9th IAA Planetary Defense Conference – PDC 2025
5–9 May 2025, Stellenbosch, Cape Town, South Africa

IAA-PDC-25-04-162

Real-time Discovery of Near-Earth Objects via Accelerated Image Analysis with AI Methods

Szabolcs Velkei[a,*], László L. Kiss[b,*], Károly Vass[a], Krisztián Sárneczky[b]

[a]*Machine Intelligence Inc, H-2015 Szigetmonostor, Horánygyöngye 102/b, Hungary*
[b]*Konkoly Observatory, HUN-REN Research Centre for Astronomy and Earth Sciences, H-1121 Budapest, Konkoly Th.M. 15-17, Hungary*

Abstract

Konkoly Observatory is conducting the most successful NEO survey project in Europe with a total number of NEOs found in the past four years in excess of 250, with three imminent impactors discovered between 2022 and 2024. Recently, supported by the European Space Agency, we started the implementation of a new search technique that is using machine learning algorithms to accelerate real-time image analysis with the scope of finding extreme trailed images of the smallest and nearest NEOs passing by. We have created a custom deep-learning model that was trained on a large dataset of astronomical images and their associated annotations. In addition to the real observations from the Piszkéstető Mountain Station of the Konkoly Observatory, we have also created a huge synthetic photorealistic training dataset to improve the precision and accuracy of the neural network. As a result, the model successfully learnt to recognise patterns and features in the images that are indicative of NEOs and space debris. The main goal was to have an optimized deep learning model to perform this analysis in real-time, providing quick and reliable detection that is made possible by the AI-based robust image-artifact decomposition for false positive suppression. The outcome of this project is a service that can quickly and accurately detect NEOs and space debris on astronomical images, potentially increasing the number of discoveries and improving the speed and reliability of the discovery process. The system has been evaluated using a set of rigorous tests and is benchmarked against existing methods. We provide valuable insights into the feasibility of using deep learning techniques for this type of image analysis problem and will lay the groundwork for future work in this field.

Keywords: Near-Earth Objects, discovery, machine learning

1. Introduction

Recent decades have underscored the significance of identifying and tracking Near-Earth objects (NEOs), given their potential to pose substantial risks to our planet. NEOs are defined as objects whose orbits bring them within 1.3 astronomical units (au) of the Sun, enabling close approaches to Earth and, in rare cases, direct impacts.

The historical records highlight the destructive potential of NEO impacts. The Chicxulub impactor, linked to the Cretaceous–Paleogene extinction event, was a NEO estimated at approximately 10 km in diameter [1]. More recently, the 1908 Tunguska event in Siberia, attributed to an object roughly 50 meters in size, devastated over 2,000 square kilometers of forest, illustrating that even much smaller bodies can have severe regional consequences [2]. The latest 2013 Chelyabinsk meteor provides

*Szabolcs Velkei, László L. Kiss
Email addresses: szabolcs.velkei@mi.services (Szabolcs Velkei), kiss.laszlo@csfk.org (László L. Kiss), karoly.vass@mi.services (Károly Vass), sarneczky.krisztian@csfk.org (Krisztián Sárneczky)
[1]Machine Intelligence Inc, H-2015 Szigetmonostor, Horánygyöngye 102/b, Hungary
[2]Konkoly Observatory, HUN-REN Research Centre for Astronomy and Earth Sciences, H-1121 Budapest, Konkoly Th.M. 15-17, Hungary

a critical example of the risks posed by small-sized impactors. On February 15, 2013, an asteroid approximately 20 meters in diameter entered Earth's atmosphere over the Chelyabinsk region of Russia, exploding at an altitude of 23–26 kilometers with an estimated energy release of 440–500 kilotons of TNT equivalent [3]. The resulting shockwave shattered windows in over 7,200 buildings and caused injuries to more than 1,500 people, primarily from flying glass.

Comprehensive risk assessment and mitigation efforts require a complete inventory of NEOs and precise knowledge of their orbital parameters. While the largest NEOs have been extensively cataloged, the majority of smaller bodies - particularly those larger than 1 meter - remain undetected, with estimates suggesting that up to 99% of this population is still unknown [4]. The challenge arises from the fact that smaller NEOs are only observable during close approaches, when their apparent brightness and rapid sky motion briefly make them detectable.

During such high-proper-motion passages, NEOs can traverse the focal plane at rates of several degrees per hour, producing elongated streaks that may extend across hundreds of pixels in a single exposure. Systematic identification of these streaks within large volumes of astronomical imaging data is critical for discovering new NEOs and refining the orbits of known objects. Enhanced detection capabilities, particularly for faint and fast-moving NEOs, are therefore essential for advancing both planetary defense and our understanding of the Near-Earth environment.

2. Detection of asteroids

The detection of asteroids - slow or fast moving point shaped or linear trails left by NEOs in astronomical images - has traditionally relied on classical image processing techniques, particularly point- and line-detection algorithms. While these approaches can highlight candidates, they are prone to confusion with other artifacts of similar scale and are particularly sensitive to noise, which requires the use of very high thresholds in these algorithms.

In recent years, machine learning has begun to surpass traditional algorithms in the domain of astronomical image analysis, including streak detection. Some approaches, such as those described by [5], employ machine learning models not for direct detection but for classification of candidates preselected by conventional algorithms, thereby streamlining the vetting process and reducing the need for manual inspection. However, these models often cannot localize streaks or provide information on their morphology, requiring subsequent manual annotation and parameter estimation.

A more direct approach leverages convolutional neural networks (CNNs), which are inherently suited to extracting spatial patterns from image data. After appropriate training, CNNs can learn to differentiate between genuine asteroid streaks and other linear features, enabling automated and robust detection. This methodology has been demonstrated on a variety of datasets, including Hubble Space Telescope imagery [6] and simulated data from the Euclid mission [7]. Notably, recent work has shown that CNNs trained on synthetic streaks can achieve high detection efficiency even for faint and short trails.

