



Universiteit
Leiden
The Netherlands

Identifying tidal disruption events with an expansion of the FLEET machine-learning algorithm

Gomez, S.; Villar, V.A.; Berger, E.; Gezari, S.; Velzen, S. van; Nicholl, M.; ... ; Alexander, K.D.

Citation

Gomez, S., Villar, V. A., Berger, E., Gezari, S., Velzen, S. van, Nicholl, M., ... Alexander, K. D. (2023). Identifying tidal disruption events with an expansion of the FLEET machine-learning algorithm. *The Astrophysical Journal*, 949(2). doi:10.3847/1538-4357/acc535

Version: Publisher's Version
License: [Creative Commons CC BY 4.0 license](#)
Downloaded from: <https://hdl.handle.net/1887/3715869>

Note: To cite this publication please use the final published version (if applicable).



Identifying Tidal Disruption Events with an Expansion of the FLEET Machine-learning Algorithm

Sebastian Gomez¹ , V. Ashley Villar^{2,3,4} , Edo Berger^{5,6} , Suvi Gezari¹ , Sjoert van Velzen⁷ , Matt Nicholl⁸ , Peter K. Blanchard⁹ , and Kate. D. Alexander^{9,10}

¹ Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA; sgomez@stsci.edu

² Department of Astronomy & Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA

³ Institute for Computational & Data Sciences, The Pennsylvania State University, University Park, PA 16802, USA

⁴ Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

⁵ Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138-1516, USA

⁶ The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, USA

⁷ Leiden Observatory, Leiden University, P.O. Box 9513, 2300 RA Leiden, The Netherlands

⁸ Birmingham Institute for Gravitational Wave Astronomy and School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, UK

⁹ Center for Interdisciplinary Exploration and Research in Astrophysics and Department of Physics and Astronomy, Northwestern University, 1800 Sherman Avenue, 8th Floor, Evanston, IL 60201, USA

¹⁰ Department of Astronomy/Steward Observatory, University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721-0065, USA

Received 2022 October 19; revised 2023 March 13; accepted 2023 March 14; published 2023 June 2

Abstract

We present an expansion of FLEET, a machine-learning algorithm optimized to select transients that are most likely tidal disruption events (TDEs). FLEET is based on a random forest algorithm trained on both the light curves and host galaxy information of 4779 spectroscopically classified transients. We find that for transients with a probability of being a TDE, $P(\text{TDE}) > 0.5$, we can successfully recover TDEs with $\approx 40\%$ completeness and $\approx 30\%$ purity when using their first 20 days of photometry or a similar completeness and $\approx 50\%$ purity when including 40 days of photometry, an improvement of almost 2 orders of magnitude compared to random selection. Alternatively, we can recover TDEs with a maximum purity of $\approx 80\%$ and a completeness of $\approx 30\%$ when considering only transients with $P(\text{TDE}) > 0.8$. We explore the use of FLEET for future time-domain surveys such as the Legacy Survey of Space and Time on the Vera C. Rubin Observatory (Rubin) and the Nancy Grace Roman Space Telescope (Roman). We estimate that $\sim 10^4$ well-observed TDEs could be discovered every year by Rubin and ~ 200 TDEs by Roman. Finally, we run FLEET on the TDEs from our Rubin survey simulation and find that we can recover $\sim 30\%$ of them at redshift $z < 0.5$ with $P(\text{TDE}) > 0.5$, or ~ 3000 TDEs yr^{-1} that FLEET could uncover from the Rubin stream. We have demonstrated that we will be able to run FLEET on Rubin photometry as soon as this survey begins. FLEET is provided as an open source package on GitHub: <https://github.com/gmzsebastian/FLEET>.

Unified Astronomy Thesaurus concepts: Black hole physics (159); Supernovae (1668); Surveys (1671)

1. Introduction

Tidal disruption events (TDEs) occur when a star gets too close to a supermassive black hole and is subsequently torn apart by the tidal forces of the black hole (Hills 1975; Rees 1988). Following this disruption, about half the stellar debris is expected to return toward the black hole and circularize into an accretion disk, beginning a phase of accretion when a bright optical transient can be observed (Gezari et al. 2009; Guillochon et al. 2009). To date, about 70 TDEs have been discovered across the electromagnetic spectrum, with a wide variety of observational features (Auchettl et al. 2017; Mockler et al. 2019; van Velzen et al. 2020; Gezari 2021; van Velzen et al. 2021; Hammerstein et al. 2023; Nicholl et al. 2022). Some optically discovered TDEs exhibit hydrogen and helium emission, others show only helium (Gezari et al. 2012; Arcavi et al. 2014), and some additionally have nitrogen and oxygen lines (Blagorodnova et al. 2019; Leloudas et al. 2019). Van Velzen et al. (2021) defined three classes of TDEs: TDE-H (hydrogen only), TDE-

He (helium only), and TDE-H+He (Bowen lines in combination with H and/or He). There is at least one TDE that evolved from TDE-H+He to TDE-He (Nicholl et al. 2019). While some TDEs show X-ray emission in excess of their optical luminosity, some are X-ray dim (Holoien et al. 2016; Auchettl et al. 2017), and others alternate between these two states (Gezari et al. 2017). Additionally, radio observations suggest that some TDEs produce fast outflows, including relativistic jets, while others do not (Zauderer et al. 2011; Bower et al. 2013; van Velzen et al. 2013; Alexander et al. 2016; van Velzen et al. 2016; Alexander et al. 2020; Cendes et al. 2022).

