



Tails: Chasing Comets with the Zwicky Transient Facility and Deep Learning

Dmitry A. Duev¹ , Bryce T. Bolin^{1,2} , Matthew J. Graham¹ , Michael S. P. Kelley³ , Ashish Mahabal^{1,4} , Eric C. Bellm⁵ ,
Michael W. Coughlin⁶ , Richard Dekany⁷ , George Helou² , Shrinivas R. Kulkarni¹ , Frank J. Masci² ,
Thomas A. Prince¹ , Reed Riddle⁷ , Maayane T. Soumagnac^{8,9} , and Stéfan J. van der Walt¹⁰

¹ Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA; duev@caltech.edu

² IPAC, California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91125, USA

³ Department of Astronomy, University of Maryland, College Park, MD 20742, USA

⁴ Center for Data Driven Discovery, California Institute of Technology, Pasadena, CA 91125, USA

⁵ DIRAC Institute, Department of Astronomy, University of Washington, 3910 15th Ave. NE, Seattle, WA 98195, USA

⁶ School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA

⁷ Caltech Optical Observatories, California Institute of Technology, Pasadena, CA 91125, USA

⁸ Lawrence Berkeley National Laboratory, 1 Cyclotron Rd., Berkeley, CA 94720, USA

⁹ Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot 76100, Israel

¹⁰ Berkeley Institute for Data Science, University of California, Berkeley, Berkeley, CA 94720, USA

Received 2021 January 11; revised 2021 February 23; accepted 2021 February 25; published 2021 April 8

Abstract

We present Tails, an open-source deep-learning framework for the identification and localization of comets in the image data of the Zwicky Transient Facility (ZTF), a robotic optical time-domain survey currently in operation at the Palomar Observatory in California, USA. Tails employs a custom EfficientDet-based architecture and is capable of finding comets in single images in near real time, rather than requiring multiple epochs as with traditional methods. The system achieves state-of-the-art performance with 99% recall, a 0.01% false-positive rate, and a 1–2 pixel rms error in the predicted position. We report the initial results of the Tails efficiency evaluation in a production setting on the data of the ZTF Twilight survey, including the first AI-assisted discovery of a comet (C/2020 T2) and the recovery of a comet (P/2016 J3 = P/2021 A3).

Unified Astronomy Thesaurus concepts: [Astroinformatics \(78\)](#); [Astronomy data analysis \(1858\)](#); [Convolutional neural networks \(1938\)](#); [Comets \(280\)](#)

1. Introduction

Comets have mesmerized humans for millennia, frequently offering, arguably, some of the most spectacular sights in the night sky. Containing the original materials from when the solar system first formed, comets provide a unique insight into the distant past of our solar system. The recent discovery of the first interstellar comet 2I/Borisov by amateur astronomer Gennadiy Borisov predictably sparked much excitement and enthusiasm among astronomers and the general public alike (e.g., Bolin et al. 2020; Fitzsimmons et al. 2019; Guzik et al. 2020). Such objects could potentially provide important information on the formation of other stellar systems. It is a very exciting time to look for comets: the large-scale time-domain surveys that are currently in operation, such as the Zwicky Transient Facility (ZTF; Bellm et al. 2019a; Graham et al. 2019), Panoramic Survey Telescope & Rapid Response System (Pan-STARRS; Chambers et al. 2016), or Asteroid Terrestrial-impact Last Alert System (ATLAS; Tonry et al. 2018), and the upcoming ones such as BlackGEM (Bloemen et al. 2016) and Vera Rubin Observatory/Large Synoptic Survey Telescope (LSST; Ivezić et al. 2008) offer the richest data sets ever available to mine for comets.

Traditional comet detection algorithms rely on multiple observations of cometary objects that are linked together and used to fit an orbital solution. To the best of our knowledge, the previous attempts to take the comet’s morphology in the optical image data into consideration in the detection algorithms have not led to reliable and robust results.

In this work, we present Tails—a state-of-the-art open-source deep-learning-based system for the identification and localization

of comets in the image data of ZTF.¹¹ Tails employs an EfficientDet-based architecture (Tan et al. 2019) and is thus capable of finding comets in single images in near real time, rather than requiring multiple epochs as with traditional methods.

1.1. The Zwicky Transient Facility

The ZTF¹² is a state-of-the-art robotic time-domain sky survey capable of visiting the entire visible sky north of -30° decl. every night. ZTF observes the sky in the g , r , and i bands at different cadences depending on the scientific program and sky region (Bellm et al. 2019a; Graham et al. 2019). The 576 megapixel camera with a 47° field of view, installed on the Samuel Oschin 48 inch (1.2 m) Schmidt Telescope, can scan more than 3750° per hour, to a 5σ detection limit of 20.7 mag in the r band with a 30 s exposure during a new moon (Masci et al. 2019; Dekany et al. 2020).

