would influence the quality of the PSF, and only fully isolated point-like sources were kept. The SNR threshold was reduced in steps until either 40 suitable sources were found, or the threshold reached SNR=50. The selected point source images were then fed through the *psfStarChooser* routine, which fits a Moffat profile to each source. Sources found to have outlying Moffat α or β values were rejected. The remaining sources were used to generate the PSF on a per-image-per-chip basis. These PSFs are then used by TRIPPY to inject signals consistent with a moving TNO into the imaging.

For each chip, artificial TNOs were implanted. With TRIPPY, each artificial moving source was implanted at the correct position and trailed according to the source’s rate of sky motion, and the exposure time of the image in question. The brightness was scaled to account for zeropoint variations between each image as determined by the LSST pipeline. Gaussian shot-noise was added considering the

gain of the appropriate analogue-to-digital converter (ADC) for each chip region, and the instrumental brightness of the source. No effort was made to adjust the brightness of the implanted sources to account for the small systematic offset in the zeropoints caused by the lack of colour term during photometric calibration. This process results in images that contain a large sample of point sources that change locations between exposures in a way that is consistent with motion of TNOs. These artificial TNO implanted images were used to train and test our CNN model.

2.4.4 Building Image Subsets

Each of 1584 FITS images with artificial sources consists of an array of 2048×4612 pixels, which is too large for model training and contains too many unrelated features. To create a more manageable data set we created sets of 64×64 -pixel sub-image pairs, where each sub-image in the pair is of the same sky location but taken at a different time. As the camera was not moved between exposures each sub-image in the pair is from same CCD within the mosaic. To cover all time separation, a pair list of index combinations from 1 to 44, with the first value in the pair always being lower, was created. (i.e. (1, 2), (1, 3), ..., (43, 44)) And, the order of the list was randomly shuffled. We then selected images for the pairs by looping over this sequence, ensuring that the full range of time separations were explored. These sub-image pairs provide a two-channel input, with the first channel being the earlier observation and the second channel the later observation of the pair. The individual exposure time was 205 seconds and the total observing period was 3 hours, resulting in the time separations between two channels ranging from a few minutes to 3 hours.

We selected 500 random sky locations within each CCD and then constructed all possible time-series combinations of those 500 sub-images locations as our training/validation data sets. Before training, we randomly split this data set into training (80%) and validation (20%). This provides a training set that does not include the full background image, preventing the model from learning the underlying stellar background, as demonstrated by our test data sets below. These time and position pairs determine the two FITS files to use and the centre coordinates of the sub-images, therefore $500 \times 44 \times 43 \div 2 = 473,000$ coordinate centre and exposure number combinations are constructed to build the full time series of data for each CCD. (see Figure 2.2). This amount of combinations corresponds to approximately 3.9G pixel values as each pair contains $2 \times 64 \times 64$ pixels.

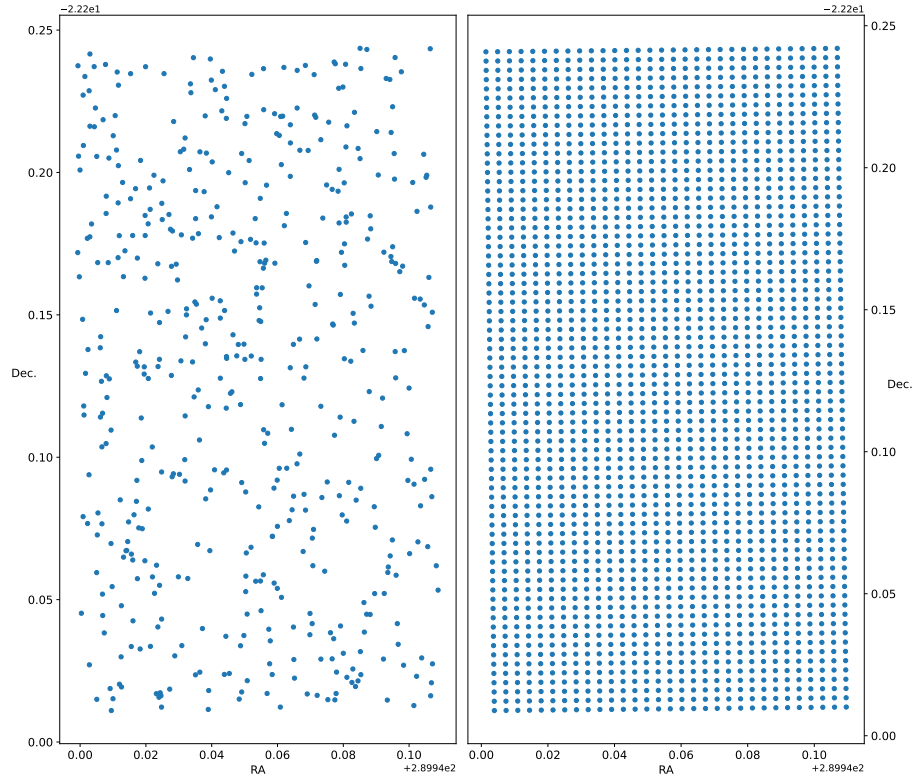


Figure 2.2: This figure shows sky coordinates of centres where images are cut into 64×64 pixel sub-images. The left plot has 500 randomized cutting positions, and the positions were different for each FITS file pair, and the right plot has evenly spaced cutting positions with 63-pixel gaps. The left one is to have a randomized training/validation set with a less volume, and the right one is to retrieve the most planted and non-planted TNOs by searching the whole area.

To generate the sub-images, the pixel coordinate on the first image pair is converted to sky coordinates and the second sub-image is centred at that sky location, ensuring the two sub-images are of the same location on the sky. For example, if the first element of time pairs was (2, 33), we then generated 500 pairs of sub-images from the 2nd and 33rd FITS files (which are separated by about 2 hours) with the centres of the sub-images taken from the randomized grid list locations in the 2nd image and the same centres of the sub-images in the 33rd image. The randomized centres changed for each time pair.

Each data set presented here is labelled by the number of the CCD from which it has been drawn. For example, the sub-images made from the 20th CCD are labelled “C20 images.” A global training set was constructed by the concatenation of the C05, C10, and C20 data sets (combined data from multiple CCDs). This global data set

provides artificial TNOs that had been made with a broader range of PSF models. Providing a number of different PSF shapes to guard against the model being trained to recognize the shape/structure of the artificial sources rather than the motion of the sources. The combined data sets “C051020” was used as the main data set for training and testing the performance of our model.

We also created sub-samples for each CCD group of images based on selecting for training only those images that contained a TNO brighter than some limiting magnitude or images that did not contain a TNO. For example, the data set labeled “C20M23” contains only those 64×64 sub-images that contain a TNO brighter than the 23th magnitude as well as an equal number of sub-image pairs that do not contain a TNO.

For the implementation section, sub-images from the whole grid were used. In the Figure 2.2, the sampling difference between the training sets and testing sets is shown. Unlike the randomized sub-image centre positions on the left plot, the gaps between centre positions are constantly 63 pixels and the sub-images cover the whole CCD plane. The step between sub-images was set at 1 pixel smaller than the 64×64 pixel sub-image to avoid an edge problem, which is a problem of having a moving object right at the edge of a sub-image. This full data set is used to evaluate the detection efficiency of our algorithm.

Table 2.2: Examples of CFHT Megacam datasets names used in this research

	5th CCD	10th CCD	20th CCD	5th+10th+20th
23th Magnitude	C05M23	C10M23	C20M23	C051020M23
25th Magnitude	C05M25	C10M25	C20M25	C051020M25
27th Magnitude	C05M27	C10M27	C20M27	C051020M27

In summary, for training sets, we used FITS files from three CCDs (The 5, 10, and 20th CCD) to generate two-channel sub-images at randomized positions on the image. The sub-image sets from the nth CCD were named as “CnM27”, and its sub-samples containing only bright ones are named as “CnM23” or “CnM25”, while the 23, 25, and 27 means the maximum magnitude of planted objects in the sub-image sets. A data set made from concatenating C05, C10, and C20 were made as well and was named as “C051020.” Table 2.2 provides the names of data sets, and Figure 2.3 shows the examples of sub-images containing TNOs with magnitudes between 22 and 27.

Testing requires presenting to the model data that has not been used for training and validation. Our testing data set consists of sub-images from all other CCDs not used in the training/validation process.




Figure 2.3: This figures shows 6 pairs of sub-images. Examples of sub-images used as our inputs. The orange circle marks the location of the synthetic TNOs. From the left upper pair to the lower right pair, 6 pairs of positive-negative sub-images are shown. For the magnitude 22, 23, and 24 cases one can see the bright source in the positive sub-image that is absent in the negative sub-image taken from a different epoch. The magnitude 25, 26 and 27 sources are progressively more difficult to distinguish. The artificial sources range from magnitude 22 to 27.

2.4.5 Labels

Supervised learning requires labels attached to input data. The input data consists of the sub-images described in section 2.4.4, and the label for each sub-image indicates the presence of an artificial TNO in that sub-image. Sub-images containing an artificial TNO were labeled “1” while sub-images that do not have a TNO were labeled “0.” The labelling for classification was binary. For example, the first two images on the upper left of the Fig 2.3 are from the same sky coordinates.

The left image has a TNO of $m \approx 22$, and the right image does not have a TNO. Therefore, the pair of images would have (1, 0) as the label of classification.

In addition to a classification model, we trained regression models to determine the location and magnitude of the TNO to be used after the classification step. To achieve this, each sub-image had labels holding the position and brightness of the TNO in the image. This augmented information was not used by the classification model and is only made available to train a regression model, which provides the position of the TNO within images classified as containing an object. Since images were 64×64 pixel size, those images were labeled with x and y positions values between 0 and 64 and with apparent magnitude values between 22 and 27. As a result, each sub-image on the first channel and the second channel was labeled with (p1, p2) for the existence of a TNO, and (x1, y1, x2, y2) for x and y coordinates of the TNO, and (m) for the magnitude of the TNO.

2.5 Model Selection

The performance of predicting the existence of a moving source in sub-images varies by the architecture and according to the training and validation set used. Here we present a summary of results from our exploration of various combinations of models, training and validation sets.

2.5.1 Architecture comparison

Our initial examination looked at that the overall performance of our set of machine learning models. To measure performance, we need to determine the rate at which the model detects moving sources in our data. We intentionally added moving sources that are well below the single exposure detection threshold as we wish to explore if our detection process can reach deeper into the data than more traditional approaches. If we used sources that are well below the detection threshold in measuring the performance of the model, however, we are not providing a realistic evaluation of the accuracy of our models. The data sets for performance comparison only included sources brighter the $r < 25$ magnitude (which is approximately the single-image detection threshold of our data set). Our performance testing is used to select, from those networks tested, the neural network that is most effective for our particular image classification and position determination problem.

Classification

We present the performance and characteristics of the networks we tested in terms of: the number of parameters, recall, precision, F1 score, number of true positives (TP); number of false positives (FP), number of false negatives (FN); and number of true negatives (TN), see in Table 2.5.1. All Metrics shown on the table were metrics when the classification threshold was set to be 0.5 (p values from the classification model must be larger than 0.5 to classify the images into positives). Among these metrics, the recall indicates how well the models retrieve (detected) the artificial sources in the data set, recall is the primary metric of interest for our detection system. Precision indicates how well the models exclude false positives out of all positively classified cases. The precision is slightly less important than the recall because the false positives can be filtered using post-processing and visual inspection at later stages of this method (section 2.6). However, the precision is still an important metric because a low precision rate would create a high visual inspection burden. Those two metrics introduced above can be expressed as the harmonic mean, the F1 score, Which can be used as a single number to express both recall and precision metrics. The metrics are defined as:

$$recall = \frac{TP}{TP + FN}, \quad precision = \frac{TP}{TP + FP}, \quad \text{and} \quad F1 = \frac{2}{\frac{1}{recall} + \frac{1}{precision}}. \quad (2.1)$$

There was a strong variation in the performance of the different models. FCNN

Table 2.3: Classification performance of tested networks, based on the validation set of C051020M25. Larger is better for recall, precision, and F1, and the metrics are explained in section 2.5.1.