The present study introduces a novel machine learning pipeline designed to enhance the detection capabilities of very faint streaks of NEOs - including both natural objects and artificial ones (satellites, rocket bodies, space debris) - within a real-time detection framework. The pipeline is purposefully optimized to deliver high performance in both cloud-based environments and edge architectures, enabling onboard processing directly on satellites or spaceborne platforms without the dependency on ground-based data transmission, thereby supporting autonomous decision-making in orbital environments.

3. The temporal streak detection

Unlike conventional preprocessing approaches, which typically rely on a single image or a difference image focused on transient phenomena, the NEODetect system employs a preprocessing pipeline based on sequences of three images aligned via stellar registration. For both conceptual clarity and ease of visualization, the measurements are represented as three-channel, 16-bit BGR images, where the B, G, and R channels correspond to observations taken at epochs T-1, T0, and T+1, respectively. Beyond standard dark, flat, bias calibration and iterative median filtering of the background, no additional preprocessing steps are applied.

The principal objective is to enable the detection of streaks at the lowest possible signal-to-noise ratios (SNR). To avoid introducing interpolation artifacts, image alignment is performed using only pixel shifts. Inevitably, subpixel displacements during acquisition or variations in atmospheric seeing introduce differences in the images of otherwise static sources (e.g., stars, galaxies). In difference imaging, such discrepancies may persist and cannot be fully eliminated through subtraction. Consequently, a

neural network trained solely on difference images cannot account for the origin of these signals. In contrast, the NEODetect input data preserves all such information in its original form, thereby providing the neural network with the opportunity to learn and distinguish various phenomena arising during the imaging process (e.g., cosmic rays, CCD amplification artifacts, diffraction spikes, seeing variations).

Following semantic segmentation by the neural network, all regions across the three temporal frames that may correspond to NEO candidates are identified. The initial step is to merge candidate detections across the image plane axes. Due to the high-noise environment, detected streaks often appear as discontinuous fragments rather than contiguous features. These fragments are typically collinear and share a common orientation, as determined by line fitting. After merging, the system searches across the temporal axis for collinear, directionally consistent segments, taking into account both the acquisition times and the exposure fill factor.

If collinear candidate fragments are identified, share a consistent direction, and exhibit a coherent temporal sequence with the T0 segment always centrally positioned, the NEO candidate is added to a provisional list. The imposition of chronological alignment - as a constraint - results in the automatic exclusion of 99.2% of false positives: in the test dataset, 71,995,819 of 72,575,480 candidates were classified as pixel noise or cosmic ray artifacts. The remaining 579,661 candidates, however, still represent an excessive number of false positives. To further minimize this figure, an additional criterion called the *temporal consequence filter* is enforced: there must exist at least one subsequent triplet of images containing a consistent continuation of the candidate within the measurement sequence, i.e., an independent NEO candidate whose position, direction, and apparent speed are compatible with the preceding detection. Only detections that satisfy this condition are reported and stored in the database.

4. Biased pattern recognition and Pareidolia

Humans possess an inherent tendency to perceive patterns, even in the absence of genuine structure. While this trait may have conferred evolutionary advantages, in the context of artificial intelligence, it often manifests as the phenomenon of hallucination. If, however, the adverse effects, namely the false positive detections, can be effectively suppressed, this propensity can be leveraged to identify faint signals that would remain undetected. The temporal consequence filtering strategy described in the preceding section is well suited to this purpose. To encourage the neural network to adopt a form of computational pareidolia, it is necessary to include extremely faint signals as true positive samples in the training dataset, even when these are well below standard detection thresholds. Accordingly, the training database was constructed to contain a substantial number of signals with SNR values as low as 0.25, where SNR is defined as:

$$\mathrm{SNR} = 10 \log_{10} \left(\frac{\mu_{\mathrm{signal}}}{\sigma_{\mathrm{noise}}} \right) \tag{1}$$

5. Training Data

The images used for the training dataset are real K88 preprocessed images taken from 2022-2024. The details of these and the historic background of the Piszkéstető NEO survey are presented by Sárneczky et al. in our companion paper (2025, this volume) [8]. The size of the binned images are 5280x5280 pixels, equivalent to an angular area of 8.76 deg^2. The training dataset exhibits significant class imbalance due to the scarcity of authentic NEO samples. Two primary strategies were considered to facilitate convergence in the machine learning process: reducing the perceptual environmental data or augmenting the sample set. The former approach was discarded as it would compromise the neural network's capacity to detect imaging anomalies and artifacts and its robustness to typical noise levels. There are two ways to enrich the samples: Augmentation of existing samples, and creation of synthetic samples. To have full control over what the neural network can learn, we chose the latter solution, and a synthetic dataset was generated utilizing the terminal frame from each real observation sequence yielding 23,952 tiles (1280x1280) used the following parameters: apparent proper motion between 2″-450″/min, CCD readout time between 0.1-1.5 exposure time ratio, signal level 0.25-10 SNR, FWHM 0.7-2 binned pixels. This dataset was divided into two equal subsets to accommodate different apparent motion profiles:

1. High-speed NEOs ranging from 50″ to 450″/min, at a density of 676 objects per 25 tiles,
2. Slow-moving NEOs, with a higher density of 65,536 objects per 25-tile pack.