Finding more TDEs rapidly and efficiently will lead to a better understanding of their nature, origin, and evolution. The challenge is that these are rare transients, representing only $\sim 0.5\%$ of all spectroscopically classified transients from magnitude-limited surveys (e.g., Fremling et al. 2020). Nevertheless, the field of time-domain astronomy is growing at an increasing rate, with surveys such as the Zwicky Transient Facility (ZTF) discovering thousands of transients a month (Bellm et al. 2019). With existing resources, it is only possible to spectroscopically classify $\sim 10\%$ of all optical transients discovered. This will become even more challenging when the Legacy Survey of Space and Time on the Vera C. Rubin Observatory (Rubin) commences in 2024 and increases the

transient discovery rate by ~ 2 orders of magnitude (Ivezic et al. 2019).

Machine-learning (ML) algorithms can be used to select promising TDE candidates for spectroscopic follow-up. Some general-purpose ML classifiers that attempt to predict the classes of optical transients already exist, but none are trained to classify TDEs using real observational data. For example, SuperRAENN (Villar et al. 2020) and Superphot (Villar et al. 2019; Hosseinzadeh et al. 2020) have been trained on real data from the Pan-STARRS1 Medium Deep Survey (PS1 MDS) and use either a recurrent autoencoder neural network or a random forest (RF) algorithm, respectively, to predict the classes of five types of supernovae (SNe Ia, Ibc, II, IIn, and SLSN-I). However, TDEs are not one of the supported classes in these classifiers, since the PS1 MDS training set only included two spectroscopically classified TDEs. The RAPID (Muthukrishna et al. 2019) and Avocado (Boone 2019) classifiers, on the other hand, are trained to distinguish among at least 12 different transient classes, including TDEs, but were trained on simulated data from the Photometric LSST Astronomical Time-series Classification project (Kessler et al. 2019), and their effectiveness on real data is yet to be verified. Lastly, the Automatic Learning for the Rapid Classification of Events (ALeRCE) broker (Sánchez-Sáez et al. 2021) uses a two-stage RF algorithm that first classifies the general nature of the transient as either an active galactic nucleus (AGN), SN, variable star, asteroid, or bogus and then proceeds to refine the class into 15 different types of transients, but TDEs are not one of these classes.

Given the absence of robust TDE photometric classifiers, we present a new version of Finding Luminous and Exotic Extragalactic Transients (FLEET), an ML algorithm originally developed to target follow-up of likely Type I superluminous SNe (SLSNe) based on their predicted probability of being SLSNe (Gomez et al. 2020a). Over the past 2 yr of operations, we have proven the efficacy of FLEET in finding SLSNe and surpassed our expectations of performance. We managed to achieve a peak purity of $\approx 80\%$ and discovered 21 of the 50 SLSNe found worldwide since we deployed FLEET in 2019 November (Gomez et al. 2022). Here we implement a similar approach and use the light curve and host galaxy information of a given transient to target likely TDEs without focusing on the classification of other transient classes. Since the goal of FLEET is to help make the most efficient use of telescope time, we choose to optimize for purity as opposed to completeness, in order to yield the purest possible sample of TDEs.

Unlike existing algorithms, FLEET is trained on real data and makes use of both light-curve information and contextual host galaxy parameters to predict the likelihood of a transient being a TDE without the need to know its redshift. FLEET is designed to be fast and capable of classifying thousands of transients within a few hours on a personal computer. Such ML algorithms as FLEET will prove not only useful but necessary as more transient surveys come online. We estimate the expected number of TDEs that could be detected by both Rubin and the Nancy Grace Roman Space Telescope (Roman; Spergel et al. 2015) and explore the possibility of using FLEET to target transients from these surveys and maximize our efficiency in recovering TDEs.

The structure of the paper is as follows. In Section 2, we outline the sources of the data used for training FLEET; in Section 3, we describe the underlying algorithm; in Section 4,

we explore possible selection effects in FLEET; in Section 5, we describe the use of FLEET in the ZTF, Rubin, and Roman time-domain surveys; and finally, we conclude in Section 6. FLEET is provided as a Python package on GitHub¹¹ and Zenodo (Gomez et al. 2020b), as well as included in the Python Package Index under the name `fleet-pipe`.

2. Data Sources

To build a training sample for FLEET, we gathered a list of all spectroscopically classified transients from the Transient Name Server (TNS).¹² We restrict the sample to only include transients from the ZTF (Bellm et al. 2019) that have at least two g -band and two r -band detections. These are the minimum data required to be able to fit a model to their light curves. We further restrict the sample to only include transients that lie within the footprint of the Pan-STARRS1 3π (PS1/ 3π) survey (Chambers & Pan-STARRS Team 2018), for the purpose of identifying their host galaxies. The resulting sample is composed of 4779 transients, with the following distinct labels from the TNS: 2983 SNe Ia, 749 SNe II, 187 SLSNe-I, 157 SNe IIn, 143 CVs, 105 SNe Ic, 89 SNe IIP, 80 SNe Ib, 68 SNe Iib, 52 SLSNe-II, 45 TDEs, 35 SNe Ic-BL, 26 SNe Ibc, 23 AGNs, 19 SNe Ibn, and 18 variable stars. We provide a list of the 45 TDEs used for this training sample in Table 2.