Architecture	Para.	Rec.	Prec.	F1	TP	FP	FN	TN
FCNN	2.2M	0.00	0.00	0.00	0	0	120045	119955
AlexNet-like	2.8M	0.79	0.88	0.83	94414	12986	25631	106969
VGG-like	0.86M	0.82	0.90	0.86	98915	10824	21130	109131
ResNet50	23M	0.93	0.93	0.93	112156	8166	7929	111749
ResNet50V2	24M	0.90	0.96	0.93	108105	4877	11980	115038
MobileNet	3.2M	0.91	0.92	0.92	109631	9978	10454	109937
MobileNetV2	2.3M	0.91	0.89	0.90	109708	14212	10335	105745

failed to perform the image recognition task by classifying all image pairs to negatives. The CNN architectures are made for 2D image recognition tasks and even a simple architecture such as AlexNet and VGG could be trained on our data set and classify the data set with recalls of approximately 80%.

After we checked the capability of CNN-based networks on our data set, we proceeded with more complicated networks. The ResNet50 and MobileNet architectures were our choices, and they were automatically modified for our input shape. They showed better performance than AlexNet and a VGG with recalls of higher than 90%. ResNet50V2 and MobileNetV2 showed similar performance to ResNet50 and MobileNet. ResNet, MobileNet and their V2 networks provided us good enough performance to achieve our goals.

The recall and precision values reported in Table 2.5.1 are computed for a classification threshold of $p > 0.5$. Recall and precision values are traded-off against the likelihood value used for classification. For example, if we set a higher threshold, the precision would go up, but the recall would go down. The threshold can be adjusted to achieve the highest model utility for the experimental setting. In other words, if TNOs detection per image is not critical but false positives are problematic we might set a higher threshold and regard precision to be more important than recall. If, however, we need to maximize detection even at the expense of a higher visual inspection burden, we can set a lower threshold and regard recall to be more important than precision. The impact on the overall detection efficiency caused by adjusting the threshold is illustrated in Figure 2.4.

Table 2.4: Regression performance on C051020M25 validation set. ‘pos.’ and ‘mag.’ in front of MAE and RMSE refer to positional (pixel) and magnitude uncertainties respectively. Lower values indicate better performance.

Architecture	Param.	pos. MAE	mag. RMSE	pos. MAE	mag. RMSE
FCNN	4.3M	10.5	16.0	1.8	3.0
AlexNet-like	4.0M	1.6	5.3	0.2	0.3
VGG-like	1.4M	5.8	6.9	0.2	0.3
ResNet50	24M	0.8	3.7	0.1	0.2
ResNet50V2	24M	0.8	3.9	0.1	0.2
MobileNet	3.2M	0.91	3.8	0.1	0.2
MobileNetV2	2.3M	3.2	7.1	0.3	0.5

Regression

The performance of predicting positions of TNOs also varied by architecture. The performance and characteristics of the regression model in terms of the number of parameters, the mean absolute error (MAE), the root mean square error (RMSE) of source positions and brightness are given in Table 2.4. The FCNN guessed the location of the TNO almost randomly, but VGG showed an improvement in all metrics especially on magnitude prediction, and AlexNet could predict the positions far better than we expected. With this promising performance of CNN-based networks, we proceeded to explore ResNet and MobileNet. ResNet50, ResNet50V2, and MobileNet predicted coordinates with an MAE less than one pixel and were found to be the best architectures among those tested. Interestingly, MobileNetV2 was worse at position prediction than MobileNet.

Out of the architectures we have tested, we have chosen the MobileNet to be our architecture. MobileNet exhibits sufficiently high recall and precision when was used as a classifier, and low MAE and RMSE when used for regression. ResNet50 showed slightly better performance, but the difference was minimal but ResNet50 has 10 times the number of parameters and is more susceptible to over-fitting. Additionally, MobileNet is highly efficient due to the smaller number of parameters. The computational cost of predicting with MobileNet is relatively cheap compared to ResNet [83]. Any number of advanced CNN-based networks could be suitable to our problem and continuing to explore these various architectures may reveal even better performance, but for our purposes the MobileNet network provides a good balance between performance and complexity. For the remainder of this manuscript we adopted MobileNet.

2.5.2 Training Data

After selecting a model we examined the nature of the training data and how the training data influences the model performance.

Single CCD

First, we train and validate on a data set from a single CCD. For example, a training set based on 80% of the entire C05M25 set and a validation set based on 20% of the C05M23 set. By validating on the brightest data set we are exploring the networks

Table 2.5: Classification performance on validation sets (same chip) and test sets (different chips) at p threshold=0.5. The models are trained on M25 data sets. The models were validated/tested on a balanced M23 data set, which is a subset of the M25 validation/test sets, for a better presentation of their effectiveness on brighter objects.

Model	Validation		Test 1		Test 2		Test 3		Test 4		Test avg.	
	Rec.	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.	Prec.
C05	0.96	0.90	0.91	0.84	0.90	0.88	0.66	0.68	0.66	0.66	0.78	0.77
C10	0.91	0.90	0.90	0.82	0.92	0.88	0.72	0.68	0.71	0.66	0.81	0.76
C20	0.90	0.90	0.62	0.68	0.63	0.65	0.93	0.89	0.89	0.87	0.77	0.77

ability to detect the moving source when sufficient signal is present in the data. Likewise, C10 and C20 models were also trained and validated. The “validation” column of Table 2.5.2 shows the performance of each model on its validation set. The validation performance is fairly high for all sets.

However, when these models were tested on data sets from other chips, the classifiers were significantly less effective. For instance, the C05 model was effective to classify the C08 data set and C13 data set, but it was not effective for C21 data set and C34 data set. The network appears to be training on characteristics specific to a particular CCD (perhaps the structure of the PSF). To improve the model performance requires an expanded training set.

Multiple CCD

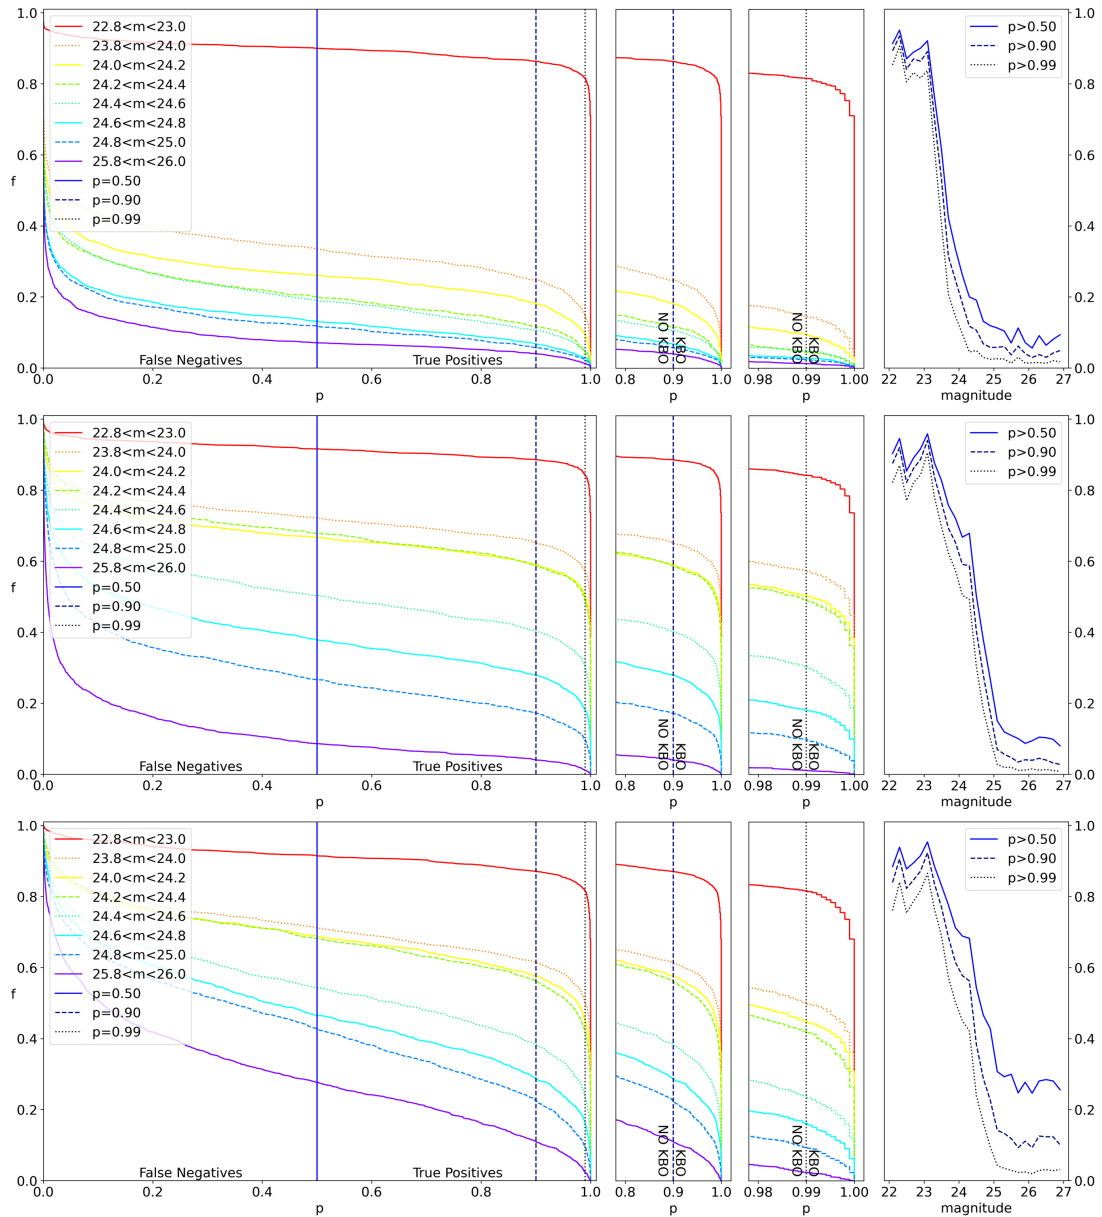


Figure 2.4: Cumulative plots of detection efficiency of the MobileNet model trained using different input data sets. Upper: Trained on data with sources brighter than $m=23$. centre: $m=25$. Lower: $m=27$. On the left side, each panel shows classification probability (p) versus fraction (f) of sources with a moving source correctly classified. The vertical blue/navy/black lines show various thresholds that can be used to select images that may contain a moving source. Thresholds of 0.5, 0.9, and 0.99 are shown, and the middle panels show the model behavior near the 0.9 and 0.99 thresholds. The panels in the right column show the fraction of sources correctly classified as a function of the brightness of the source, also known as the completeness curve.