The ZTF partnership has been running a specialized survey, the Twilight Survey (ZTF-TS) that operates at solar elongations down to 35 degrees with an r -band limiting magnitude of 19.5 (Ye et al. 2020; Bellm et al. 2019b). ZTF-TS has so far resulted in the discovery of a number of Atira asteroids (orbits interior to Earth’s) as well as the first inner-Venus object, 2020 AV2 (Ip et al. 2020). Motivated by the success, ZTF-TS will be expanded in Phase II of the project, which commenced in 2020 December.

Comets become more easily detectable when close to the Sun as they become brighter and start exhibiting more

¹¹ <https://github.com/dmitryduev/tails>

¹² <https://ztf.caltech.edu>

pronounced coma and tails. Furthermore, it has been shown that the most detectable direction of approach of an interstellar object is from directly behind the Sun because of observational selection effects (Jedicke et al. 2016) and the fact that this direction has a greater cross section for asteroids to bend around and pass into the visibility volume (Engelhardt et al. 2017; Do et al. 2018).

Tails automates the search for comets with detectable morphology. While trained and evaluated on a large corpus of ZTF data, in this work we focus on Tails’ performance when applied to the ZTF-TS data.

2. Tails: A Deep-learning Framework for the Identification and Localization of Comets

Deep learning (DL) is a subset of machine learning that employs artificial multi-layer neural networks (McCulloch & Pitts 1943). DL systems are able to discover, in a highly automated manner, efficient representations of the data, simplifying the task of finding the meaningful sought-after patterns in them. We refer the reader to a brilliant introduction into DL given in Géron (2019).

DL systems often reach near-optimal performance for a given task and are able to learn even very complicated, highly nonlinear mappings between the input and output spaces. The art of building applied DL systems involves two major challenges: finding a suitable network architecture and, more importantly, constructing a large, labeled representative data set for the network training. In the case of comet detection, the training set must reflect the possible variations across different seeing conditions, filters, sky location, CCDs, and include data artifacts caused by, for example, cross-talk or telescope reflections.

2.1. Data Set

To build a seed sample for labeling, we first identified all potential observations of known comets conducted with ZTF from 2018 March 5 to 2020 March 4, based on their predicted position and brightness. The code for accomplishing that is based on the Python libraries `pypride` (Duev et al. 2016) and `solarsyslib` (Jensen-Clem et al. 2018) and uses the comet ephemerides obtained from the Minor Planet Center (MPC)¹³ for a coarse search, followed by a JPL Horizons¹⁴ (Giorgini et al. 1996) query for precision.

To provide more contextual information, epochal image data are supplemented by properly aligned reference images of the corresponding patches of sky and difference (epochal minus reference) images generated with the ZOGY algorithm (Zackay et al. 2016), all produced by the ZTF Science Data System at Caltech’s IPAC (Masci et al. 2019). Finally, we generate image triplet cutouts 256 by 256 pixel in size, which in angular measure translates into 4.3 by 4.3 at ZTF’s pixel scale of 1.01/pixel.

We selected over 60,000 individual observations with the total comet magnitude ranging from 10 to 23 (as reported by JPL Horizons; see Figure 1), out of which about 20,000 were sourced for manual annotation. This resulted in an initial sample of 3000 examples with identifiable morphology.




Figure 1. Distribution of over 60,000 individual observations of comets as a function of the predicted total magnitude (as reported by JPL Horizons) used in the seed sample.

We also compiled a set of approximately 20,000 negative examples consisting of point-like cometary detections, patches of sky with no identified transient or variable sources, CCD-edge cases, and a wide range of real (point-source) transient and bogus (e.g., artifacts due to bright stars, optical ghosts, and “demonstrators”) samples from the Braai data set (Duev et al. 2019).

To expand the data set, we then assembled a standard ResNet-based (He et al. 2015) classifier for comet identification. With this basic classification model, we ran several rounds of an active-learning-like procedure, where we would first train the classifier, evaluate it on the whole data set, sample both confident predictions and the cases close to the classifier’s decision boundary, manually inspect and label those examples, and add them to the training set. Roughly 2000 positive and 2000 negative examples were added to the training set via this method.

The resulting training data set contains about 5000 positive and 22,000 negative examples (see Figure 2). Each triplet in the set has been assigned a label $[p_c, x, y]$, where p_c marks the presence of a comet in the image and $x, y \in [0, 1]$ is the relative positions of the comet’s center of mass, as reported by JPL Horizons. For positive examples this translates into $[1, x_{\text{JPL}}, y_{\text{JPL}}]$, and for negative ones $[0, ?, ?]$, where the question marks denote that these do not affect the loss in this case.