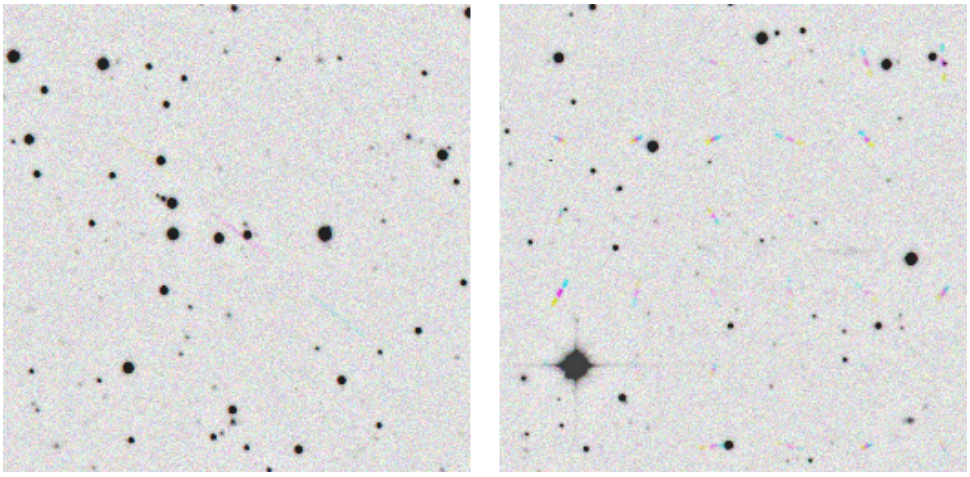

Figure 1: Inverted images of fast- and slow-moving synthetic samples in the temporal training dataset

This balanced approach resulted in the creation of approximately 323,000 fast-moving and 31.4 million slow-moving synthetic NEOs. Notably, NEOs with an SNR of 0.25 are not directly detectable in single images and require advanced software tracking techniques for identification.

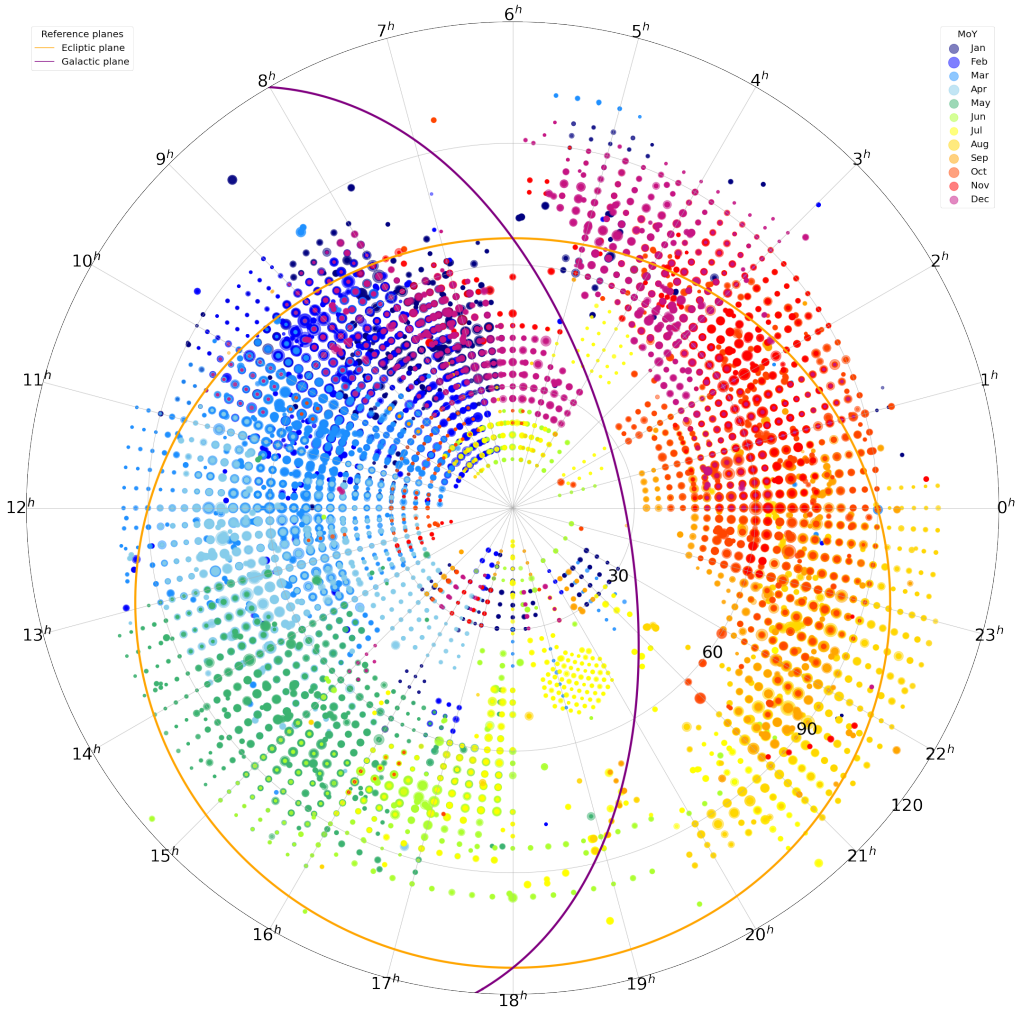

Figure 2: The celestial distribution of the synthetic training image tiles

6. Network architecture

A dedicated neural network was developed and trained using the unique temporal input, leveraging Machine Intelligence Designer (MInD) - the proprietary deep learning development environment of Machine Intelligence Zrt[3]. We utilized the environment's automated hyperparameter and structural search functionalities to identify the optimal architecture of our spatiotemporal U-Net network - which type is successfully used for different segmentation and detection in the past [9].