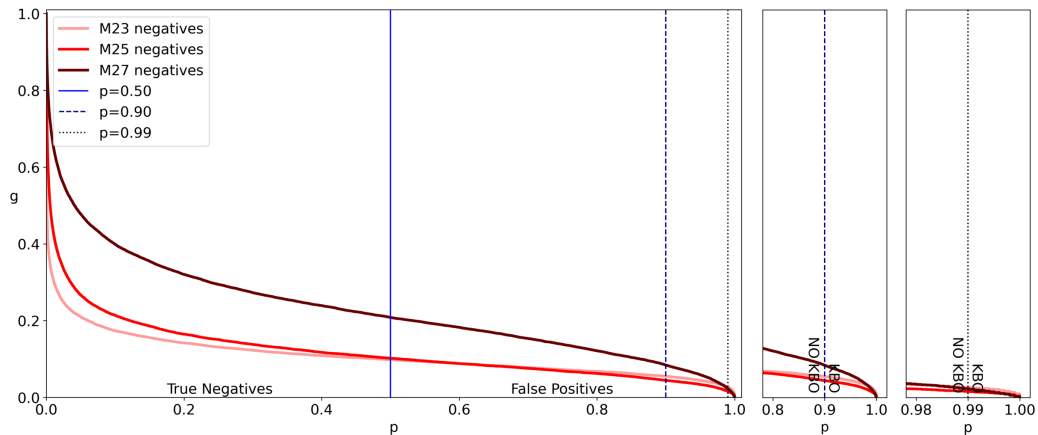


Figure 2.5: Cumulative plots of the model probability (p) versus the fraction of images classified as containing a TNO (g). All image pairs used for this plot are negatives, i.e. do not contain any artificial moving sources. Thus, g is the number of false positives compared to the total number of negative sub-images. There is very little difference between the M23 and M25 false-positive rates. The M27 model, however, returns false positives at double the rate of the M23/M25 models when $p=0.5$.




Figure 2.6: ROC curves from testing the C051020M25 classifier on image pairs. The x-axis is the false positive rate that shows how frequent false alarms are compared to the total number of negatives (Less is better). The y-axis is the true positive rate that shows how many positives are correctly classified as positives compared to the total number of positives (More is better). These ROC curves show the general performance of the classifier when applied with various p thresholds. The more the curve is skewed to the left-upper area, the better the classifier performs. Area under curve (AUC) indicates the performance of the classification model (More is better). The ROC curves have AUCs of 0.87, 0.82, and 0.74 depending on the maximum magnitude limits of 23, 24, and 25. The diamond markers indicate points at $p=0.99$ classification threshold.

C051020, the combined data set described in section 2.4.4, provides a more diverse training set. The model trained on a single chip data set showed recall and precision mostly below 80%. (see Table 2.5.2) When a MobileNet model was trained on C051020M25 data set, however, the test results on C08M23, C13M23, C21M23 and C34M23 data sets were evenly and sufficiently high at approximately 90%.² (see Table 2.5.2) Using the larger and more diverse training set appears to have removed our over-fitting problem. For the remainder of the manuscript we considered only the MobileNet model trained on the C051020 data set.

Magnitude Limit on the Training Data

We have analyzed the M23, M25, M27 models to determine which magnitude limit on the training data gives the best result. In Figure 2.4, the completeness of each model trained on a combined test set C08132134 is shown. M23 model was effective to classify a test set up to the 23th magnitude objects, and quickly became ineffective faint-ward of approximately 23.5: the model trained on bright sources could not predict the existence of a moving source significantly fainter than in the training set. M25 model could classify images with objects brighter than the 23rd magnitude, similar to the M23 model. But, the recall drops quickly fainter-ward to approximately 24.3 instead of approximately 23.5. M27 model showed similar aspects to the M25 model. Because the model was trained with a data set including up to the 27th magnitude objects, it could retrieve some of dim objects, but by using the M27 model to classify images, there was a problem that the fraction of false positives increases (see Figure 2.5). The fraction of false positives compared to the total number of negative sub-images of the M27 Model, the fraction g on the Figure 2.5, was approximately 21% when $p=0.5$, indicating a significant population of false positives would be included after the classification. Therefore, C051020M25 model can correctly classify images including faint sources as well as the M27 model but with fewer false positives than the M27 model. The general performance of this model is shown with ROC curves on the figure 2.6. We concluded that the ~ 24.5 th magnitude is the limit we can successfully classify objects with this method, and we decided to use the C051020M25 MobileNet model going forward.

²There is no particular reason why the 5th, 10th, 20th CCD data sets are chosen for training, and the 8th, 13th, 21th, 34th CCD data set are chosen for testing, the choice of which CCDs to use for training and which to use for testing was made arbitrarily.

Table 2.6: Classification performance of the combined data set models on validation sets (same chip) and test sets (different chips) at p threshold=0.5. The table displays the performance of the models trained on an M25 data set when validated/tested on a balanced M23 data set. The suffixes “R” and “M” after the model names denote ResNet and MobileNet.

Model	Valid.		Test 1		Test 2		Test 3		Test 4		Test avg.	
	Rec.	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.	Prec.
C051020R	0.98	0.94	0.91	0.89	0.92	0.92	0.93	0.90	0.89	0.88	0.91	0.90
C051020M	0.96	0.94	0.90	0.90	0.92	0.91	0.93	0.89	0.89	0.88	0.91	0.90

Table 2.7: Regression performance of the combined data set model on the validation set (same CCDs) and test sets (different CCDs) at threshold = 0.5. The table displays the performance of the models trained on an M25 data set when validated/tested on a balanced M23 data set. The suffixes “R” and “M” after the model names denote ResNet and MobileNet.

Model	Valid.		Test 1		Test 2		Test 3		Test 4		Test avg.	
	pos.	mag.	pos.	mag.	pos.	mag.	pos.	mag.	pos.	mag.	pos.	mag.
C051020R	0.41	0.08	1.9	0.2	1.6	0.2	1.3	0.1	1.3	0.2	1.5	0.18
C051020M	0.5	0.1	1.9	0.2	1.6	0.2	1.3	0.2	1.9	0.2	1.7	0.2

2.6 Implementation

Given the trained and selected models, we now examine how these models can be used to discover SSOs. The C051020M25 MobileNet classifier was our choice of classifier, and C051020M25 MobileNet regression model was our choice of predicting positions and magnitudes. Requiring $p > 0.99$ for our classification model correctly filters out $\sim 98\%$ of images without a moving source (see Figure 2.5.) Even the 2% leftover from the CNN classification, however, results in a large number of false positives. To further reduce the false positives, we investigated two kinds of post-processing methods.

2.6.1 Approach 1: Linear Fitting Approach

We first tried to use both classification model and the regression model. In this study, the classification model measures the likelihood that each image in a pair belongs to the “has a TNO” class. The classification model provides likelihoods p_1

and p_2 between 0 and 1 and the user then selects at what level of likelihood an image is classified as containing a TNO. The choice of likelihood values is dependent on the goals and intentions of the specific project. A project desiring to provide a very pure sample may require a threshold of $p > 0.99$, while a project attempting to detect as many TNOs as possible might have a lower threshold value. For example, if the threshold is 0.5, images with values above 0.5 will be classified as having a TNO but that may create a high false-positive rate. By adjusting the threshold like this, we decided how strictly we want to filter the proposed detection. Classification models and regression models were trained on each data set, and the models were stored separately. After that, by adjusting the thresholds of the models, how performance metrics such as recall and precision change were identified and recorded.

Recall and precision can be used as criteria for classifier performance. For example, if a classifier has a recall of 0.9 and a precision of 0.8, this means that the filtered data set will miss 10% of the moving sources and 20% of the predicted positives are actually negatives.

In reality, TNOs are scarce on the sky at the flux limits of this experiment (about 1 per CCD) and very few sub-images will have a real TNO. This results in the number of negatives classified as positive (false positives) being far greater than the number of true positives, unless the p threshold is very near to 1. This creates the problem that the number of false positives make TNO exploration practically impossible. From this sparseness of TNO, the necessity of another post-processing step emerged.

When the classification model declares ($p > 0.5$) that a moving source is present in the sub-image we pass the two sub-images to the regression model which provides estimates of the (x_1, y_1, x_2, y_2) position of the moving source (see Table 2.4 for estimates of the precision of these estimates). A distant solar system object, such as a TNO, will exhibit linearly sky motion over the short time period (a few hours) during which our data have been acquired (see Fig. 2.1 for the ability to identify linearly moving objects). Multiple source detection within a small area of sky are likely to be of the same object and the rates of motion inferred by grouping these source measurements should be consistent with expectations for a solar system body. By requiring that each source measurement in a group is consistent with expected linear motion, and rejecting from the group those sources that are inconsistent with this behaviour, we are able to list the potential moving objects.

In detail, the linear motion of TNOs were detected in the following steps. First, for every single source in the data set, a list of sources that may be grouped as a

single object, based on proximity on the sky, were built. This results in the grouping a few dozens nearby source measurements. A maximum-likelihood based fit to linear motion, with iterative rejection of those sources that are more than $4\text{-}\sigma$ from fit, was then performed. Once an acceptable fit was determined, those source measurements that are within 0.5 arc-seconds of the linear model were accepted as belonging to the same candidate object (see Figure 2.7). Candidate objects with more than three source measurements (a triplet) and that exhibited rates of sky motion of between 0.5 and 15 arc-seconds per hour, reasonable rates for solar system bodies between 10 au and 300 au when observed at opposition, were then selected for visual inspection.




Figure 2.7: Blue dots are regression-predicted positions of a group of sources that the classification model gave $p > 0.5$ probability of being a moving source. Yellow crosses are the predicted locations a moving source would be at, in each exposure, based on a maximum likelihood-based linear fit to the regressor-based locations. The linear model provides an excellent match to this group of sources which are then declared as moving object candidates.

The linear fitting steps above filter out many of the false positives from section 2.6.1 by only accepting sources that can be linearly connected into linked sources and with linear rates of motion that are consistent with expectations for solar system bodies. We call the accepted linked sources a candidate “track”. The maximum number of source measurements in a single track is 44, as there were 44 images in our input time series.

After this fitting process, there is still a possibility that non-SSO images are passed

as a SSO candidate. This occurs when there is a defect, such as saturation bleed, in the image and this defect moves slightly from exposure to exposure in a way that mimics an SSO. Hence, a visual inspection is made for the final candidate selection. The inspection consists of examining the scatter plot of the candidate track (For example, Figure 2.7) and confirming or rejecting the existence of a moving object in the image in the associated sub-images.

After a brief visual inspection, we found that the MAE of the regression model results for faint objects ($m \geq 23$) is too large to enable accurate linear fitting. Moreover, while some TNOs were successfully detected, we could only find very few real SSOs or TNOs because of confusion caused by densely populated planted sources. Therefore, we explored another way to exploit the CNNs.

2.6.2 Approach 2: Scoring Approach

This time, we attempted to use only the classification model. The output of the classification model is (p_1, p_2) . For each series of sub-images at the same coordinates, we gave them a score by summing up p_1 or p_2 . For example, there are 44×43 p 's from measuring (1st=1st exposure, 2nd=2nd exposure), (1st, 3rd), (1st, 4th), ..., (2nd, 3rd), (2nd, 4th), ..., (43th, 44th) image pairs with the classification model at fixed coordinates (x_1, y_1) . We excluded p 's less than 0.99, and hence most of false positives can be avoided. Additionally, the value of p 's were linearly mapped using the function: $\tilde{p} := (x - 0.99) \times 100$. This function transformed the values from 0.99 to 0, from 0.991 to 0.1, from 0.992 to 0.2, ..., from 0.999 to 0.9, and from 1 to 1. We sum the mapped \tilde{p} 's for each time-series sub-image sets. This summed probability value, or a score, expresses a likelihood of having a moving source in the sub-image set. In other words, the score is the likelihood to detect a moving source in the sub-images fixed at the same sky coordinate throughout the 44 images. We ran this scoring on all sub-image sets, and excluded sets with a low score.