2.2. Deep Neural Network Architecture and Training

Tails adopts a custom architecture (see Figure 3) based on EfficientDet D0 (Tan et al. 2019), a variant of a state-of-the-art architecture designed for object detection—a computer-vision technique for the identification and location of objects in image data.

This architecture delivers best-in-class object detection efficiency and performance across a wide range of resource constraints. This is achieved by using EfficientNet—state-of-the-art backbone networks for feature extraction, a weighted bidirectional feature pyramid network (BiFPN), which allows easy and fast multiscale feature fusion, and a compound scaling method that simultaneously and uniformly scales the resolution, depth, and width for all backbone, feature, and location/class prediction networks (Tan et al. 2019).

The use of a BiFPN, which effectively represents and processes multiscale features, makes this architecture particularly well suited for the problem of morphology-based comet identification and localization.

A batch of triplet image stacks ($n_b, 256, 256, 3$) in size, where n_b is the number of stacks in the batch, is passed through an EfficientNet B0 backbone (Tan & Le 2019). The extracted

¹³ <https://minorplanetcenter.net/iau/MPCORB/CometEls.txt>

¹⁴ <https://ssd.jpl.nasa.gov/horizons.cgi>

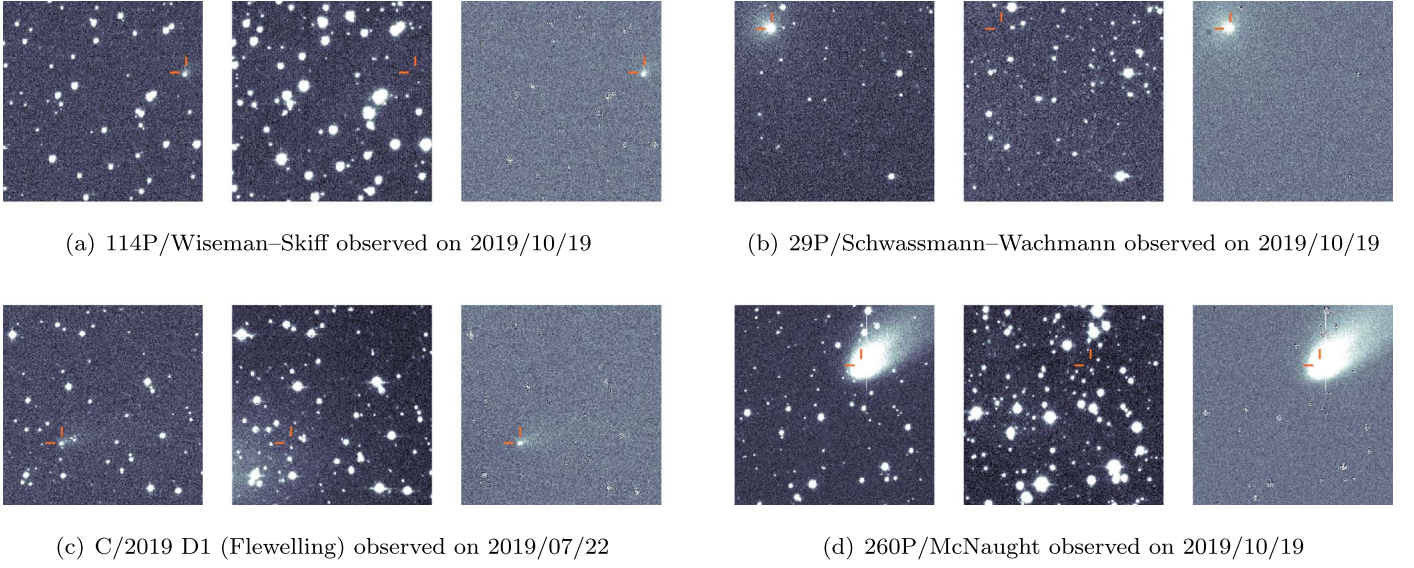


Figure 2. Comet examples from the training set. The first image in a triplet is the epochal science image, the second is the reference image, and the third is the ZOGY difference image.