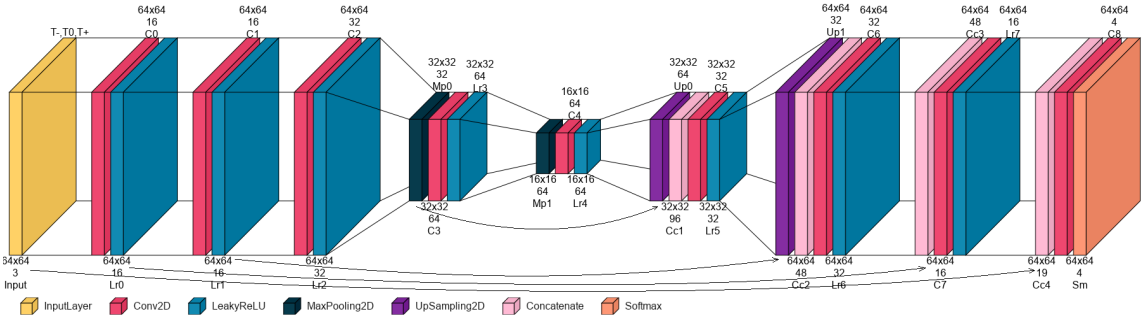

Figure 3: The NEODetect architecture

This search process, through appropriate weighting, seeks a compromise between robust solutions and computational efficiency. The objective was to maximize the detection of signals within the SNR 0.25-10 range while minimizing computational requirements, thereby ensuring both cost-effectiveness and potential future deployment in spaceborne environments on radiation-tolerant (RT) or radiation-hardened by design (RHBD) edge devices. To this end, we developed network variants with computational demands of 20, 100, and 500 Gops, each capable of processing generic 2k x 2k images.

The largest network variant, intended for terrestrial systems, is capable of running on both GPU and CPU platforms, with a typical processing speed of 1.3 seconds per megapixel per 4th Gen Intel Xeon Gold CPU core. On our test server, which operates with 8-core CPU Docker containers with 40 GB RAM, the processing of the K88 (Piszkéstető) 25-megapixel observations requires only 4 seconds, demonstrating that real-time detection is easily achievable.

The network was trained on 1088×1088 randomly cropped sections extracted from 1280×1280 tiles in the training dataset. The training regimen comprised three distinct phases:

1. Warmup phase: A 2-epoch initialization period with a low learning rate (1e-4) to stabilize early parameter updates.
2. Optimization phase: Utilization of the AMSGrad variant of the Adam stochastic optimizer over 20 epochs, with a learning rate of 0.005 and a batch size of 16. This phase prioritized rapid convergence while maintaining gradient stability.
3. Stabilization phase: A final 3-epoch refinement stage employing a reduced learning rate (1e-4) with decay to 1e-5, specifically targeting the low-SNR subset of the training data to enhance detection sensitivity for faint signals.

The phased approach balanced convergence speed with parameter stability, particularly critical for maintaining performance on low-SNR features.

7. Results

The results beyond the test dataset were summarized during a two-month beta test. The primary objective of this test was to evaluate the NEODetect service under operational conditions. To this end, we deployed an on-premise solution at one site, while the remaining solutions were deployed via cloud services for three additional locations. During this period, we enhanced the robustness of the runtime environment and developed several new features in response to user feedback. Our ongoing goal is to make the service accessible to any observer, thereby broadening participation in real-time NEO detection. Our exhaustive evaluation, encompassing both historical archives and contemporary test phases,

[3]Machine Intelligence Designer (MInD), https://mi.services/mind/

confirms that NEODetect has successfully fulfilled its core objectives. The system demonstrated exceptional efficiency in detecting objects within velocity ranges previously overlooked in NEO surveys. A significant ancillary outcome emerged in the form of robust space debris and artificial object detection. Statistical analyses revealed unanticipated strengths in the system's capability to identify other targets, suggesting its potential to unlock novel opportunities for enhancing space debris monitoring and cataloging efforts. As illustrated in Figure 4, NEODetect identified thousands of previously unidentified Earth-orbiting objects during the evaluation period. These findings underscore its dual utility in both planetary defense and space domain awareness, addressing critical gaps in current observational frameworks.

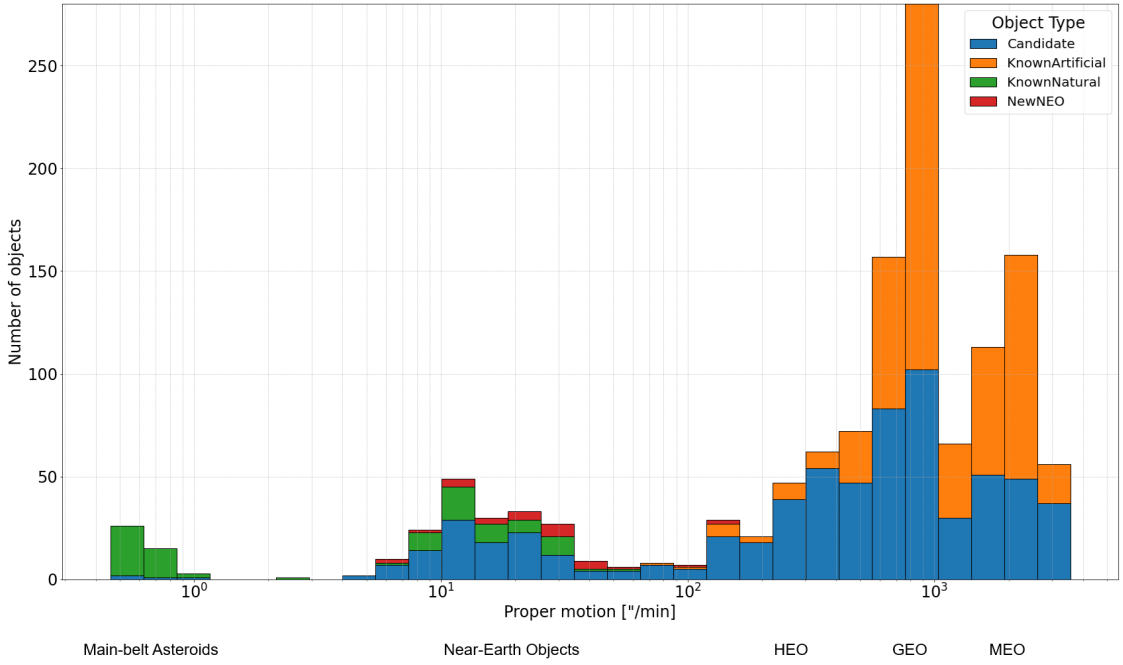

Figure 4: Histogram of the detected objects by their apparent proper motion.