We also use the average of the original (unscaled) values of p for each exposure of a sub-image to filter the candidate list. Each exposure sub-image is classified 43 separate times (once for each time that exposure appears in a sub-image pair). The classification value, P , assigned to a given exposure sub-image is the average of the 43 classification values for that exposure sub-image. For example, the sub-image for exposure i is paired with the 43 other exposure sub-images and presented to the CNN classifier. For each pair of sub-images, the classifier returns 2 values, one for each

of the sub-images. We then take the value of P as the average of all values of p 's that were returned for a given exposure (43 values). We observed that the P rarely goes above 0.85 if a moving object is not clearly present during visual inspection. Therefore, we excluded sub-image series that do not have a P larger than 0.85. After this step, only series that have a clear P spike remained. In short, a time-series sub-image set should pass two conditions: $\sum^{all\ pairs} \tilde{p} \geq 80$ and $\exists P \in A : a \geq 0.85$ to be considered as a candidate. This additional post-processing step significantly decreased the number of false positives.

Candidates with more than a certain score were very likely to include a visible and obvious moving object. Therefore, we might be able to claim that very highly scored sub-image series (For example, $score \geq 1000$) include a SSO, without additional visual inspection. Note that the maximum possible score was $44 \times 43 = 1892$. This automated feature can be very useful for a future survey as the detections of the survey can be reported rapidly.

Moreover, sub-image series with a high score (For example, $80 \leq score < 1000$) can be further visually inspected to determine if they contain SSOs. For the inspection, we made GIF animations with stretched and normalized pixel values and chose the time between each frame to be 125ms. As additional information to the image, the inspector was given the number of the current image frame (1 to 44) and the p average value attached to each sub-image series. In this way, the inspector can estimate when the moving source appears during the inspection. This additional information is quite helpful when the moving light source is faint.

2.7 Results and Discussion

We present the results from two post-processing approaches, show the peak detection efficacy, and finally show the real SSOs we found.




Figure 2.8: Analysis of linear fitting approach on the C08132134 test set. (Left) Number of false positive tracks as a function of the minimum length of the track (n) to accept a track as a detection candidate. The vertical lines and end points are at $n=3$, $n=6$, $n=11$, and $n=22$. Setting the minimum track length requirement to these values would result in 14435, 1514, 387, and 98 candidates to vet, respectively. (Right) A completeness graph showing artificial TNO retrieval rates as a function of magnitude for various track length. “Found” objects did not go through the human inspection, therefore the actual retrieval rate will be slightly lower. Figure 2.4 is based on each single image pair classification result while the completeness shown here is for the final candidates after classification, and linear fitting.

2.7.1 Results with the Linear Fitting Approach

This subsection describes the results from the linear fitting approach. Please refer to the section 2.6.1 for the description of this approach.

The number of planted TNOs found strongly depended on the magnitude of objects. On the whole, there were 2,910 planted sources added to C08, C13, C21, and C34. In our final candidate list we retrieved 1,454 of these sources, with classification $p > 0.5$ and track length $n > 3$. The ratio represents only half the success in retrieving planted sources in terms of the number of retrieved TNOs. However, we detected a sharp decrease in detection ratio as the magnitude of TNOs increases. If we divide the planted objects into magnitude bins and then analyze how many of

them survived through the classification model and post-processing algorithms, we observed that around 80% of objects were detected when they were brighter than the 24th magnitude. However, the detection rate dropped rapidly after the magnitude of 24.5 (Figure 2.8, left side). We can see that this method is quite effective at finding moving objects in astronomical images when the moving objects are sufficiently bright. TNOs dimmer than the 25th magnitude were rarely retrieved as the SNR is low after the magnitude (refer Fig. 2.3). This behaviour is similar to the single image pair result shown on Fig 2.4.

The longer the required track length, n , the more secure the detection but the fewer number of detection were made. Without any post-processing methods, the number of sub-image pairs is $\sim 960,000$ at $p > 0.9$ and $\sim 1,200,000$ at $p > 0.5$. This would result in an impossible number of candidates for visual inspection. The number of candidates diminishes quickly as we require n , the length of the tracks to be 3 or greater. We marked that a track was correctly classified as a positive in this post-processing algorithm when the actual position of a TNO was inside 3 pixels of the measured position of a TNO, and detection occurs more than or equal to n times. With a requirement of a minimum length $n=3$, we retrieved 1,454 TNOs out of 2910 TNOs. With a minimum requirement of a length $n=6$, we retrieved 1,220 TNOs. With a minimum requirement of a length $n=11$, we retrieved 859 TNOs. With a minimum requirement of a length $n=22$, we retrieved 240 TNOs (Figure 2.8, right side). The longer tracks enable confidence in the detection, but at a large cost of missing many potential TNOs.

To determine the right number of tracks to be inspected, we analyzed the relationship between the length of tracks and the number of tracks for false positive cases. On the left side of the figure 2.8, the number of tracks is very large with the minimum requirement of $n=3$. However, the number decreases rapidly as we increase the minimum length. For instance, With $n=6$, there were ten times fewer candidates compared to $n=3$. We selected $n=11$ as the practical requirement for further inspection. At that point, the number of false positives decreases to below 400 candidates, which translates to approximately 100 candidates per CCD. Requiring $n=22$ does not yield good results at finding planted sources, even though inspecting candidates will be easier since the number of false positives per CCD is around 25.

With the $n=11$ minimum length requirement chosen, this method was effective to find artificial sources up to $m=23.9$ as it achieves approximately “50% retrieval

rate” at the magnitude. For our data set, it corresponds to $\text{SNR}=7.2$. Note that the 50% retrieval rate is not the absolute 50%. Because stars and galaxies block some parts of the skies, some moving sources are intrinsically impossible to find (refer subsection 2.7.3). This performance was not as good as we expected from utilizing a deep learning model, even before the visual inspection.

Upon inspecting false positive candidates manually with $n=11$ requirement or $n=22$ requirement, most of those tracks were definitely not non-planted SSOs, but rather caused by noises and wrong measurements on the positions. Only very few of them included an SSO, and only one of them was a very slow object (approximately 3 arcsecond/hour).

2.7.2 High Precision Detection with the Probability Cumulative Approach




Figure 2.9: Analysis of scoring approach on the C08132134 test set (Left) Number of false positive sub-image series as a function of the minimum score of the sub-image series (n) to accept a sub-images series as a detection candidate. The vertical lines are at $s=80$ and $s=1000$. Setting the minimum score requirement to these values resulted in 110 and 1 false-positive candidates. (Right) A completeness graph showing artificial TNO retrieval rates as a function of magnitude for two kinds of score requirements and the manual inspection requirement. Figure 2.4 is based on each single image pair classification result while this completeness graph is to show to the final number of TNOs found after post-processing.

Here we describe the results from the scoring approach. Please refer to the section 2.6.2 for the description of this approach, and Fig. 2.9 for the overall performance of the approach. For this approach, we marked that a sub-image series was correctly classified as a positive in this scoring approach when a moving object was in any sub-image of the series.

First of all, sub-image series with a very high score ($s=1000$ requirement) are confirmed to always have a moving object in them for our test set. Out of 889 such sub-image series, 888 of them had an artificial TNO. The only false positive in the set was a non-planted real slow-moving object (See Figure 2.11). From this result,

candidates above a certain score can be indeed reported as slow SSOs without further inspection. This capability of automatic detection can lead us to a total automation of SSO searches in future.

The large majority of sub-image series with a score of 80 or higher were found to be true positives. Out of the 2085 sub-image series with high scores, 1975 were confirmed to contain an artificial TNO. It means 95% of GIFs produced with this method would have a moving source. This high precision makes the manual inspection worth the time and effort and can be recommended as the final step of this discovery process.

The numbers of false positives were analyzed, and we find that this approach is effective at filtering on false positives while retaining true positives. The number of false positives with the $s=80$ was 110 for the C08132134 data set, and this number 110 is similar to the 98 candidates given in the linear fitting approach with the length $n=22$ requirement, which was the most strict requirement (See the left side of the Figure 2.8 for a comparison). At the similar number of false positives, the number of successfully retrieved TNOs from the linear fitting approach was only 240, while the number of successfully retrieved TNOs from the scoring approach was 1390 for the $s=80$ requirement. It is a significant improvement in the recall while keeping the number of false positive similar.

This method was effective at finding moving sources up to $m=24.7$. For our data set, it corresponds to $\text{SNR}=3.4$.

There were several advantages of the scoring approach over the linear fitting approach. This method excludes more false positives than the linear fitting approach. post-processing steps were straightforward and took significantly less computing time. The implementation of this method does not depend on the regression model, and hence we only need to train and use one model. Furthermore, we can select the score threshold based on the research goal. For example, bright and slow objects can be automatically claimed as an SSO without manual inspection. Even when a moving object is dim, it is easier to spot the difference when the the moving object is shown in animations rather than inspecting each image separately. Moving objects in a certain range of exposure duration appear sharp in our vision. [84] With these advantages, we achieved the high precision and high magnitude limit described above.

2.7.3 Peak Detection Efficacy

Even with the brightest TNOs in a data set, the recall does not reach 100%. For example, in Figure 2.4, the recall is 80-90% for the 22nd magnitude objects at the classification threshold of $p=0.99$. The peak efficiency is determined by the level of crowding from stars and galaxies in the image. Figure 2.10 exhibits two common obstacles to source detection: bright sources and bad columns. These types of obstacles hinder finding a moving object from the sub-images. Therefore, even with the best detection process one does not anticipate detecting 100% of the moving sources present in a series of images.




Figure 2.10: Sub-images containing artificial sources that were not detected by our process. The orange circles indicate the position at which the artificial source was added to the image. (Left) A bright stationary object (a star) in the background is obscuring the moving object. (Right) Bad columns on the detector prevent the detection of the moving object. For these objects, the classification model assigns a low probability to the presence of a moving source. For those that pass the classification step, the regression model often fails to predict the correct position of the moving source, and the candidate is rejected during the tracking process.

2.7.4 Discovery of real SSOs




Figure 2.11: A sub-image series accepted as being a detected solar system object. The images in the grid are arranged in rows and columns, at the same sky coordinates, depicting the progression of time. The time interval between each image is 248 seconds. The images are arranged in a sequential order from left to right, and from top to bottom. This non-planted candidate returned the highest score from our method. Pixel scale is 0.185 arcsecond per pixel. Based on the rate of motion of ~ 3 arcsecond per hour, this object is likely to be a Centaur or a TNO.




Figure 2.12: Sub-image series examples accepted as being detected SSOs. The images are arranged in sequential order from left to right. The first row: A dim oval-shaped light source moves from the center-left to the center-right almost horizontally. The moving object is easier to spot when images are made into an animation. The second row: A bright oval-shaped light source moves from the center-left to the lower right. Based on their rate of motion, these objects are likely to be asteroids.

The visual inspection with the scoring approach was done for non-planted areas of our data sets to examine its capability to find real SSOs. Out of 835 highly scored sub-image series ($s \geq 80$), 424 sub-image series had a non-planted SSO that we could visibly and confidently confirm. Some of them appeared more than once in another sub-image series as they move across the sky, and the 424 sub-image series containing a moving source included 196 unique SSOs. The candidates that passed this visual inspection may then be confirmed by follow-up observations.