The light curves of all of the transients used for training come from ZTF and were obtained from the ALeRCE broker (Förster et al. 2021). While we only require four detections for a transient to be included in the training set, 90% of them have at least eight detections, and 50% have at least 22 detections in either the g or r band. We also query a $1'$ region around each transient in the PS1/ 3π (Chambers & Pan-STARRS Team 2018) and Sloan Digital Sky Survey (SDSS) catalogs (Alam et al. 2015; Ahumada et al. 2020) and include the photometry, radii, and separations to all of the nearby objects to the transient. Finally, we correct all photometry for Galactic extinction using the Schlafly & Finkbeiner (2011) estimates of $E(B - V)$ and the Barbary (2016) implementation of the Cardelli et al. (1989) extinction law.

3. FLEET Algorithm

FLEET was originally designed to target SLSNe from transient alert streams. Here we expand the capabilities of FLEET and train it to find TDEs. We provide a short description of the processing steps included in FLEET, but for a full description of the algorithm, see Gomez et al. (2020a).

3.1. Host Galaxy

TDEs are known to happen in the nuclei of galaxies (e.g., Gezari 2021), as is the case for all of the TDEs in our training sample. Additionally, the presence of TDEs is known to be correlated with galaxy type, since TDEs appear to be significantly overrepresented in poststarburst galaxies (e.g., French et al. 2016; Graur et al. 2018). Moreover, the presence of TDEs is therefore correlated with galaxy color, which largely tends to reside in the “green valley” of the color-magnitude diagram of galaxies (e.g., Hammerstein et al. 2021c). For these reasons, we test how well the relative separation, g - and r -band magnitudes, and half-light radius of a

¹¹ <https://github.com/gmzsebastian/FLEET>

¹² <https://www.wis-tns.org/>

galaxy serve to discriminate likely TDEs from other types of transients.

In order to obtain the properties of a host, we first need to associate a given transient with its most likely host galaxy. To do this, we query a $1'$ region of the PS1 and SDSS catalogs around each transient to identify its most likely host galaxy. First, we assign a probability of being a galaxy to every object in the field to rule out stars. We estimate this probability using a custom k nearest-neighbors algorithm trained on data from the Canada–France–Hawaii Telescope Legacy Survey (Hudelot et al. 2012), which uses the `CLASS_STAR` classifier flag in `SExtractor` to separate stars from galaxies, relying on a multilayer feed-forward neural network (Bertin & Arnouts 1996). Then, we calculate the probability of chance coincidence, P_{cc} , of every galaxy using the method of Bloom et al. (2002) described in Berger (2010) and select the galaxy with the lowest P_{cc} as the host galaxy of the transient. This method is computationally fast and accurate for $\sim 95\%$ of the associated transients, determined from manual vetting of the training set in Gomez et al. (2020a). This is slightly lower than more complex algorithms, for example, the success rate of $\sim 99\%$ reported by the DELIGHT algorithm (Förster et al. 2022) and $\sim 97\%$ from the GHOST classifier (Gagliano et al. 2021). The issue of host galaxy association is less ambiguous for TDEs, which, thanks to their nuclear nature, have a mean $P_{cc} \sim 10^{-3}$, and their host association is therefore effectively always successful, correctly identifying 100% of the hosts for the TDEs in our training sample. Compare this to the full population of transients in our training set, which have a mean $P_{cc} \sim 2.5 \times 10^{-3}$, a less confident value but one that still implies a very high likelihood of association.

3.2. Light Curve

In addition to host galaxy information, we use the light curves of transients to predict which are the most likely to be TDEs. TDEs are known to be relatively blue at early times, have broad long-lasting light curves, and fade without showing significant changes in their color (e.g., Mockler et al. 2019; van Velzen et al. 2021; Hammerstein et al. 2023; Nicholl et al. 2022). Therefore, deriving parameters from their light curves should provide critical information to separate TDEs from other transients.

We are motivated to classify TDEs early, which will allow us to trigger multiwavelength follow-up observations to both be able to robustly model their light curves and search for the early-time UV or X-ray emission found in some TDEs (e.g., Chornock et al. 2014; Wevers et al. 2019). Additionally, we aim for FLEET to be computationally efficient, and able to fit the light curves of thousands of transients at a time. Therefore, we fit the light curves of the transients with a simple, computationally efficient model,

$$m = e^{W(t-\phi)} - A \times W(t - \phi) + m_0, \quad (1)$$

where W is the effective width of the light curve, A modifies the decline time relative to the rise time, m_0 is the peak magnitude, and ϕ is a phase offset relative to the time of the first observation. We provide two versions of the model. One is aimed at finding TDEs at early times using only the first 20 days of data after discovery and has a fixed value of $A = 0.6$. We chose this value for A because it is the mean value for all transients with full light curves, with a mean standard deviation




Figure 1. Light curves of the TDE AT 2019ehz fit with the model described in Equation (1). The dashed lines show the fit using only data up to 20 days after detection (with a fixed value of $A = 0.6$), while the solid lines are the result of fitting the data up to 40 days after detection (with A as a free parameter). The former is part of our rapid classifier, while the latter is part of our late-time classifier. The absolute magnitudes are calculated based on the known TDE redshift of $z = 0.074$.

of 0.18. Specifically for TDEs, the value of A and its standard deviation is $A = 0.55 \pm 0.17$, similar to the value for the full population of transients. The second model uses the first 40 days of data and includes A as a free parameter. This model is able to more confidently predict a transient class but at the expense of triggering follow-up at a later phase. In Figure 1, we show an example of both models fit to a TDE light curve, with independent fits to the g - and r -band light curves.