Furthermore, due to weak signal training, the recognition rate is above 51.5% even for signals with an SNR strength of 0.3 and quickly approaches nearly 100%.

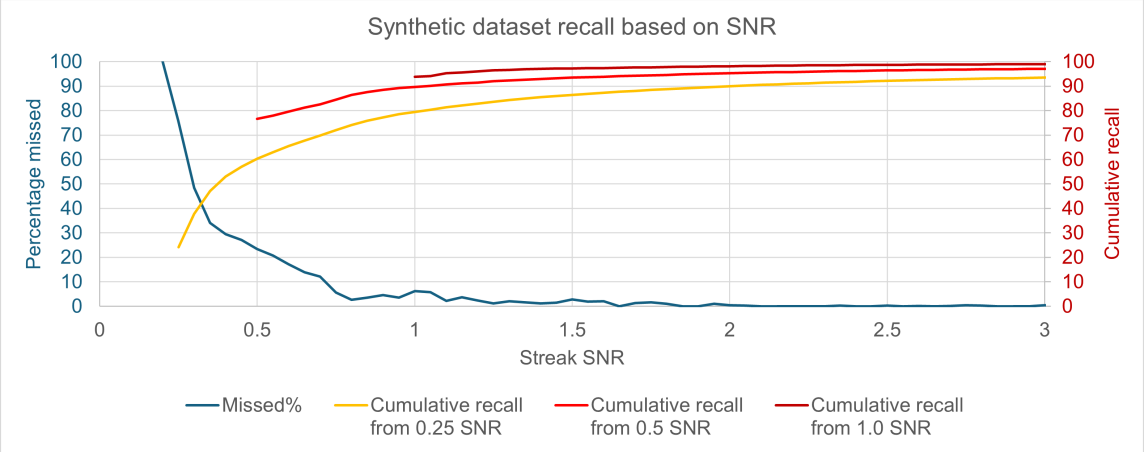

Figure 5: Test dataset synthetic NEO recovery rate based on the NEO candidates' SNR.

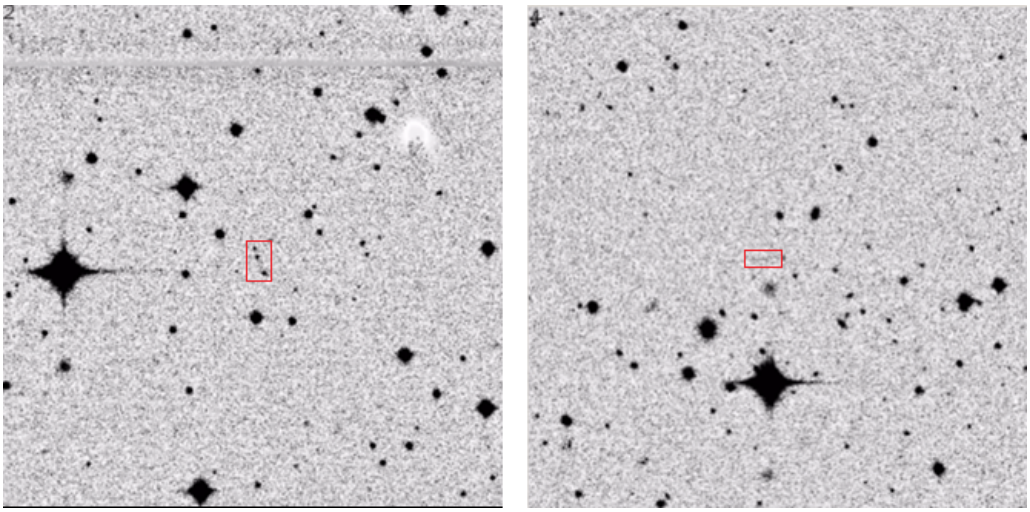

Figure 6: Examples of confirmed streaks from K88 Piszkéstető images. Left: a NEO overlapping with stars. Right: a NEO with SNR of only 0.35