The majority of discovered objects had sky motion in excess of 15 arcseconds per hour, but a handful of slow moving objects, which could be TNOs were also found. Note that the low number of discoveries on slow objects is not because the method is more sensitive to finding fast objects; it is rather because asteroids in a particular magnitude range outnumber TNOs: The method was actually more sensitive to smaller and slower objects because they spent more time in the sub-image frame and resulted in very high scores. One of the slow-moving objects is shown on Figure 2.11. We also report examples of fast-moving objects on Figure 2.12.

The initial intention of this study was to find TNOs with deep learning techniques, but after post-processing steps we found that this method could also find other SSOs such as asteroids. The model was trained using image pairs that have small time (minimum of 4 minutes) and large time (maximum of 3 hours) offsets between the images. This enabled the model to be sensitive to objects with fast rates of motion

(sources well offset between two sub-images with a small time offset) and slower rates of motion (sources well offset between two sub-images with a large time offset). The fast non-planted moving sources also exhibited trailing, indicating that the final model was not particularly sensitive to the shape of the PSF used in creating the artificial sources.

2.8 Summary and Conclusion

This study applied CNN-based networks to detect SSOs in MegaCam/MegaPrime images of CFHT. For our test data set, we used a series of 44 exposures (each approximately 205 seconds) that had been acquired as part of a search for satellites of Saturn in July 2019. Artificial moving sources were added to the original images and used to train a variety of network architectures. Of the architectures examined, we found that CNNs and particularly ImageNet algorithms were particularly adept at classifying images as containing or not containing a moving source. There were 7 architectures compiled: the FCNN for comparison, the modified AlexNet, the modified VGG, ResNet50, ResNet50V2, MobileNet and MobileNetV2. Out of those, we determined that ResNet50V2 and MobileNet have the best performance for our study. We selected MobileNet for moving source detection because it has fewer parameters, making it lightweight.

We also found that MobileNet is more effective when trained using examples from multiple CCDs rather than a single CCD. Models constructed by training on a single CCD were ineffective at finding moving sources on other CCDs. This may be due to the training having latched onto the image shape, or some other characteristic specific to the CCD and the artificial source addition process, rather than object motion.

The sub-images pairs were further accessed by post-processing steps such as the linear fitting approach or the scoring approach. The linear fitting approach was an approach to group nearby moving object positions measured by the trained regression model, and remove outliers to build postprocessed candidates. The linear fitting approach was not very effective at retrieving planted TNOs without significant visual inspection. Alternatively, a scoring approach was developed to return a likelihood (score) of having a moving object in a sub-image series at a fixed sky position. When a score from a sub-image series was significantly high, exceeding 1000, the sub-image series could be claimed as an automatic detection without further manual inspection; otherwise, if the score was relatively high, exceeding 80, the sub-image series was

subjected to manual inspection for the presence of a moving source. The manual inspection was worth the time as 95% of the sub-image series classified as positives were true positives. The approach diminished the number of candidates substantially and effectively retrieved planted TNOs as low as SNR=3.6, and it was sensitive to SSO over a broad range of motion rates. Passing the image pairs into the classifier and analyzing scores, we discovered planted and non-planted SSOs in an automated manner.

Using this approach we have demonstrated that the two-channel MobileNet architecture can be effectively trained to enable the detection of moving sources in astronomical images. Using the process above resulted in the detection of slow and fast-moving solar system objects, even though the model was only trained using slow-moving sources. The future study can involve adapting this method to other data sets, and exploring other deep learning techniques to improve the results.

Chapter 3

Additional concepts for a CNN-based discovery system.

In this chapter, I present some additional analysis that explores the nature of our data set, machine learning, and concepts for future work. The section 3.1 describes an attempt to classify images with TNOs using only one image, and the section 3.2 summarizes the ideas to enhance the discovery method.

3.1 Discovery of a TNO with just one image: Why it works and why it should not

Dual-channel training refers to providing the model with two sub-images from two different times, and single-channel training refers to training the model with only a single sub-image (no time-series information). We set dual-channel training as the default setting for this research because we assumed at least two images are required to detect the difference between images. To test this hypothesis, we tested a single-channel classification process. The single-channel process is revealing of overfitting.

I found that, to my surprise, the validation step for the MobileNet architecture did not perform significantly worse in the single-channel case. There was no significant difference in performance between the single-channel training and dual-channel training on their validation sets. On the validation set, single-channel trained models could achieve approximately 0.9 of precision and recall. However, when examining the single-channel model using the test sets (where the model is trained on one CCD and then detects sources on a different CCD), the

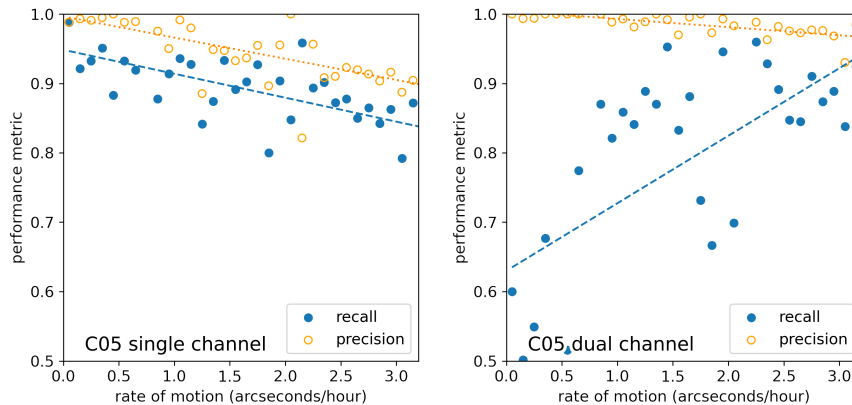


Figure 3.1: The graphs show the relationship between the performance of models and motion rates of TNOs. The left graph represents the case where the model is trained on each image separately, and the right graph represents the case where the model is trained on pairs of images that depict the movement of TNOs. As the rate of motion increases, the recall of the dual-channel model also increases. Unlike single-channel training, dual-channel training enables the model to recognize the movement of the source.

single-channel trained models failed. The average recall on the test set was only 0.636, and the performance was inadequate for TNO detection. Validation was found to be an insufficient criterion for model evaluation.

Single images of TNOs could detect the object if the sources are extended due to motion during the exposure. Such trailing might present differently in different CCDs (due to various in the PSF) but might still enable models to be trained to detect TNOs on single images. Perhaps the single-channel validation performance was achieved by the detection of image shape. Is the single-channel approach detecting ‘trailed’ points-sources?

The single-channel trained models did not correctly classify (detecting) images with relatively fast-moving objects, even when considering just validation sets. Figure 3.1 is the analysis of recall and precision with respect to the rate of motion of TNOs. On the left side of the Figure, the higher the rate of motion, the lower the performance. Suppose the single-channel model was detecting movement via source trailing. In that case, relatively fast-moving objects should have higher recall and precision, but the opposite was found to be true. The single-channel models cannot distinguish between moving and non-moving celestial bodies. The right plot of Figure 3.1 shows the precision for the dual-channel model is consistently higher than the single-channel training, and the recalls increased rapidly the faster the source

moves. TNOs with more evident movement are detected more easily, just like conventional methods. The dual-channel approach has resulted in training that detects object motion. The single-channel trained model appears to have been classified based on the precise shapes of point sources of the background stars compared to the fine structure present in shapes of the artificially generated TNO point sources. The single-channel model does not detect motion but does reveal subtle differences between the shapes of point sources on different CCDs. In conclusion, using at least two time-series images was the minimum requirement of this method, and my analysis demonstrates that the dual-channel approach detects changes between the sub-images, becoming more effective as those changes become more significant.

3.2 Ideas to improve the discovery methods

Images should be labelled correctly When an artificial source appears at the edge of a sub-image, part of the light source goes beyond the sub-image and is seen on spatial adjacent sub-images. During the labelling of images for training, only the sub-image containing the centroid of the moving source was labelled as holding a moving source. Spatially adjacent images that might contain some fraction of the moving source signal were labelled as negatives. Incorrectly labelled training data can negatively impact the effectiveness of the resulting model. In addition, when the analysis process found these kinds of edge detection on “negative” labelled images, they became labelled as possible real TNOs (positively classified sub-images not labelled as containing an artificial TNO), increasing the number of TNO candidates to vet but not increasing the number of real detections. To mitigate this issue, I have excluded these sub-images during the vetting process, but this incorrect labelling results in a more significant human effort. Sub-images with a partial moving source should be labelled as positives from the first place or excluded from the training set to avoid confusion. An improvement in the labelling algorithm will be developed in future work.

Try other deep learning architectures While MobileNet and ResNet are solid architectures, others, such as EfficientNet and Inception, are available. However, I could not test those architectures because the sub-image pair shape does not match the architecture of the Keras application at this time. The documentation on the

application notes that the input shape on EfficientNet and Inception should be three channels, and the image height and width should be no smaller than 75 pixels. Even though those architectures do not guarantee higher performance on our data set, future work should examine more architectures from Keras or modify the available architectures to achieve higher performance.

Make data sets with more channels This idea could be the most important one. The experimental data set consists of 44 epochs, but the model only used two epochs per step; the model training is not based on the entire history of the moving object. Providing the complete (or at least a more extensive) set of epochs, i.e. more channels of sub-images, should enable this method to find even fainter TNOs at lower SNR per image. Giving all frames at once would be equivalent to a video classification and object segmentation problem with some mix of RNNs and CNNs. This approach could be as effective as the shift-and-stack approach but with lower computation demands.

Using trained models on images without artificial TNOs We had around 800 artificial TNOs in each 2048×4612 pixel image. A large number of artificial sources was beneficial for making training sets. Still, it harmed the ability to find real moving objects. It provided the possibility that some real moving sources in negatively labelled sub-image pairs could (weakly, as there are only very few real TNOs to be found in a given image) confuse the training. For example, suppose there is a real moving object in a sub-image and an artificial object. In that case, the regression model may measure the position of the artificial object, not the real one. Moreover, the measured positions of real moving sources are easily obscured by the many artificial sources (see Figure 2.1), impacting the ability to make SSO tracklets effectively. Therefore, future work should involve applying the trained models on clear images free of any artificial TNOs to make discoveries easier.

The algorithm improvements described above should be investigated before implementing this approach in a production system. However, I have already demonstrated (see Chapter 2) that the CNN approach is highly effective at detecting solar system objects in astronomical imaging.

Chapter 4

Conclusion

The goal of this research was to explore a machine-learning approach to the detection of TNOs. The approach does not rely on data engineering or feature extraction (data agnostic) while minimizing the number of false negatives and false positives. I examined a deep learning approach to the discovery of TNOs and compared it with previous catalogue search methods [45].

I explored a fully connected neural network and a selection of CNN architectures for the TNO detection and classification tasks. In this work, I confirmed that CNNs can be used to detect moving sources, TNOs, in astronomical images. The CNN model based on the MobileNet architecture filtered stationary objects and found TNOs as faint as $m_r \approx 25$ in 205-second CFHT MegaPrime r-band image pairs, reaching a peak detection efficiency ($\sim 85\%$) for sources brighter than $m > 23$ (see Figure 2.4). As seen using the ROC curve (Figure 2.6) each image pair achieved an AUC of 0.87 for bright sources ($m < 23$), and AUC of 0.74 for sources with $m < 25$. After post-processing of the CNN classification outputs, visual inspection of the candidate detections revealed a manageable fraction (less than 1 in 10) of false positives demonstrating the feasibility of CNNs for TNO discovery. My search of the data using the CNN approach revealed a number of previously undetected small solar system objects including some TNOs.