3.3. Feature Optimization

FLEET uses the RF implementation from the `scikit-learn` Python package (Pedregosa et al. 2011) to predict the probability of a transient being a TDE, $P(\text{TDE})$. Our training set used to train this classifier is made of 4779 spectroscopically classified transients, described in Section 2. We obtain the host galaxy and light-curve information described in Sections 3.1 and 3.2 for each transient and optimize which features have the most predictive accuracy.

The training set of all classified transients is unbalanced, meaning not every class of transient has the same number of events. To prevent the classifier from being biased toward predicting the more common classes and to improve predictive accuracy, we oversample all transients to have a population size equal to that of the largest transient class (i.e., 2983 SNe Ia) using the Synthetic Minority Over-sampling Technique (Chawla et al. 2002). This algorithm generates random samples of new features drawn along vectors joining every pair of existing objects in feature space until the desired number of samples has been reached. In Gomez et al. (2020a), we determined the optimal grouping of transient classes to be nuclear, SLSN-I, SLSN-II, SN II, SN IIb, SN IIc, SN Ia, SN Ibc, and star. The only difference in the class labels between the TDE classifier and the original SLSN classifier is that the former “nuclear” class has been split into distinct “AGN” and “TDE” classes. Allowing the classifier to group transients into these distinct classes improves its overall accuracy. Nevertheless, we are not concerned with the individual classifications of other transients, since they are eventually compressed into a binary “TDE” versus “non-TDE” classification. The uncertainties for all predictions from the RF algorithm presented here

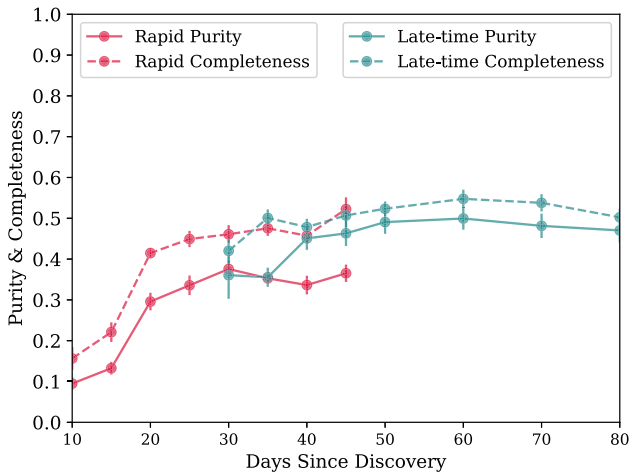


Figure 2. Purity and completeness as a function of days of photometry included in the models. The dashed lines represent completeness, while the solid lines represent purity. The rapid model is shown in red, while the late-time model is shown in teal.

represent the 1σ scatter of 25 different realizations of each model, generated using a different initial random seed.

FLEET is able to classify thousands of transients within a few hours on a personal computer. Classifying a new transient takes on the order of 10–20 s. Once the required catalog data and an image of the host galaxy from PS1/3 π have been downloaded (~ 100 KB each), the required time to reclassify a transient is about 5–10 s.

We optimize several hyperparameters of the RF classifier, including the number of days of photometry to consider, the features to be included in each model, and the depth of the RF trees, each one optimized using a grid search described in more detail in this section. The two metrics we optimize for are “completeness” and “purity.” Completeness is defined as the total number of true-positive TDEs divided by the total number of TDEs, and purity is the total number of true-positive TDEs divided by the sum of true-positive TDEs and false-positive TDEs.

Our predictive accuracy improves as a transient evolves and more light-curve data are used but at the expense of identifying the transient at a later phase. We optimize the number of days of photometry to include, in both the rapid and late-time classifiers, testing the performance of FLEET using a grid of 10–80 days of photometry in steps of 5 days for up to 50 days, then steps of 10 for the longer timescales (Figure 2). We find that for the rapid classifier, including 20 days of photometry is the minimum required to produce reasonable results, since including only 15 days of photometry reduces the purity and completeness by about half. Conversely, we find that including 40 days of photometry in the late-time classifier significantly increases the completeness from $\approx 35\%$ to $\approx 50\%$ compared to using only 35 days. These thresholds are reasonable when we compare to the typical rise times of the transients in our training set, 95% of which peak $\lesssim 30$ days after discovery. Due to the nature of the ZTF survey, some transients will have multiple observations in a single day in a specific band. Therefore, a measure of the number of data points included is not as constraining as the days of photometry included. We find that including more than 40 days of photometry does not significantly improve either the purity or completeness. Including more than ~ 70 days of photometry begins to

degrade our metrics, since these late-time data might be of lower quality with higher scatter or include phenomena such as flattening or light-curve bumps that are not accounted for in our model.

Additionally, we optimize the set of features to include in the classifier and determine the optimal set that yields the highest purity. We list this optimal set of features below, where we use one W , A , and Δm for each g and r band.

1. W : width parameter of the light curve.
2. A : multiplicative modifier for the decline rate of the light curve (only used for late-time model).
3. R_n : transient–host separation normalized by half-light radius of the host.
4. Δt : time of peak brightness minus time of discovery.
5. Δm : host magnitude minus peak observed transient magnitude.
6. $(g-r)_P$: light-curve color at peak.
7. $(g-r)_L$: late-time light-curve color 40 days after peak.