Figure 3. Tails architecture: a custom EfficientDet D0-based network (Tan et al. 2019). A batch of duplet or triplet image stacks of size $(n_b, 256, 256, [2|3])$, correspondingly, is passed through an EfficientNet B0 backbone, where n_b is the number of stacks in the batch. The extracted features from the last five blocks/levels of the backbone network are passed through a bidirectional feature pyramid network (BiFPN). The resulting five output tensors denoted in colored circles are fed into the head network, which outputs the probability of the image containing a comet p_c and its predicted relative position (x, y) .

features from the last five blocks/levels of the network are passed through the BiFPN. The resulting five output tensors denoted by the colored circles in Figure 3 are fed into the head network, which outputs the probability of the image containing a comet p_c and its centroid's predicted relative (x, y) position.¹⁵

We defined the loss function as


$$L = w_c \cdot L_c + w_p \cdot L_p, \quad (1)$$

¹⁵ We note that standard object detection algorithms typically output bounding boxes and corresponding object class probabilities, i.e., sets of $(4+n_{\text{classes}})$ numbers. Our approach allowed us to simplify the head network architecture and both simplify and speed up the assembly of the training data set, bypassing the unnecessary complexity and potential inaccuracy of drawing boundary boxes around known comet detections.


where L_c denotes the binary cross-entropy function for the label c (1—there is a comet in the image, 0—there is no comet) and the predicted probability p_c . If $[p_c] = 1$, L_p is computed as an L_1 loss for the relative position (x, y) and its prediction (x_p, y_p) with a small L_2 regularizing term (with $\epsilon = 10^{-3}$), and w_c and w_p denote the weights of the two terms, respectively,

$$\begin{aligned} L_c &= \sum c \cdot \log(p_c) + (1 - c) \cdot \log(1 - p_c) \\ L_p &= \sum [p_c] \cdot (|x - x_p| + |y - y_p| \\ &\quad + \epsilon \cdot \sqrt{|x - x_p|^2 + |y - y_p|^2}). \end{aligned} \quad (2)$$

We employed the Adam optimizer (Kingma & Ba 2014), a batch size of 32, and a 81%/9%/10% training/validation/test data split. For data augmentation, we applied random horizontal

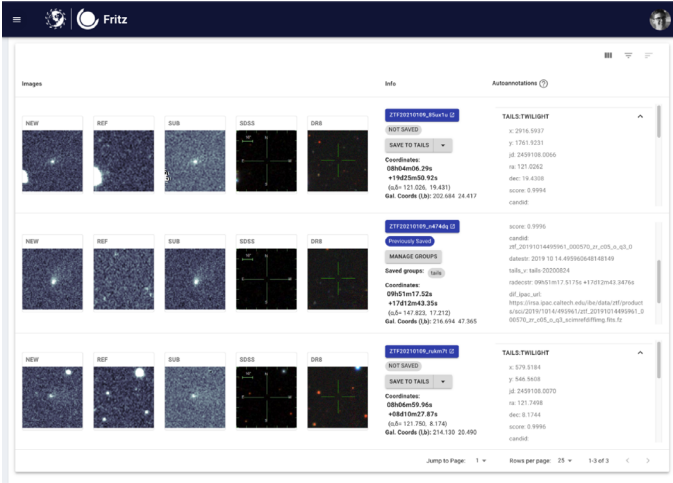


(a) False positive rate (FPR) and false negative rate (FNR) as a function of score p_c . FPR and FNR balance out at around 1.7% for a score threshold of 0.5.

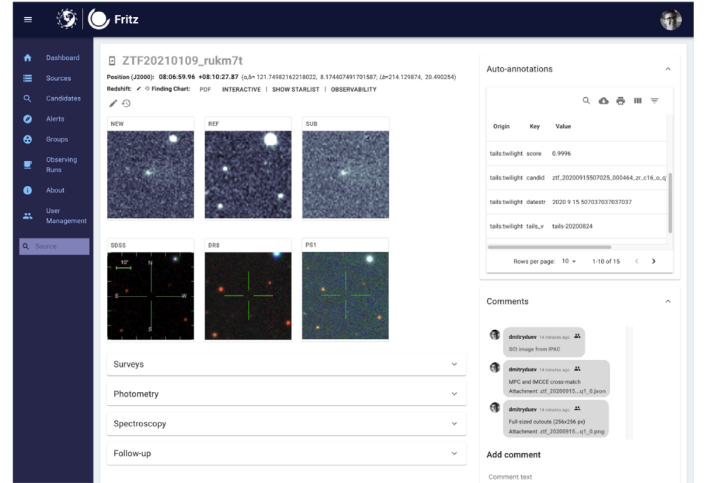


(b) RMSE of the predicted comet position versus the reported on JPL Horizon for the 650 positive samples from the test set.

Figure 4. Test set performance of Tails. The set contains 1400 negative and 650 positive examples.