This research has implications for TNO discovery programs. The immediate significance of this research is that I found several solar system objects in our data set that had previously not been detected (although this data set had previously been searched using a shift-and-stack technique). This data-agnostic CNN-based search approach can be applied to survey data sets that have been acquired in a fashion that is inconsistent with classical catalogue-based approaches. One such

data set is the r-band imaging from the Canada-France Imaging Survey [85]. We can also enhance the training of the model using new data sets via transfer learning approaches [86].

A number of enhancements to the process may also be possible. There is a possibility that my CNN could be more effective if trained on image triplets, quadruplets or longer, as opposed to only image pairs. Given the training is based on the detection of motion within the image one might expect that if we gave more images, or even the whole set of 44 images at once, as an input training set, the CNN could be trained to be even more effective. The drawback of such an approach would be that the model would be less applicable to existing data collections that are unlikely to have such long imaging sequences. However, a sequence of three images is a common observing strategy. Therefore, I think it is worth training on three image time series instead of the current two image pairs. For a longer input, 3D CNN might be a better choice as it can learn from both spatial and temporal input [87].

Alternatively, different deep learning methods such as recurrent neural networks, long short-term memory, or convolutional recurrent neural networks [88, 89] could be tested for TNO searches. While CNN is more fixed to a certain dimension of input, the aforementioned methods can be trained and used on varying lengths of input. If those methods work well on TNO surveys, there is a possibility that the performance of a trained model increases and post-processing/vetting would no longer be necessary. These additional approaches will be pursued in future investigations.

Advancement in telescopes and cameras allows us to observe TNOs at an increasing rate, and therefore advancement in our software for the automated detection of TNOs is required to match with the growing number of detectable TNOs. The automated method should be able to detect TNOs as effectively as the conventional method, catalog-based searches, or better. The ideal automated method should be able to report candidates without any visual inspection. Less visual inspection is better, no visual inspection is the best!

Concurrent attempts are made to develop similar deep learning-based TNO discovery, but my approach is trying to discover TNOs in a fully automated way. For example, the DECam Ecliptic Exploration Project team combined the shift-and-stack method with CNN to achieve a limiting magnitude of $r \sim 26.3$ ($\text{SNR} \sim 5$), and have made a preliminary report of more than 2,000 candidate TNOs being detected in their data set [90]. Similarly, the New Horizons Team is searching

for their next encounter target using the shift-and-stack method with a CNN approach assisting with visual vetting, they have detected a larger density of TNOs around 80 au compared to previous estimations [91] but have not yet found an encounter target. The usage of deep learning methods, especially CNNs, is already making TNO discoveries easier and achieving detections at lower thresholds of S/N. Combining CNN and shift and stack method has been proven to be effective. However, we may find that feeding non-processed data sets does better than feeding shift and stack processed data. This thesis showed that TNO searches with a deep learning method are possible, and my future study will focus on discovering TNOs with the method.

Bibliography

- [1] Matthias J. Gruber, Bernard D. Gelman, and Charan Ranganath. States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron*, 84(2):486–496, 2014.
- [2] Celeste Kidd and Benjamin Y. Hayden. The psychology and neuroscience of curiosity. *Neuron*, 88(3):449–460, 2015.
- [3] J. K. Davies, J. McFarland, M. E. Bailey, B. G. Marsden, and W. H. Ip. The Early Development of Ideas Concerning the Transneptunian Region. In M. A. Barucci, H. Boehnhardt, D. P. Cruikshank, A. Morbidelli, and Renee Dotson, editors, *The Solar System Beyond Neptune*, page 11. University of Arizona Press, 2008.
- [4] J. M. Petit, J. J. Kavelaars, B. J. Gladman, R. L. Jones, J. Wm. Parker, A. Bieryla, C. Van Laerhoven, R. E. Pike, P. Nicholson, M. L. N. Ashby, and S. M. Lawler. The Canada-France Ecliptic Plane Survey (CFEPS)—High-latitude Component. *AJ*, 153(5):236, May 2017.
- [5] J. V. Scotti, R. S. McMillan, T. H. Bressi, E. J. Christensen, R. E. Hill, R. A. Kowalski, E. C. Beshore, G. J. Garradd, A. R. Gibbs, A. D. Grauer, S. M. Larson, R. H. McNaught, and G. V. Williams. 2007 DU112. *Minor Planet Electronic Circulars*, 2008-P59, August 2008.
- [6] Rosemary E. Pike, Jayatee Kanwar, Mike Alexandersen, Ying-Tung Chen, and Megan E. Schwamb. Characterizing the Discovery of a New Trans-Neptunian Object Binary in a Trailed Point-spread Function Search. *PSJ*, 2(4):159, August 2021.
- [7] B. Sicardy, J. L. Ortiz, M. Assafin, E. Jehin, A. Maury, E. Lellouch, R. Gil Hutton, F. Braga-Ribas, F. Colas, D. Hestroffer, J. Lecacheux, F. Roques,

- P. Santos-Sanz, T. Widemann, N. Morales, R. Duffard, A. Thirouin, A. J. Castro-Tirado, M. Jelínek, P. Kubánek, A. Sota, R. Sánchez-Ramírez, A. H. Andrei, J. I. B. Camargo, D. N. da Silva Neto, A. Ramos Gomes, R. Vieira Martins, M. Gillon, J. Manfroid, G. P. Tozzi, C. Harlinton, S. Saravia, R. Behrend, S. Mottola, E. García Melendo, V. Peris, J. Fabregat, J. M. Madiedo, L. Cuesta, M. T. Eibe, A. Ullán, F. Organero, S. Pastor, J. A. de Los Reyes, S. Pedraz, A. Castro, I. de La Cueva, G. Muler, I. A. Steele, M. Cebrián, P. Montañés-Rodríguez, A. Oscoz, D. Weaver, C. Jacques, W. J. B. Corradi, F. P. Santos, W. Reis, A. Milone, M. Emilio, L. Gutiérrez, R. Vázquez, and H. Hernández-Toledo. A Pluto-like radius and a high albedo for the dwarf planet Eris from an occultation. *Nature*, 478(7370):493–496, October 2011.
- [8] M. J. Brucker, W. M. Grundy, J. A. Stansberry, J. R. Spencer, S. S. Sheppard, E. I. Chiang, and M. W. Buie. High albedos of low inclination Classical Kuiper belt objects. *Icarus*, 201(1):284–294, May 2009.
- [9] A. Farkas-Takács, Cs. Kiss, E. Vilenius, G. Marton, T. G. Müller, M. Mommert, J. Stansberry, E. Lellouch, P. Lacerda, and A. Pál. “TNOs are Cool”: A survey of the trans-Neptunian region. XV. Physical characteristics of 23 resonant trans-Neptunian and scattered disk objects. *AAP*, 638:A23, June 2020.
- [10] W. M. Grundy, K. S. Noll, and D. C. Stephens. Diverse albedos of small trans-neptunian objects. *Icarus*, 176(1):184–191, July 2005.
- [11] H. Boehnhardt, D. Schulz, S. Protopapa, and C. Götz. Photometry of Transneptunian Objects for the Herschel Key Program ‘TNOs are Cool’. *Earth Moon and Planets*, 114(1-2):35–57, November 2014.
- [12] H. L. Johnson and A. J. Gardiner. The Magnitude and Color of Mars During the 1954 Opposition. *PASP*, 67(395):74, April 1955.
- [13] C. Snodgrass, B. Carry, C. Dumas, and O. Hainaut. Characterisation of candidate members of (136108) Haumea’s family. *AAP*, 511:A72, February 2010.
- [14] W. M. Grundy, R. P. Binzel, B. J. Buratti, J. C. Cook, D. P. Cruikshank, C. M. Dalle Ore, A. M. Earle, K. Ennico, C. J. A. Howett, A. W. Lunsford, C. B. Olkin, A. H. Parker, S. Philippe, S. Protopapa, E. Quirico, D. C. Reuter, B. Schmitt, K. N. Singer, A. J. Verbiscer, R. A. Beyer, M. W. Buie, A. F. Cheng,

- D. E. Jennings, I. R. Linscott, J. Wm. Parker, P. M. Schenk, J. R. Spencer, J. A. Stansberry, S. A. Stern, H. B. Throop, C. C. C. Tsang, H. A. Weaver, G. E. Weigle, and L. A. Young. Surface compositions across Pluto and Charon. *Science*, 351(6279):aad9189, March 2016.
- [15] G. Randall Gladstone and Leslie A. Young. New Horizons Observations of the Atmosphere of Pluto. *Annual Review of Earth and Planetary Sciences*, 47:119–140, May 2019.
- [16] W. M. Grundy, M. K. Bird, D. T. Britt, J. C. Cook, D. P. Cruikshank, C. J. A. Howett, S. Krijt, I. R. Linscott, C. B. Olkin, A. H. Parker, S. Protopapa, M. Ruard, O. M. Umurhan, L. A. Young, C. M. Dalle Ore, J. J. Kavelaars, J. T. Keane, Y. J. Pendleton, S. B. Porter, F. Scipioni, J. R. Spencer, S. A. Stern, A. J. Verbiscer, H. A. Weaver, R. P. Binzel, M. W. Buie, B. J. Buratti, A. Cheng, A. M. Earle, H. A. Elliott, L. Gabasova, G. R. Gladstone, M. E. Hill, M. Horanyi, D. E. Jennings, A. W. Lunsford, D. J. McComas, W. B. McKinnon, R. L. McNutt, J. M. Moore, J. W. Parker, E. Quirico, D. C. Reuter, P. M. Schenk, B. Schmitt, M. R. Showalter, K. N. Singer, G. E. Weigle, and A. M. Zangari. Color, composition, and thermal environment of Kuiper Belt object (486958) Arrokoth. *Science*, 367(6481):aay3705, February 2020.
- [17] J. M. Petit, J. J. Kavelaars, B. J. Gladman, J. L. Margot, P. D. Nicholson, R. L. Jones, J. Wm. Parker, M. L. N. Ashby, A. Campo Bagatin, P. Benavidez, J. Coffey, P. Rousselot, O. Mousis, and P. A. Taylor. The Extreme Kuiper Belt Binary 2001 QW₃₂₂. *Science*, 322(5900):432, October 2008.
- [18] S. D. Benecchi, K. S. Noll, W. M. Grundy, and H. F. Levison. (47171) 1999 TC₃₆, A transneptunian triple. *Icarus*, 207(2):978–991, June 2010.
- [19] Charles H. Lineweaver and Marc Norman. The Potato Radius: a Lower Minimum Size for Dwarf Planets. *arXiv e-prints*, page arXiv:1004.1091, April 2010.
- [20] B. Gladman, M. Holman, T. Grav, J. Kavelaars, P. Nicholson, K. Aksnes, and J. M. Petit. Evidence for an Extended Scattered Disk. *Icarus*, 157(2):269–279, June 2002.
- [21] A. C. Becker, K. S. Arraki, A. Rest, W. M. Wood-Vasey, and B. G. Marsden. 2004 VN112. *Minor Planet Electronic Circulars*, 2007-S29, September 2007.