The half-light radius in the r band of the host galaxy is obtained from the SDSS catalog or from the PS1/3 π catalog if the object is not in SDSS. The $g-r$ color is measured from the light-curve fits, as opposed to the photometry, since we found this to be a more uniform method that allows us to measure an estimate for the color even for transients with sparse data coverage. We tested the effects of including other features in the classifier, including the photometric redshift of the host galaxy from SDSS, the $g-r$ color measured from the brightest observed photometry data points, the host separation, the host radius, a χ^2 estimate for the goodness of fit of the model, and the probability of chance coincidence P_{cc} . We found that including these features either hurt or did not improve our metrics. The optimal set of features was performed by drawing 200 random selections of features of random size from the full sample of 13 features. We then selected the best five sets and added or removed each individual feature to fine-tune the optimization and reach the final selection listed above.

The importance of each included feature is not defined independently of other features; if two features are correlated, then their relative importance could be affected. We use the permutation importance method described in Breiman (2001) to calculate the correlated importance of each feature and show the results for the late-time classifier in the top panel of Figure 3. In the bottom panel of Figure 3, we show the correlation between features and find that, with the exception of a strong correlation between Δm_g and Δm_r , the individual features are mostly independent. We find the most important features to be R_n and $(g-r)_P$ for both the rapid and late-time classifiers, as well as $(g-r)_L$ for the late-time classifier.

In Figure 4, we plot the phase space of our most important features that distinguish TDEs from the other transients used for training. The figures show how R_n and $(g-r)_P$ can help separate TDEs from other transients. A low R_n value is obviously a good discriminant, since TDEs happen in the nuclei of galaxies. The $(g-r)_P$ can help separate TDEs from AGNs and SNe, since TDEs tend to be bluer. Lastly, $(g-r)_L$ becomes important when selecting TDEs, since their colors tend to remain fairly constant as they evolve, unlike other transients (e.g., van Velzen et al. 2021; Hammerstein et al. 2023; Nicholl et al. 2022).

Finally, we fit for the optimal depth of the RF trees. We optimize the tree depth by running a grid search from a depth of

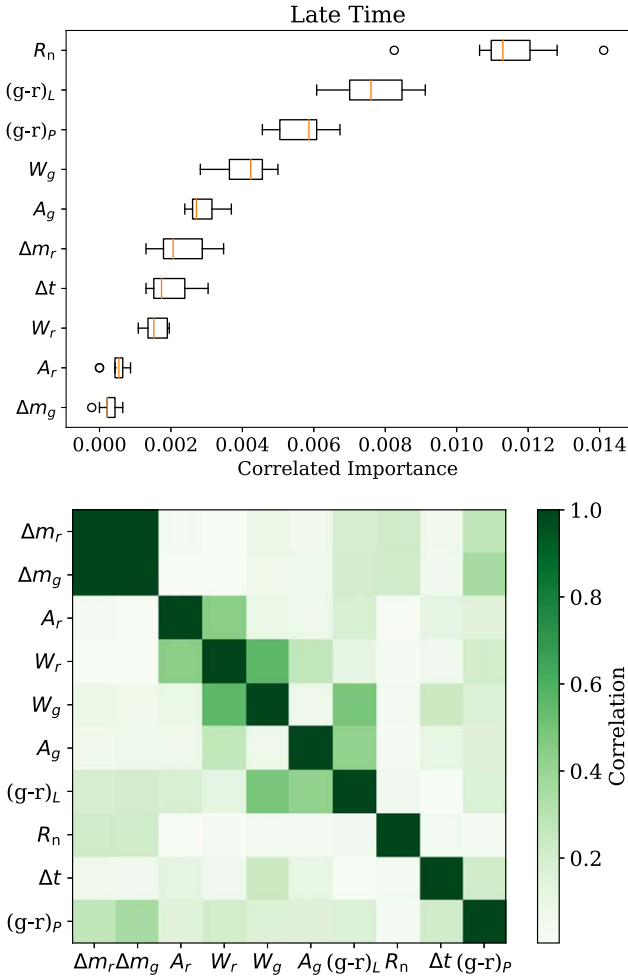


Figure 3. Top: correlated importance for the features used in the late-time version of the classifier. We find that the host separation R_n and transient colors $(g-r)_P$ and $(g-r)_L$ appear to be the most important features for predictive accuracy. Bottom: correlation matrix for the same features. Except for Δm_g and Δm_r , most features show a low correlation.

5–20 in steps of 3, then selecting the three values with the highest purity and repeating the same test using a finer grid with steps of 1. We find an optimal depth of 14 for the rapid classifier and 17 for the late-time classifier.

For completeness, we test how a classifier that only uses light-curve information would perform. We find that using light-curve information alone reduces the peak purity of the late-time classifier to a relatively reasonable $\approx 25\%$ at a completeness of $\approx 15\%$. On the other hand, using light-curve information alone reduces the purity and completeness of the rapid classifier to very low values, below $\approx 10\%$.