(a) Candidate scanning page. The users can inspect the candidates and save vetted objects to one or more groups. Candidates that are not saved to any group within 7 days are removed from Fritz.



(b) Source page. It aggregates and displays in an interactive manner all kinds of information related to an object that exists on Fritz, such as photometry, spectroscopy, auto-annotations, comments, finder charts, follow-up requests, and other data.

Figure 5. Screenshots of the Fritz user interfaces used for Tails candidate inspection and vetting.

and vertical flips of the input data; no random rotations and translations were added. We note that the test/validation sets did not contain augmented data from the training set. We used standard techniques to maximize training performance: if no improvement in validation loss was observed for 10 epochs, the learning rate was reduced by a factor of 2, and training was stopped early if no improvement was observed for 30 epochs.

The EfficientNet's weights were randomly initialized.¹⁶ We first set $w_c = 10$, $w_p = 1$ to allow for a fast convergence of the feature-extracting part of the network. To fine-tune the performance, we trained Tails on a balanced data set setting $w_c = 1.1$, $w_p = 1$ and monitored the validation loss for early stopping, then bumped $w_p = 2$ and monitored the validation positional rms error, and finally, we added the omitted negative examples and again monitored the validation loss for early stopping.

The resulting classifiers were put through the same active-learning-like procedure as was employed in the initial data set assembly, using several months of ZTF-TS data.

¹⁶ We experimented with pretrained weights, however that neither helped the network to reach convergence faster, nor did it affect the final performance. We believe this is likely due to the fact that astronomical images are very different from those in commonly used data sets.

3. Tails Performance

Evaluated on the test set, with a score p_c threshold of 0.5, Tails demonstrates false-positive (FPR) and false-negative rates (FNR) of 1.7%, and a ~ 1 –2 pixel median rms error of the predicted comet centroid position versus that acquired from JPL Horizon (see Figure 4).


The ZTF instrument's CCD mosaic has 16 individual $6k \times 6k$ science CCDs. The raw ZTF image data are split into four readout quadrants per CCD and all processing is conducted independently on each CCD readout quadrant. We tessellate each $3k \times 3k$ CCD-quadrant image into a 13×13 grid of overlapping 256×256 pixel tiles and evaluate Tails on those.¹⁷

Tails has been deployed in production since late 2020 June. We have implemented a sentinel service¹⁸ that processes the incoming data in real time and posts the plausible candidates to Fritz¹⁹, the ZTF Phase II open-source science data platform (van der Walt et al. 2019; Duev et al. 2019;


¹⁷ Standard fully convolutional approaches often used in computer vision proved to be overkill in this case.

¹⁸ See <https://github.com/dmitryduev/tails>.


¹⁹ <https://github.com/fritz-marshall/fritz>




(a) 2I/Borisov observed on 2019/10/05




(d) Bogus detection next to a bright star




(b) 2I/Borisov observed on 2019/10/15



(e) Bogus detection due to a brightened satellite trail



(c) C/2010 U3 (Boattini) observed on 2020/08/11



(f) Bogus detection due to a telescope reflection

Figure 6. Candidates identified by Tails. Panels (a), (b), and (c) on the left show the detections of real comets. Typical false positives are shown on panels (d), (e), and (f) on the right. For each image triplet, the left pane shows the epochal science exposure, the middle pane shows the reference image of the corresponding patch of sky, and the right pane shows the ZOGY difference image.

Table 1
Orbital Elements of C/2020 T2 Provided by the MPC

Element	Value
e	0.9934213
Incl.	27.87307
Peri.	150.38279
Node	83.04834
q	2.0546940
T	2021 July 11.14638 TT

Kasliwal et al. 2019), for further manual inspection and vetting. The candidates are auto-annotated with the detailed information on the detection such as the score, CCD, and sky positions, and cross-matches with known solar system objects. Figure 5 shows screenshots of the Fritz user interfaces used in the process.

It takes about 5 hr to run inference on a typical set of nightly ZTF-TS data (~ 45 30 s exposures) on an *e2-highcpu-32* virtual machine instance (32 vCPU, 32 GB memory, SSD disk) on the Google Cloud Platform, including I/O operations.

Consistently with the expected rate of comet observations, a typical run on nightly Twilight data yields a few dozen candidates, which, given the typical number of processed tiles, gives an empirical FPR value of about 0.01%.

The scanning results are accumulated and used to expand the training set and improve Tails' performance.

We have evaluated Tails' performance on a random sample of 200 observations of known comets with identifiable morphology in 2020 July–August and found an empirical recall value of 99%.

Figure 6 shows a number of comet candidates not from the training set identified by Tails, including some of the ZTF observations of the comet 2I/Borisov. Optical artifacts resembling cometary objects are the main source of contamination.