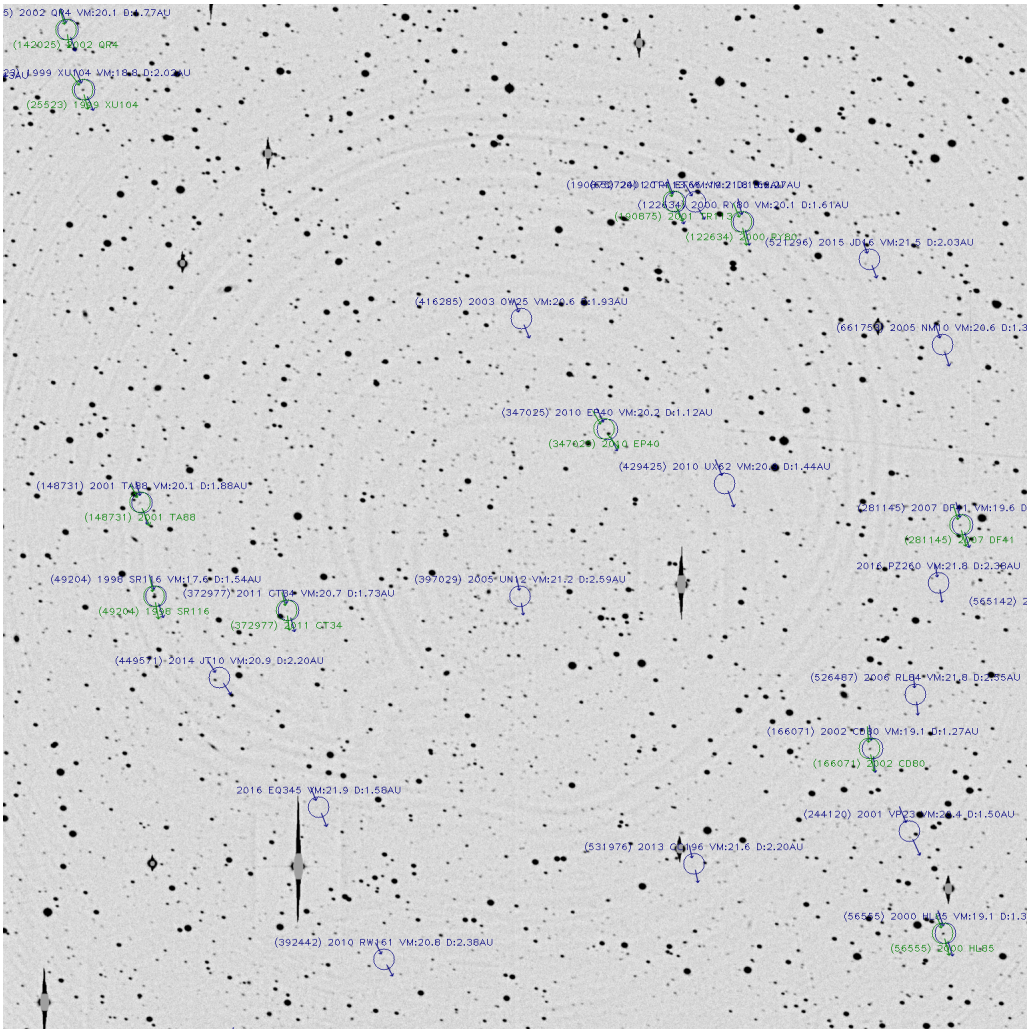

Figure 7: Annotation (blue) based on MPC database and detected NEOs (green) on long exposure image of Karl Schwarzschild Observatory, Tautenburg (50% crop)

The detection efficiency for faint objects (SNR < 0.5) is demonstrated in the accompanying image series (Figs. 8–9). These observational 2x2 binned frame crops, spanning exposure times of 26–30 seconds, showcase NEODetect's capability to resolve targets at magnitudes down to V = 19.5.

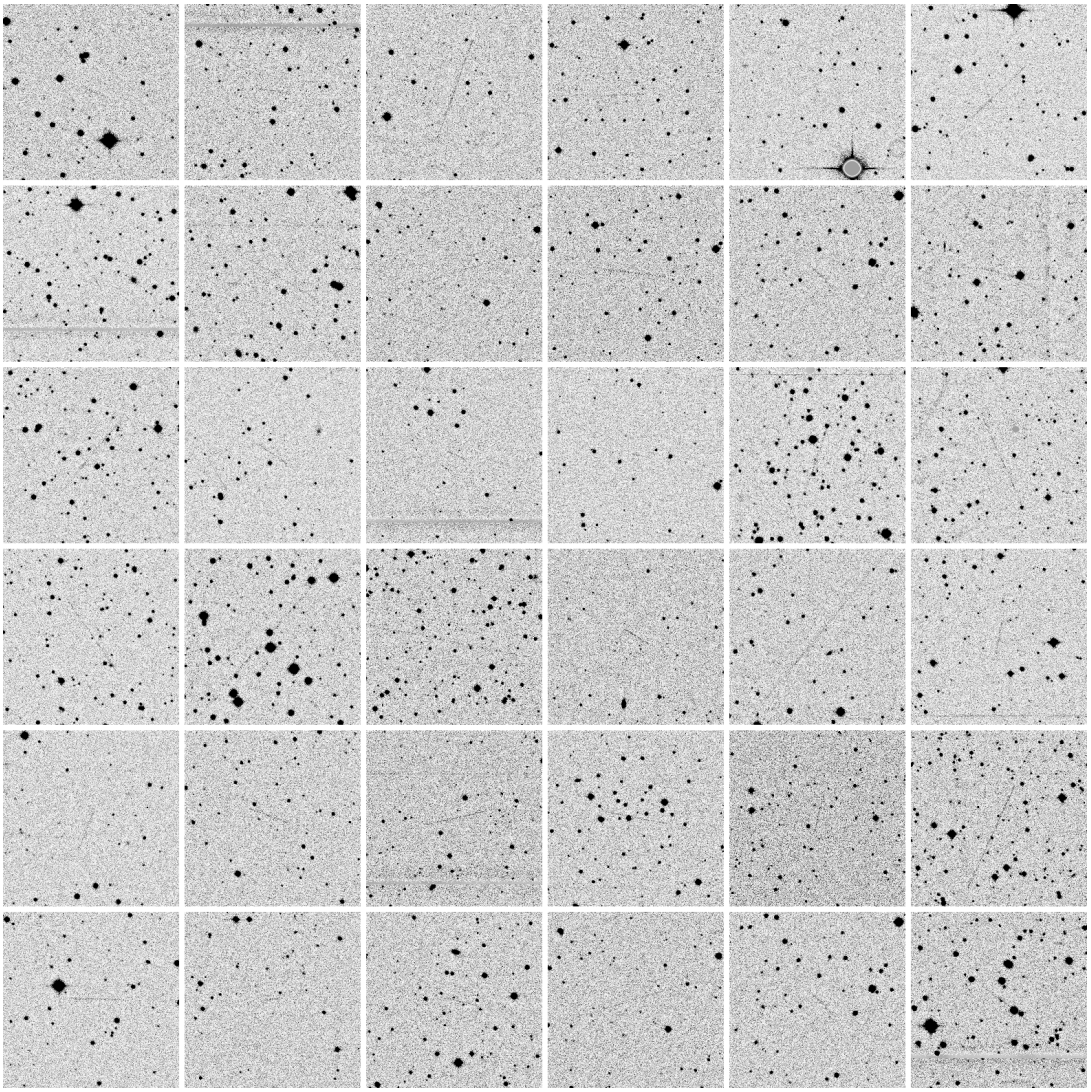

Figure 8: Very faint space debris candidates (SNR$<$0.5)