3.4. Validation

For validation of the algorithm, we implement a leave-one-out cross-validation method. This method trains the classifier using every transient except one, predicts the classification of the one transient, and then repeats this process, cycling through all transients. This allows us to robustly test our classifier without having to divide the data set into training and test sets, which would compromise the already small sample of TDEs. We use three different methods to evaluate the performance of our classifier: a confusion matrix, a purity curve, and a completeness curve. Since we are not concerned with the

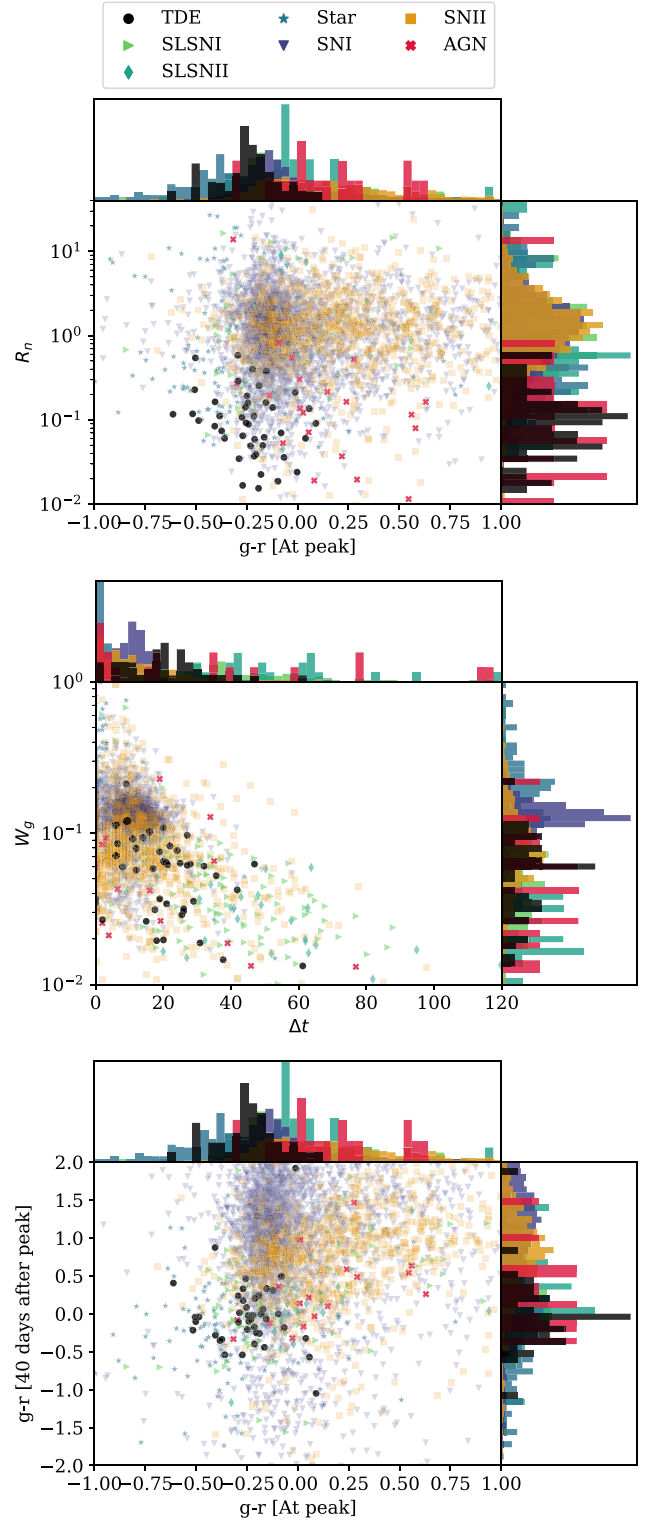


Figure 4. Phase space of the most relevant features used for the classifier, showing the various classes of transients in different colors/markers. Top: normalized host separation (R_n) vs. color of the transient during peak $(g-r)_P$. Transients without a detected host are not included in this plot. Middle: width of the light curve in the g band (W_g) compared to the time to peak (Δt). Bottom: color of the transient during peak $(g-r)_P$ compared to the $(g-r)_L$ color 40 days after peak.

classification of transients other than TDEs, we collapse the individual transient classification categories listed in Section 3 into a binary “TDE” versus “non-TDE” classification. To