3.1. Discovery of Comet C/2020 T2

On 2020 October 7, Tails discovered a candidate that was posted to MPC's Possible Comet Confirmation Page (PCCP)²⁰ as ZTFDD01 (see Figure 7). It was later confirmed to be a long-period comet and designated C/2020 T2 (Palomar), marking the first DL-assisted comet discovery (Duev et al. 2020). The candidate was found in the Twilight survey data; it was at 19.3 mag in the ZTF r band. The FWHM of the object was approximately $2''.5$ – $3''$, compared to nearby background stars that have an FWHM of $\sim 2''$. The object showed a tail extending up to $5''$ in the westward direction. Table 1 summarizes the orbital elements of C/2020 T2 provided by the MPC and Figure 8 shows its orbit as of the discovery date.

To determine if Tails could have discovered C/2020 T2 before 2020 October 7, we searched the ZTF archive for all Twilight Survey data covering the ephemeris position of the

²⁰ https://minorplanetcenter.net/iau/NEO/pccp_tabular.html




Figure 7. Discovery image of C/2020 T2 (Palomar), the first DL-assisted comet discovery by Tails. Taken on 2020 October 7 with the ZTF camera on the 48 inch Schmidt telescope at Palomar. The left panel shows the epochal science exposure (256×256 pixel cutout), the middle panel shows the reference image, and the right panel shows the ZOGY difference image. East is to the left, north is down.

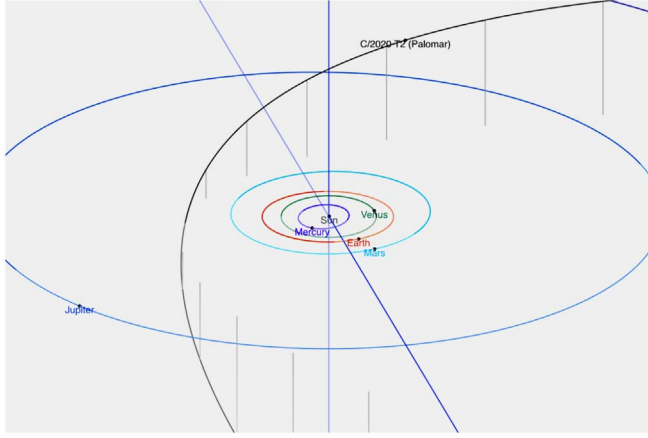


Figure 8. The orbit of comet C/2020 T2 as of 2020 October 7. Image credit: NASA/JPL-Caltech / D. Duev.

comet with the ZChecker software (Kelley et al. 2019). Eleven nights of data were found between 2020 June 11–20 (evening twilight) and October 7–21 (morning twilight). The comet was in conjunction with the Sun between the two sets, and not observable by ZTF. We measured the brightness of the coma in 4 pixel radius apertures, and aperture corrected the photometry according to the ZTF pipeline documentation. The data are shown in Figure 9. Typical seeing was $2''$ in June, and the comet was very faint ($r = 20.2$ mag), near the single-image detection limit ($r = 20.4$ – 20.9 mag, 5σ point source), and had no morphological features for Tails to pick up. Thus 2020 October 7 was really the first opportunity for Tails to discover the comet.

3.2. Recovery of Comet P/2016 J3 = P/2021 A3 (STEREO)

A comet candidate was identified by a combination of Tails and the ZTF Moving Object Detection Engine (Masci et al. 2019) on 2020 January 4 UTC and was submitted to the PCCP as ZTF0Ion (see Figure 10). It was later identified as a recovery of comet P/2016 J3 (STEREO) and given the designation P/2016 J3 = P/2021 A3 (STEREO; Bolin 2021). P/2021 A3 was identified in the evening Twilight survey data at r 19.3 mag with a clearly extended appearance scoring 0.9 with a coma