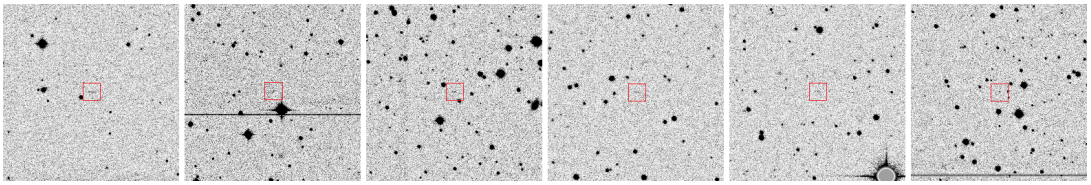

Figure 9: Very faint slow NEO candidates (SNR$<$1.0)

8. Comparison with other method

Synthetic tracking (ST) is a powerful method for detecting fast-moving NEOs. By computationally shifting and stacking a sequence of short-exposure images along possible motion paths, ST effectively eliminates the loss of signal caused by object trailing, thereby significantly boosting both detection sensitivity and astrometric precision. This technique is especially valuable for smaller telescopes, as it enables them to achieve deep limiting magnitudes by integrating over longer effective exposures, making them more capable in both survey and follow-up modes for NEOs [10].

The synthetic tracking was configured for a default apparent motion range of 2–25 ″/min, limited by hardware resources (GPU cluster with 4x high-end GPU). To address faster-moving targets, the system was augmented with a 4×4 pixel binning post-processing step, expanding its effective detection range to 25–40 ″/min.

An empirical analysis of NEO detections at the K88 site between 2023 and 2025 reveals distinct detection profiles for AI-driven NEODetect and synthetic tracking methodologies:

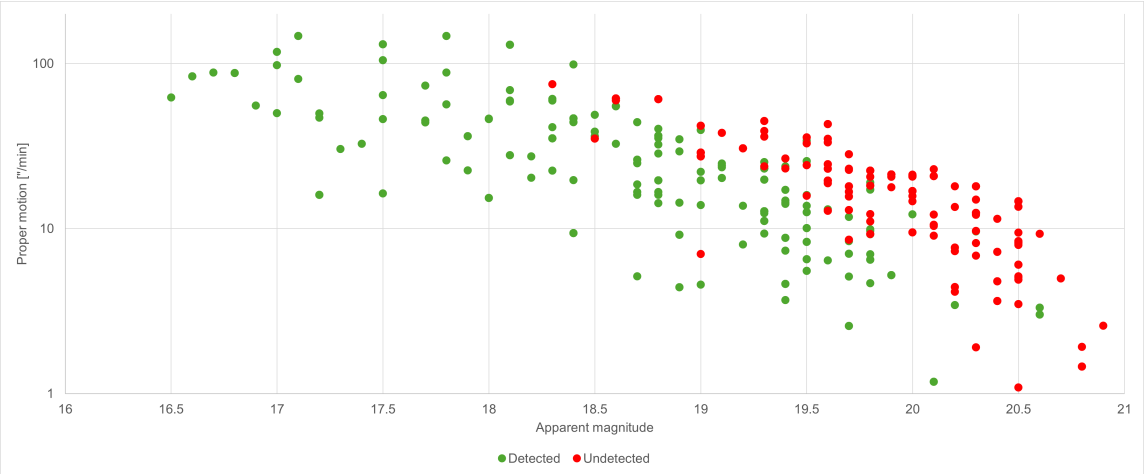

Figure 10: Green: NEODetect detections, Red: Synthetic tracking detections

Supplementing Synthetic Tracking: NEODetect does not discard the value of synthetic tracking. Instead, it can integrate synthetic tracking as a post-processing step: while NEODetect excels at real-time detection of fast-moving, linearly-traced faint objects – where rapid discovery is mission-critical – a synthetic tracking algorithm can operate in the background as a supplementary system. This hybrid approach leverages the advantages of both methodologies (real-time AI detection and post-hoc synthetic stacking) as a compromise solution, albeit at the cost of reduced sky coverage. The trade-off arises because synthetic tracking implementation necessitates a significantly higher observational repetition rate compared to NEODetect's spatio-temporally optimized processing with only 4-6 required images, thereby limiting the accessible field-of-view in practical operational frameworks.

Replacing Synthetic Tracking: NEODetect can fully replace synthetic tracking, albeit with a strategic compromise. While this eliminates the capability to detect the faintest slow-moving NEOs (which require extreme image stacking depths), it maximizes sky coverage – a critical advantage that compensates for reduced sensitivity. Our statistical analyses reveal that NEODetect identifies 53.8% of NEOs detectable via 18–24-frame synthetic tracking stacks while requiring only 4–6 frames. This efficiency enables 50–100% higher discovery rates under equivalent observational time.

For planetary defense applications, this trade-off is operationally justified. Maximizing the revisit rate of surveyed sky areas (rather than prolonged dwell times on individual fields) enhances the probability of detecting hazardous objects during their brief Earth-approach windows. Our test period confirms that NEODetect's wide-field prioritization strategy increases short-term mitigation of consequences of impact risk, compared to synthetic tracking-dominated surveys, despite marginally reduced sensitivity to ultra-faint objects. This paradigm shift aligns with the time-critical nature of planetary defense, where rapid all-sky coverage outweighs the diminishing returns of exhaustive depth-of-search optimizations.