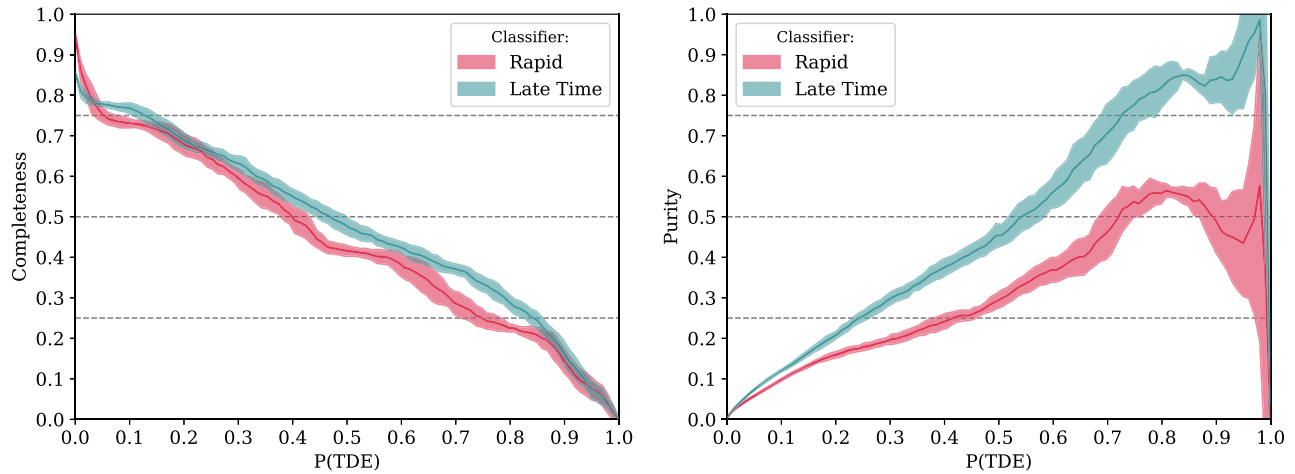


Figure 5. Left: completeness as a function of classification confidence for the rapid and late-time classifiers. Right: purity as a function of classification confidence for the rapid and late-time classifiers. The shaded regions represent 1σ uncertainties, calculated by generating each model using 25 different random seeds. We find that the completeness decreases monotonically until $P(\text{TDE}) \sim 0.8$, after which it decreases at a faster rate due to the small sample size of TDEs with $P(\text{TDE})$ above this value. For the same reason, the uncertainty in the purity increases significantly for values above $P(\text{TDE}) \sim 0.8$.

calculate the non-TDE probability for each transient, we sum the probabilities of all other transient classes.

The purity of the classifier increases and the completeness declines as we restrict the sample to events with progressively higher values of classification confidence (Figure 5). For $P(\text{TDE}) > 0.5$, the purity is $\approx 45\%$ for the late-time classifier and $\approx 30\%$ for the rapid classifier at a corresponding completeness of $\approx 50\%$ and $\approx 40\%$, respectively. This represents a factor of ~ 60 improvement over a random selection of transients, which would yield an $\approx 0.5\%$ success rate in a magnitude-limited survey (Villar et al. 2019; Fremling et al. 2020). To estimate the uncertainty in our measurements, we run each RF model 25 times with different initial random seeds to create the uncertainty regions shown in Figure 5.

In Figure 6, we show the confusion matrices, namely, the label predicted by our classifier compared to the true label of the transient, normalized in terms of both purity and completeness. We impose a confidence cut of $P > 0.75$ for either the TDE or non-TDE classes, corresponding to the peak classifier purity. We find that 64% ($N = 17$) and 45% ($N = 8$) of true TDEs in this subsample of transients with $P > 0.75$ were predicted to be TDEs by the late-time and rapid classifiers, respectively. The matrices also show that 81% ($N = 17$) and 62% ($N = 8$) of all transients predicted to be TDEs were true TDEs for the late-time and rapid classifiers, respectively.

In Figure 7, we show how the rapid and late-time classifiers perform at classifying TDEs and not misclassifying other objects as a function of classification confidence level, $P(\text{TDE})$. We find that half of the correctly identified TDEs have a $P(\text{TDE}) \gtrsim 0.6$, but some true TDEs are still misclassified with a low $P(\text{TDE}) \lesssim 0.1$. Some of the misclassified TDEs include AT 2018jbv, AT 2020ddv, AT 2020opy, and AT 2020riz, which were instead classified as likely SLSNe due to their relatively high Δm , and AT 2019gte and AT 2020neh, which were classified as likely SNe Ia due to their relatively fast light curves and lack of late-time photometry.

As an additional point of comparison, we test how the performance of FLEET compares to the performance of a more classical approach of targeting TDEs by imposing basic selection cuts on their observational parameters (e.g., van Velzen et al. 2021). We test the effectiveness of using selection cuts on the two most important parameters, $(g-r)_P$ and R_n . We

find that imposing a cut and only selecting targets with $(g-r)_P < -0.5$ mag yields the highest purity of $\sim 5\%$ at a corresponding completeness of $\sim 10\%$. Alternatively, imposing a cut of $(g-r)_P < 0$ mag yields a much higher completeness of $\sim 85\%$ but with an even lower purity of $\sim 3\%$. Similarly, we determine $R_n < 0.14$ to be the optimal threshold that yields the highest purity of $\sim 7\%$ at a corresponding completeness of $\sim 60\%$. The purity obtained from either a cut on $(g-r)_P$ or R_n is well below the $\approx 50\%$ purity we can obtain from FLEET, even with the rapid classifier. We further attempt to optimize the purity of these selection cuts by implementing them in conjunction. We find that a combined threshold of $(g-r)_P < -0.5$ mag and $R_n < 0.12$ would yield the highest possible purity of $\sim 38\%$. This purity is much closer to the estimates from FLEET but comes at the expense of an $\sim 6\%$ completeness, much lower than the $\approx 50\%$ completeness that FLEET can achieve at the same purity level. Therefore, we have shown that FLEET is expected to perform much better than imposing basic selection cuts.

3.5. Classifier Summary

Using our rapid classifier, trained using host galaxy information and the first 20 days of photometry of 45 TDEs, we predict that we can recover TDEs with $\approx 50\%$ purity at $\approx 30\%$ completeness for $P(\text{TDE}) > 0.7$. With our late-time classifier, trained instead on 40 days of photometry, we can recover TDEs with a purity of $\approx 70\%$ and a corresponding completeness of $\approx 40\%$ for the same $P(\text{TDE}) > 0.7$ threshold.

Currently, time-domain surveys report $\sim 20,000$ transients yr^{-1} , and we expect $\sim 0.5\%$ of them to be TDEs given current observational rates, corresponding to ~ 100 TDEs that could be discovered every year. Given our predictions for purity and completeness, using FLEET would result in ~ 15 – 28 TDEs discovered per year, depending on the classifier used. This is about twice as many as the 8–16 TDEs classified each year since 2019.