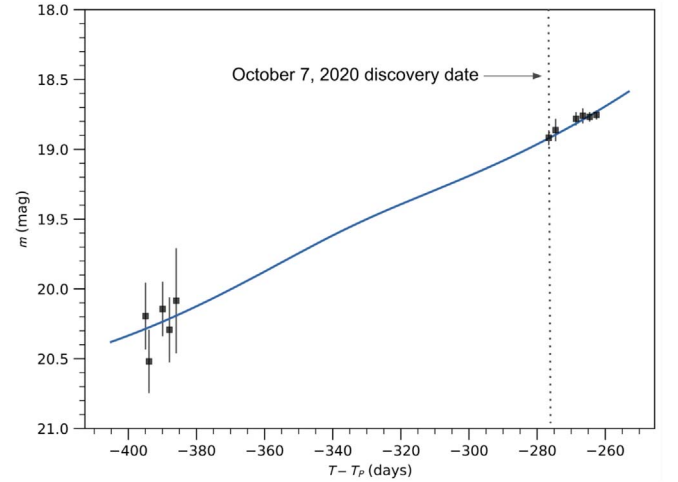


Figure 9. Photometry of comet C/2020 T2 (Palomar) derived from ZTF-TS images (r band) vs. time from perihelion. A best-fit model lightcurve is also shown: $r = 9.85 + 9.54 \log_{10}(r_h) + 5 \log_{10}(\Delta) - \Phi(\theta)$, where r_h is the heliocentric distance in astronomical units, Δ is the comet-observer distance in astronomical units, and $\Phi(\theta)$ is the phase angle correction from Schleicher et al. (1998). T_p denotes the time of perihelion passage (2021 July 11).

$\sim 10''$ wide and a tail extending past $20''$ in the northeast direction.

4. Discussion

This work demonstrates the potential of the state-of-the-art deep-learning computer-vision architecture designs when applied to the problem of astronomical source detection and localization, with a specific focus on comets.

We experimented with the input data and trained a version of Tails that instead of triplet image stacks uses duplets—epochal/reference images, omitting the ZOGY difference images. Our tests show that this version achieves essentially the same performance as the one trained on triplets without requiring image differencing, expanding the range of potential use cases of Tails.

While Tails is trained only on ZTF data, with transfer learning it can be adapted to other sky surveys, including the upcoming Vera Rubin Observatory’s LSST (Ivezić et al. 2008).




Figure 10. Recovery detection of P/2021 A3 (STEREO) identified by Tails in the ZTF *r*-band data taken on 2021 January 4. The left panel shows the epochal science exposure (256 × 256 pixel cutout), the middle panel shows the reference image, and the right panel shows the ZOGY difference image. East is to the left, north is down.

D.A. Duev would like to thank Ivan Duev for assistance with data labeling. D.A. Duev acknowledges support from Google Cloud and from the Heising-Simons Foundation under grant No. 12540303.

Based on observations obtained with the Samuel Oschin Telescope 48 inch and the 60 inch Telescope at the Palomar Observatory as part of the Zwicky Transient Facility project. ZTF is supported by the National Science Foundation under grant No. AST-1440341 and a collaboration including Caltech, IPAC, the Weizmann Institute for Science, the Oskar Klein Center at Stockholm University, the University of Maryland, the University of Washington, Deutsches Elektronen-Synchrotron and Humboldt University, Los Alamos National Laboratories, the TANGO Consortium of Taiwan, the University of Wisconsin at Milwaukee, and Lawrence Berkeley National Laboratories. Operations are conducted by COO, IPAC, and UW.

This research has made use of data and/or services provided by the International Astronomical Union’s Minor Planet Center.

The authors would like to express gratitude to the anonymous referee.

Facilities: PO:1.2 m, ZTF.

Software: astropy (Astropy Collaboration et al. 2018), Fritz (<https://github.com/fritz-marshall/fritz>), Kowalski (Duev et al. 2019), matplotlib (Hunter 2007), numpy (Harris et al. 2020), pandas (Pandas Development Team 2020), pypride (Duev et al. 2016), SEP (Barbary 2016), sbpy (Mommert et al. 2019), TensorFlow (Abadi et al. 2016), ZChecker (Kelley et al. 2019).

ORCID iDs

Dmitry A. Duev <https://orcid.org/0000-0001-5060-8733>
 Bryce T. Bolin <https://orcid.org/0000-0002-4950-6323>
 Matthew J. Graham <https://orcid.org/0000-0002-3168-0139>
 Michael S. P. Kelley <https://orcid.org/0000-0002-6702-7676>
 Ashish Mahabal <https://orcid.org/0000-0003-2242-0244>
 Eric C. Bellm <https://orcid.org/0000-0001-8018-5348>
 Michael W. Coughlin <https://orcid.org/0000-0002-8262-2924>
 Richard Dekany <https://orcid.org/0000-0002-5884-7867>
 George Helou <https://orcid.org/0000-0003-3367-3415>
 Shrinivas R. Kulkarni <https://orcid.org/0000-0001-5390-8563>

Frank J. Masci <https://orcid.org/0000-0002-8532-9395>
 Thomas A. Prince <https://orcid.org/0000-0002-8850-3627>
 Reed Riddle <https://orcid.org/0000-0002-0387-370X>
 Maayane T. Soumagnac <https://orcid.org/0000-0001-6753-1488>
 Stéfan J. van der Walt <https://orcid.org/0000-0001-9276-1891>