Table 1: Comparison with synthetic/software tracking

Feature	NEODetect	Software tracking
Detection Method	Deep learning inference	Brute-force stacking
Hardware Requirements	CPU or GPU, Dockerized	Requires GPU (cluster)
Sensitivity	High overall	Extremely high for slow movers
False Positive Suppression	Advanced AI-based	Manual vetting needed
Deployment	On-premise or cloud, easy scaling	Typically on-premise
Speed	Realtime	Offline
Electric demand	Low	Very high

9. Conclusion and Future work

We have presented the development and validation of the NEODetect pipeline, an AI-driven system for real-time detection of Near-Earth Objects (NEOs) and space debris in astronomical images. By leveraging a spatiotemporal U-Net architecture trained on synthetic and observational data, the pipeline processes triple-consecutive 16-bit frames to distinguish faint streaks from noise while suppressing false positives. Key innovations include the spatiotemporal neural network architecture, temporal consequence filtering for false positive reduction (99.2% automated rejection), and integration of edge-compatible processing for deployment on both GPU and CPU platforms.

The system demonstrated robust performance across diverse observational conditions:

- Processing Speed: 4 seconds per 25-megapixel image on a Dell Precision 7960 server's 8 Xeon Gold CPUs via OpenVINO optimization.
- Detection Sensitivity: Identification of streaks with SNR$>$0.25 and lengths$>$4 pixels, validated against 246 confirmed NEOs in operational and historical testing.
- False Positives: Multi-stage filtering reduced candidates from 72.5 million to 579,661 and finally to 2833 true candidates and only 55 false positives, with final validation achieving 99.9999% accuracy.
- Comparative analysis revealed that NEODetect detects 53.8% of slow-moving NEOs identified by software tracking method on 18 frame sequences, but at 50× lower computational cost and using only 4 frames, enabling real-time analysis during image acquisition. This result also means that when using 4-6 image-long sequences, there is a possibility of finding 1.5-2x as many NEOs in a given period as with software tracking.

Future Directions will focus on:

- Enhancing photometric characterization of detected streaks.
- Integrating real-time orbit determination services for rapid trajectory validation.
- Expanding deployment to next-generation surveys.
- Implement the core of NEODetect onto a radiation-hardened edge device to demonstrate in-orbit operation capabilities.

Acknowledgements

This project has been supported by the European Space Agency and the Hungarian Government through the Space Safety Programme (S2P) Period 2 "Competitiveness Segment under COSMIC" with Contract No. 4000144612.

We thank the team at Piszkéstető Observatory for the 12 months observation and all of our beta testers on different sites around the world:

- K88 GINOP-KHK, Piszkéstetö

- 033 Karl Schwarzschild Observatory, Tautenburg
- 106 Crni Vrh Observatory
- 126 Monte Viseggi Observatory

Facilities:

- 60/90/60cm Schmidt Telescope at Piszkéstető
- 134/200cm Alfred Jensch Schmidt Telescope at Tautenburg
- 36cm Schmidt-Cassegrain telescope at Crni Vrh

Software:

- Machine Intelligence Designer (Machine Intelligence Zrt, 2020)
- NEODetect (Machine Intelligence Zrt, 2025)
- OpenCV (Bradski, 2000)
- IO.Astrodynamics (Sylvain Guillet, 2024)
- Watney Astrometry Engine (Jussi Saarivirta, 2023)
- Tensorflow (Martín Abadi et al., 2015)
- Keras (Chollet, Francois et al., 2015)
- Astropy (Astropy Collaboration et al. 2013, 2018)
- Skyfield (Rhodes, Brandon, 2019)
- Tycho Tracker (Daniel Parrott, 2025)

References

[1] Alvarez, L. W., Alvarez, W., Asaro, F., Michel, H. V., Extraterrestrial cause for the cretaceous-tertiary extinction, Science 208 (1980) 1095–1108.

[2] Chyba, C. F., Thomas, P. J., Zahnle, K. J., The 1908 tunguska explosion: atmospheric disruption of a stony asteroid, Nature 361 (1993) 40–44.

[3] Olga, P. P., Peter, J., et al., Chelyabinsk airburst, damage assessment, meteorite recovery, and characterization, Science 342 (2013) 1069–1073.

[4] Harris, A. W., D'Abramo, G., The population of near-earth asteroids, Icarus 257 (2015) 302–312.

[5] Duev, D. A., Mahabal, A., Ye, Q., et al., Deepstreaks: identifying fast-moving objects in the zwicky transient facility data with deep learning, Monthly Notices of the Royal Astronomical Society 486 (2019) 4158–4165.

[6] Kruk, S., Martín, P. G., Popescu, M., et al., Hubble asteroid hunter, Astronomy & Astrophysics 661 (2022) 15.

[7] M. Pöntinen, M. Granvik, A. A. Nucita, L. Conversi, et al., Euclid: Identification of asteroid streaks in simulated images using deep learning, Astronomy & Astrophysics 679 (2023) 18.

[8] Sárneczky, K., Szabó, N., Kiss, L. L., Velkei, Sz., NEO and imminent impactor discoveries from Hungary: recent results and lessons learnt, in: Acta Astronautica, Stellenbosch, Cape Town, South Africa. May 05–09, 2025. Paper number XYZ.

[9] Colin, L., Michael, D. F., Rene, V., Austin, R., Gregory, D. H., Temporal convolutional networks for action segmentation and detection, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA.

[10] Zhai, C., et al., Near-earth object observations using synthetic tracking, 2024.