- [22] B. Gladman, J. Kavelaars, L. Allen, J. M. Petit, C. Vanlaerhoven, L. Jones, P. Smith, P. Nicholson, J. L. Margot, I. Smith, R. Jacobson, M. Brozovic, S. Greenstreet, C. Leung, S. Lawler, J. Parker, A. Bieryla, A. Parker, and G. V. Williams. Eleven New Tnos. *Minor Planet Electronic Circulars*, 2011-O40, July 2011.
- [23] Ron Cowen. A comet's odd orbit hints at hidden planet, April 2001.
- [24] Rodney Gomes. The Common Origin of the High Inclination TNO's. *Earth Moon and Planets*, 92(1):29–42, June 2003.
- [25] J. C. Becker, T. Khain, S. J. Hamilton, F. C. Adams, D. W. Gerdes, L. Zullo, K. Franson, S. Millholland, G. M. Bernstein, M. Sako, P. Bernardinelli, K. Napier, L. Markwardt, Hsing Wen Lin, W. Wester, F. B. Abdalla, S. Allam, J. Annis, S. Avila, E. Bertin, D. Brooks, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, C. E. Cunha, C. B. D'Andrea, L. N. da Costa, C. Davis, J. De Vicente, H. T. Diehl, P. Doel, T. F. Eifler, B. Flaugher, P. Fosalba, J. Frieman, J. García-Bellido, E. Gaztanaga, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, W. G. Hartley, D. L. Hollowood, K. Honscheid, D. J. James, K. Kuehn, N. Kuropatkin, M. A. G. Maia, M. March, J. L. Marshall, F. Menanteau, R. Miquel, R. L. C. Ogando, A. A. Plazas, E. Sanchez, V. Scarpine, R. Schindler, I. Sevilla-Noarbe, M. Smith, R. C. Smith, M. Soares-Santos, F. Sobreira, E. Suchyta, M. E. C. Swanson, A. R. Walker, and DES Collaboration. Discovery and Dynamical Analysis of an Extreme Trans-Neptunian Object with a High Orbital Inclination. *AJ*, 156(2):81, August 2018.
- [26] Shigeru Ida, John Larwood, and Andreas Burkert. Evidence for Early Stellar Encounters in the Orbital Distribution of Edgeworth-Kuiper Belt Objects. *APJ*, 528(1):351–356, January 2000.
- [27] Scott J. Kenyon and Benjamin C. Bromley. Stellar encounters as the origin of distant Solar System objects in highly eccentric orbits. *Nature*, 432(7017):598–602, December 2004.
- [28] Raúl de la Fuente Marcos and Carlos de la Fuente Marcos. An Update on the Future Flyby of Gliese 710 to the Solar System Using Gaia EDR3: Slightly Closer and a Tad Later than Previous Estimates. *Research Notes of the American Astronomical Society*, 4(12):222, December 2020.

- [29] Brett Gladman and Collin Chan. Production of the Extended Scattered Disk by Rogue Planets. *ApJL*, 643(2):L135–L138, June 2006.
- [30] Yukun Huang, Brett Gladman, Matthew Beaudoin, and Kevin Zhang. A Rogue Planet Helps to Populate the Distant Kuiper Belt. *ApJL*, 938(2):L23, October 2022.
- [31] Naoki Koshimoto, Takahiro Sumi, David P. Bennett, Valerio Bozza, Przemek Mróz, Andrzej Udalski, Nicholas J. Rattenbury, Fumio Abe, Richard Barry, Aparna Bhattacharya, Ian A. Bond, Hirosane Fujii, Akihiko Fukui, Ryusei Hamada, Yuki Hirao, Stela Ishitani Silva, Yoshitaka Itow, Rintaro Kirikawa, Iona Kondo, Yutaka Matsubara, Shota Miyazaki, Yasushi Muraki, Greg Olmschenk, Clément Ranc, Yuki Satoh, Daisuke Suzuki, Mio Tomoyoshi, Paul J. Tristram, Aikaterini Vandorou, Hibiki Yama, and Kansuke Yamashita. Terrestrial and Neptune mass free-floating planet candidates from the MOA-II 9-year Galactic Bulge survey. *arXiv e-prints*, page arXiv:2303.08279, March 2023.
- [32] Konstantin Batygin, Fred C. Adams, Michael E. Brown, and Juliette C. Becker. The planet nine hypothesis. *Physrep*, 805:1–53, May 2019.
- [33] Chadwick A. Trujillo and Scott S. Sheppard. A Sedna-like body with a perihelion of 80 astronomical units. *Nature*, 507(7493):471–474, March 2014.
- [34] Rodney S. Gomes, Jean S. Soares, and Ramon Brassier. The observation of large semi-major axis Centaurs: Testing for the signature of a planetary-mass solar companion. *Icarus*, 258:37–49, September 2015.
- [35] Michael E. Brown and Konstantin Batygin. Observational Constraints on the Orbit and Location of Planet Nine in the Outer Solar System. *ApJL*, 824(2):L23, June 2016.
- [36] Konstantin Batygin and Michael E. Brown. Evidence for a Distant Giant Planet in the Solar System. *AJ*, 151(2):22, February 2016.
- [37] Cory Shankman, J. J. Kavelaars, Michele T. Bannister, Brett J. Gladman, Samantha M. Lawler, Ying-Tung Chen, Marian Jakubik, Nathan Kaib, Mike Alexandersen, Stephen D. J. Gwyn, Jean-Marc Petit, and Kathryn Volk. OSSOS. VI. Striking Biases in the Detection of Large Semimajor Axis Trans-Neptunian Objects. *AJ*, 154(2):50, August 2017.

- [38] Michael E. Brown and Konstantin Batygin. The Orbit of Planet Nine. *AJ*, 162(5):219, November 2021.
- [39] Patryk Sofia Lykawka and Takashi Ito. Is There an Earth-like Planet in the Distant Kuiper Belt? *AJ*, 166(3):118, September 2023.
- [40] Katherine Jones-Smith and Harsh Mathur. Modified Newtonian Dynamics as an Alternative to the Planet Nine Hypothesis. *arXiv e-prints*, page arXiv:2304.00576, April 2023.
- [41] Wesley Fraser, Samantha Lawler, Edward Ashton, Ying-Tung Chen, Yukun Huang, Brett Gladman, J. Kavelaars, Jean-Marc Petit, Lowell Peltier, Rosemary Pike, Mike Alexandersen, Daniel Hestroffer, Benoit Noyelles, Chan-Kao Chang, Shiang-Yu Wang, Andrew Connolly, Bryce Kalmbach, Marielle Eduardo, Mario Juric, and Stephen Gwyn. The Classical and Large-a Distant Solar SYstem (CLASSY) Survey. In *AAS/Division for Planetary Sciences Meeting Abstracts*, volume 54 of *AAS/Division for Planetary Sciences Meeting Abstracts*, page 414.01, December 2022.
- [42] Pedro H. Bernardinelli, Gary M. Bernstein, Masao Sako, Tongtian Liu, William R. Saunders, Tali Khain, Hsing Wen Lin, David W. Gerdes, Dillon Brout, Fred C. Adams, Matthew Belyakov, Aditya Inada Somasundaram, Lakshay Sharma, Jennifer Locke, Kyle Franson, Juliette C. Becker, Kevin Napier, Larissa Markwardt, James Annis, T. M. C. Abbott, S. Avila, D. Brooks, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, F. J. Castander, L. N. da Costa, J. De Vicente, S. Desai, H. T. Diehl, P. Doel, S. Everett, B. Flaugher, J. García-Bellido, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, D. L. Hollowood, D. J. James, M. W. G. Johnson, M. D. Johnson, E. Krause, N. Kuropatkin, M. A. G. Maia, M. March, R. Miquel, F. Paz-Chinchón, A. A. Plazas, A. K. Romer, E. S. Rykoff, C. Sánchez, E. Sanchez, V. Scarpine, S. Serrano, I. Sevilla-Noarbe, M. Smith, F. Sobreira, E. Suchyta, M. E. C. Swanson, G. Tarle, A. R. Walker, W. Wester, Y. Zhang, and DES Collaboration. Trans-Neptunian Objects Found in the First Four Years of the Dark Energy Survey. *ApJS*, 247(1):32, March 2020.
- [43] Malena Rice and Gregory Laughlin. Exploring Trans-Neptunian Space with TESS: A Targeted Shift-stacking Search for Planet Nine and Distant TNOs in the Galactic Plane. *PSJ*, 1(3):81, December 2020.

- [44] Vera C. Rubin Observatory LSST Solar System Science Collaboration, R. Lynne Jones, Michelle T. Bannister, Bryce T. Bolin, Colin Orion Chandler, Steven R. Chesley, Siegfried Eggl, Sarah Greenstreet, Timothy R. Holt, Henry H. Hsieh, Zeljko Ivezić, Mario Jurić, Michael S. P. Kelley, Matthew M. Knight, Renu Malhotra, William J. Oldroyd, Gal Sarid, Megan E. Schwamb, Colin Snodgrass, Michael Solontoi, and David E. Trilling. The Scientific Impact of the Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST) for Solar System Science. *arXiv e-prints*, page arXiv:2009.07653, September 2020.
- [45] J. M. Petit, M. Holman, H. Scholl, J. Kavelaars, and B. Gladman. A highly automated moving object detection package. *Monthly Notices of the Royal Astronomical Society*, 347(2):471–480, January 2004.
- [46] E. Bertin and S. Arnouts. SExtractor: Software for source extraction. *AAPS*, 117:393–404, June 1996.
- [47] J. M. Petit, J. J. Kavelaars, B. J. Gladman, R. L. Jones, J. Wm. Parker, C. Van Laerhoven, P. Nicholson, G. Mars, P. Rousselot, O. Mousis, B. Marsden, A. Bieryla, M. Taylor, M. L. N. Ashby, P. Benavidez, A. Campo Bagatin, and G. Bernabeu. The Canada-France Ecliptic Plane Survey—Full Data Release: The Orbital Structure of the Kuiper Belt. *AJ*, 142(4):131, October 2011.
- [48] J. J. Kavelaars, Jean-Marc Petit, Brett Gladman, Michele T. Bannister, Mike Alexandersen, Ying-Tung Chen, Stephen D. J. Gwyn, and Kathryn Volk. OSSOS Finds an Exponential Cutoff in the Size Distribution of the Cold Classical Kuiper Belt. *APJL*, 920(2):L28, October 2021.
- [49] Jean-Marc Petit, Brett Gladman, J. J. Kavelaars, Michele T. Bannister, Mike Alexandersen, Kathryn Volk, and Ying-Tung Chen. The Hot Main Kuiper Belt Size Distribution from OSSOS. *APJL*, 947(1):L4, April 2023.
- [50] Z. W. Zhang, F. B. Bianco, M. J. Lehner, N. K. Coehlo, J. H. Wang, S. Mondal, C. Alcock, T. Axelrod, Y. I. Byun, W. P. Chen, K. H. Cook, R. Dave, I. de Pater, R. Porrata, D. W. Kim, S. K. King, T. Lee, H. C. Lin, J. J. Lissauer, S. L. Marshall, P. Protopapas, J. A. Rice, M. E. Schwamb, S. Y. Wang, and C. Y. Wen. First Results from the Taiwanese-American Occultation Survey (TAOS). *APJL*, 685(2):L157, October 2008.