References

- Abadi, M., Agarwal, A., Barham, P., et al. 2016, arXiv:1603.04467
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
 Barbary, K. 2016, *JOSS*, **1**, 58
 Bellm, E. C., Kulkarni, S. R., Barlow, T., et al. 2019b, *PASP*, **131**, 068003
 Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019a, *PASP*, **131**, 018002
 Bloemen, S., Groot, P., Woudt, P., et al. 2016, *Proc. SPIE*, 9906, 990664
 Bolin, B. T., Lisse, C. M., Kasliwal, M. M., et al. 2020, *AJ*, **160**, 26
 Bolin, B. T. E. A., Lin, Y.-Z., Masci, F. J., et al. 2021, MPEC 2021-A157: COMET P/2016 J3 = P/2021 A3 (STEREO), <https://www.minorplanetcenter.net/mpec/K21/K21AF7.html>
 Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv:1612.05560
 Dekany, R., Smith, R. M., Riddle, R., et al. 2020, *PASP*, **132**, 038001
 Do, A., Tucker, M. A., & Tonry, J. 2018, *ApJL*, **855**, L10
 Duev, D. A., Mahabal, A., Masci, F. J., et al. 2019, *MNRAS*, **489**, 3582
 Duev, D. A., Pogrebenko, S. V., Cimò, G., et al. 2016, *A&A*, **593**, A34
 Duev, D. A. E. A., Duev, I. D., Bolin, B. T., et al. 2020, MPEC 2020-U170: COMET C/2020 T2 (Palomar), <https://minorplanetcenter.net/mpec/K20/K20UH0.html>
 Engelhardt, T., Jedicke, R., Vereš, P., et al. 2017, *AJ*, **153**, 133
 Fitzsimmons, A., Hainaut, O., Meech, K. J., et al. 2019, *ApJL*, **885**, L9
 Géron, A. 2019, Hands-on machine learning with Scikit-Learn and TensorFlow (2nd ed.; Sebastopol, CA: O’Reilly Media, Inc.)
 Giorgini, J. D., Yeomans, D. K., Chamberlin, A. B., et al. 1996, AAS/DPS Meeting, **28**, 25.04
 Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, *PASP*, **131**, 078001
 Guzik, P., Drahus, M., Rusek, K., et al. 2020, *NatAs*, **4**, 53
 Harris, C. R., Jarrod Millman, K., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
 He, K., Zhang, X., Ren, S., & Sun, J. 2015, arXiv:1512.03385
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Ip, W. H., Bolin, B. T., Masci, F. J., et al. 2020, arXiv:2009.04125
 Ivezić, Z., Tyson, J. A., Acosta, E., et al. 2008, arXiv:0805.2366v4
 Jedicke, R., Bolin, B., Granvik, M., & Beshore, E. 2016, *Icar*, **266**, 173
 Jensen-Clem, R., Duev, D. A., Riddle, R., et al. 2018, *AJ*, **155**, 32
 Kasliwal, M. M., Cannella, C., Bagdasaryan, A., et al. 2019, *PASP*, **131**, 038003
 Kelley, M. S. P., Bodewits, D., Ye, Q., et al. 2019, in ASP Conf. Ser., 523, Astronomical Data Analysis Software and Systems XXVII, ed. P. J. Teuben et al. (San Francisco, CA: ASP), 471
 Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980
 Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, *PASP*, **131**, 018003

- McCulloch, W. S., & Pitts, W. 1943, [The Bulletin of Mathematical Biophysics](#), 5, 115
- Mommert, M., Kelley, M., de Val-Borro, M., et al. 2019, [JOSS](#), 4, 1426
- Pandas Development Team 2020, T. pandas-dev/pandas: Pandas, latest, Zenodo doi:[10.5281/zenodo.3509134](#)
- Schleicher, D. G., Millis, R. L., & Birch, P. V. 1998, [Icar](#), 132, 397
- Tan, M., & Le, Q. V. 2019, arXiv:[1905.11946](#)
- Tan, M., Pang, R., & Le, Q. V. 2019, arXiv:[1911.09070](#)
- Tonry, J. L., Denneau, L., Heinze, A. N., et al. 2018, [PASP](#), 130, 064505
- van der Walt, S. J., Crellin-Quick, A., & Bloom, J. S. 2019, [JOSS](#), 4, 1247
- Ye, Q., Masci, F. J., Ip, W.-H., et al. 2020, [AJ](#), 159, 70
- Zackay, B., Ofek, E. O., & Gal-Yam, A. 2016, [ApJ](#), 830, 27