- [51] Anita L. Cochran, Harold F. Levison, S. Alan Stern, and Martin J. Duncan. The Discovery of Halley-sized Kuiper Belt Objects Using the Hubble Space Telescope. *APJ*, 455:342, December 1995.
- [52] T. B. McCord and J. A. Westphal. Two-dimensional silicon vidicon astronomical photometer. *Applied Optics*, 11:522–526, January 1972.
- [53] Philip J Baldwin, Wendy C Evans, Damaris L Guevara, Jeff B Berner, Erica L Weir, and Allison C McCarthy. Nasa’s evolving ka-band network capabilities to meet mission demand. In *27th Ka and Broadband Communications Conference (Ka)*, 2022.
- [54] Laith Alzubaidi, Jinglan Zhang, Amjad J. Humaidi, Ayad Qasim Al-Dujaili, Ye Duan, Omran Al-Shamma, José I. Santamaría, Mohammed Abdulraheem Fadhel, Muthana Al-Amidie, and Laith Farhan. Review of deep learning: concepts, cnn architectures, challenges, applications, future directions. *Journal of Big Data*, 8, 2021.
- [55] Colin J. Burke, Patrick D. Aleo, Yu-Ching Chen, Xin Liu, John R. Peterson, Glenn H. Sembroski, and Joshua Yao-Yu Lin. Deblending and classifying astronomical sources with Mask R-CNN deep learning. *Monthly Notices of the Royal Astronomical Society*, 490(3):3952–3965, December 2019.
- [56] Md. Saikat Islam Khan, Anichur Rahman, Tanoy Debnath, Md. Razaul Karim, Mostofa Kamal Nasir, Shahab S. Band, Amir Mosavi, and Iman Dehzangi. Accurate brain tumor detection using deep convolutional neural network. *Computational and Structural Biotechnology Journal*, 20:4733–4745, 2022.
- [57] Long Chen, Shaobo Lin, Xiankai Lu, Dongpu Cao, Hangbin Wu, Chi Guo, Chun Liu, and Fei-Yue Wang. Deep neural network based vehicle and pedestrian detection for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 22(6):3234–3246, 2021.
- [58] Soad Almabdy and Lamiaa Elrefaei. Deep convolutional neural network-based approaches for face recognition. *Applied Sciences*, 9(20), 2019.
- [59] Izhar Wallach, Michael Dzamba, and Abraham Heifets. AtomNet: A Deep Convolutional Neural Network for Bioactivity Prediction in Structure-based Drug Discovery. *arXiv e-prints*, page arXiv:1510.02855, October 2015.

- [60] J. Luu, D. Jewitt, C. Trujillo, J. Chen, W. R. Brown, C. W. Hergenrother, T. B. Spahr, W. Offutt, and B. G. Marsden. 1997 CS29. *Minor Planet Electronic Circulars*, 1997-N08, July 1997.
- [61] Scott S. Sheppard, Chadwick A. Trujillo, David J. Tholen, and Nathan Kaib. A New High Perihelion Trans-Plutonian Inner Oort Cloud Object: 2015 TG387. *AJ*, 157(4):139, April 2019.
- [62] M. E. Brown, C. A. Trujillo, and D. Rabinowitz. 2003 EL_61, 2003 UB_313, and 2005 FY_9. *IAUCirc*, 8577:1, July 2005.
- [63] Michele T. Bannister, J. J. Kavelaars, Jean-Marc Petit, Brett J. Gladman, Stephen D. J. Gwyn, Ying-Tung Chen, Kathryn Volk, Mike Alexandersen, Susan D. Benecchi, Audrey Delsanti, Wesley C. Fraser, Mikael Granvik, Will M. Grundy, Aurélie Guilbert-Lepoutre, Daniel Hestroffer, Wing-Huen Ip, Marian Jakubik, R. Lynne Jones, Nathan Kaib, Catherine F. Kavelaars, Pedro Lacerda, Samantha Lawler, Matthew J. Lehner, Hsing Wen Lin, Tim Lister, Patryk Sofia Lykawka, Stephanie Monty, Michael Marsset, Ruth Murray-Clay, Keith S. Noll, Alex Parker, Rosemary E. Pike, Philippe Rousselot, David Rusk, Megan E. Schwamb, Cory Shankman, Bruno Sicardy, Pierre Vernazza, and Shiang-Yu Wang. The Outer Solar System Origins Survey. I. Design and First-quarter Discoveries. *AJ*, 152(3):70, September 2016.
- [64] Wesley Fraser, Mike Alexandersen, Megan E. Schwamb, Michaël Marsset, Rosemary E. Pike, J. J. Kavelaars, Michele T. Bannister, Susan Benecchi, and Audrey Delsanti. TRIPPy: Trailed Image Photometry in Python. *AJ*, 151(6):158, June 2016.
- [65] Rodrigo A. Ibata, Alan McConnachie, Jean-Charles Cuillandre, Nicholas Fantin, Misha Haywood, Nicolas F. Martin, Pierre Bergeron, Volker Beckmann, Edouard Bernard, Piercarlo Bonifacio, Elisabetta Caffau, Raymond Carlberg, Patrick Côté, Rémi Cabanac, Scott Chapman, Pierre-Alain Duc, Florence Durret, Benoît Famaey, Sébastien Fabbro, Stephen Gwyn, Francois Hammer, Vanessa Hill, Michael J. Hudson, Ariane Lançon, Geraint Lewis, Khyati Malhan, Paola di Matteo, Henry McCracken, Simona Mei, Yannick Mellier, Julio Navarro, Sandrine Pires, Chris Pritchet, Celine Reylé, Harvey Richer, Annie C. Robin, Rubén Sánchez-Janssen, Marcin Sawicki, Douglas Scott, Vivien Scottez, Kristine

- Spekkens, Else Starkenburg, Guillaume Thomas, and Kim Venn. The Canada-France Imaging Survey: First Results from the u-Band Component. *APJ*, 848(2):128, October 2017.
- [66] Yann LeCun, D Touresky, G Hinton, and T Sejnowski. A theoretical framework for back-propagation. In *Proceedings of the 1988 connectionist models summer school*, volume 1, pages 21–28, 1988.
- [67] Leon Bottou. *Stochastic Gradient Descent Tricks*, volume 7700 of *Lecture Notes in Computer Science (LNCS)*, pages 430–445. Springer, neural networks, tricks of the trade, reloaded edition, January 2012.
- [68] Kunihiko Fukushima, Sei Miyake, and Takayuki Ito. Neocognitron: A neural network model for a mechanism of visual pattern recognition. *IEEE transactions on systems, man, and cybernetics*, 13(5):826–834, 1983.
- [69] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [70] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *arXiv e-prints*, page arXiv:1409.0575, September 2014.
- [71] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications, 2017.
- [72] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.
- [73] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012.
- [74] Abien Fred Agarap. Deep learning using rectified linear units (relu). *CoRR*, 1803.08375, 2018.

- [75] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [76] Haohan Wang, Bhiksha Raj, and Eric P. Xing. On the origin of deep learning. *CoRR*, abs/1702.07800, 2017.
- [77] Benjamin Graham. Fractional max-pooling. *CoRR*, abs/1412.6071, 2014.
- [78] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [79] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *ArXiv*, abs/1611.03530, 2017.
- [80] Olivier Boulade, Laurent G. Vigroux, Xavier Charlot, Pierre Borgeaud, Pierre-Henri Carton, Jean de Kat, Jean Y. Rousse, Yannick Mellier, Pierre Gigan, David Crampton, and Christopher L. Morbey. Megacam: the next-generation wide-field imaging camera for CFHT. In Sandro D’Odorico, editor, *Optical Astronomical Instrumentation*, volume 3355 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 614–625, July 1998.
- [81] E. A. Magnier and J. C. Cuillandre. The Elixir System: Data Characterization and Calibration at the Canada-France-Hawaii Telescope. *PASP*, 116(819):449–464, May 2004.
- [82] James Bosch, Yusra AlSayyad, Robert Armstrong, Eric Bellm, Hsin-Fang Chiang, Siegfried Eggl, Krzysztof Findeisen, Merlin Fisher-Levine, Leanne P. Guy, Augustin Guyonnet, Željko Ivezić, Tim Jenness, Gábor Kovács, K. Simon Krughoff, Robert H. Lupton, Nate B. Lust, Lauren A. MacArthur, Joshua Meyers, Fred Moolekamp, Christopher B. Morrison, Timothy D. Morton, William O’Mullane, John K. Parejko, Andrés A. Plazas, Paul A. Price, Meredith L. Rawls, Sophie L. Reed, Pim Schellart, Colin T. Slater, Ian Sullivan, John D. Swinbank, Dan Taranu, Christopher Z. Waters, and W. M. Wood-Vasey. An Overview of the LSST Image Processing Pipelines. In Peter J. Teuben, Marc W. Pound, Brian A. Thomas, and Elizabeth M. Warner, editors, *Astronomical Data Analysis Software and Systems XXVII*, volume 523 of *Astronomical Society of the Pacific Conference Series*, page 521, October 2019.

- [83] Mingjie Sun, Jianguo Li, and Changshui Zhang. Extreme Value Preserving Networks. *arXiv e-prints*, page arXiv:2011.08367, November 2020.
- [84] D. C. Burr and M. J. Morgan. Motion deblurring in human vision. *Proceedings: Biological Sciences*, 264(1380):431–436, 1997.
- [85] Nicholas J. Fantin, Patrick Côté, Alan W. McConnachie, Pierre Bergeron, Jean-Charles Cuillandre, Stephen D. J. Gwyn, Rodrigo A. Ibata, Guillaume F. Thomas, Raymond G. Carlberg, Sébastien Fabbro, Misha Haywood, Ariane Lançon, Geraint F. Lewis, Khyati Malhan, Nicolas F. Martin, Julio F. Navarro, Douglas Scott, and Else Starkenburg. The Canada-France Imaging Survey: Reconstructing the Milky Way Star Formation History from Its White Dwarf Population. *APJ*, 887(2):148, December 2019.
- [86] Behnam Neyshabur, Hanie Sedghi, and Chiyuan Zhang. What is being transferred in transfer learning? *arXiv e-prints*, page arXiv:2008.11687, August 2020.
- [87] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):221–231, 2013.
- [88] Alex Sherstinsky. Fundamentals of Recurrent Neural Network (RNN) and Long Short-Term Memory (LSTM) network. *Physica D Nonlinear Phenomena*, 404:132306, March 2020.
- [89] Baoguang Shi, Xiang Bai, and Cong Yao. An End-to-End Trainable Neural Network for Image-based Sequence Recognition and Its Application to Scene Text Recognition. *arXiv e-prints*, page arXiv:1507.05717, July 2015.
- [90] Kevin Napier, Hsing-Wen Lin, David Gerdes, Fred Adams, Anna Simpson, Matthew Porter, Larissa Markwardt, Gabriel Gowman, Hayden Smotherman, Pedro Bernardinelli, Mario Juric, David Trilling, Ryder Strauss, Colin Chandler, Chadwick Trujillo, Cesar Fuentes, Matthew Holman, Scott Sheppard, and DEEP Collaboration. The Apparent and Absolute Magnitude Distributions of the Kuiper Belt from Single-Night Detections in DEEP. In *AAS/Division for Planetary Sciences Meeting Abstracts*, volume 54 of *AAS/Division for Planetary Sciences Meeting Abstracts*, page 501.06, December 2022.

- [91] W. C. Fraser, S. B. Porter, H. W. Lin, K. Napier, R. J. Spencer, Jj. Kavelaars, A. J. Verbiscer, F. Yoshida, T. Terai, T. Ito, D. Gerdes, S. D. Benecchi, S. A. Stern, S. Gwyn, M. W. Buie, L. Peltier, K. N. Singer, P. C. Brandy, New Horizons Lorri Team, and New Horizons Ggi Science Team. Approaches to Detecting Kuiper Belt Objects for NASA's New Horizons Extended Mission: Digging Into the Noise. In *LPI Contributions*, volume 2806 of *LPI Contributions*, page 2361, March 2023.