4. Selection Effects

We explore possible selection effects that might result from using FLEET for target selection. For testing, we classify the 32 TDEs presented in Nicholl et al. (2022) using the late-time

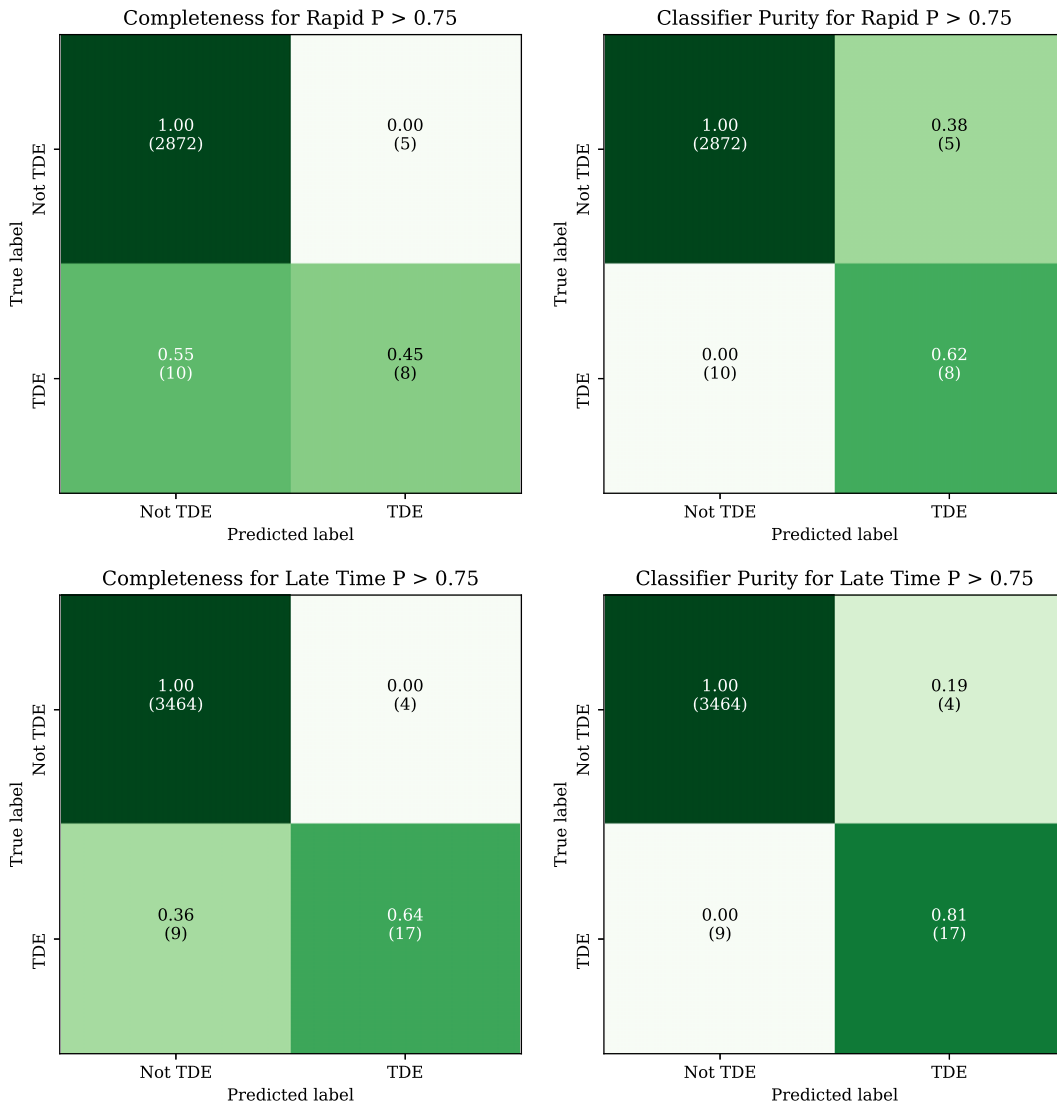


Figure 6. Using a sample of only transients with a classification probability of $P(\text{TDE}) > 0.75$ or $P(\text{non-TDEs}) > 0.75$, we produce a confusion matrix that indicates a purity of 62% for the rapid classifier and 81% for the late-time classifier. The corresponding completeness values for the same classifiers are 45% and 64%. These values represent the raw number of TDEs used in the classifier and are not weighted by the number of samples in each class.

FLEET classifier and compare the observational parameters of the TDEs with $P(\text{TDE}) > 0.5$ with those of all TDEs. To quantify if the parameter distributions are different, we implement a two-sample Kolmogorov–Smirnov (K-S), where a K-S metric of $D = 0$ indicates that the two samples are drawn from the same distribution, and $D = 1$ means there is no overlap between the distributions. In Figure 8, we show how the two populations compare and find no obvious bias against most parameters. We determine K-S metrics (and p values) of 0.14 (0.93) for redshift, 0.12 (0.98) for apparent magnitude, 0.21 (0.52) for R_n , 0.21 (0.52) for Δm , 0.16 (0.85) for $(g-r)_p$, 0.11 (0.99) for W_g and W_r , and 0.19 (0.64) for $(g-r)_L$, meaning that the two populations differ at the $\sim 10\%$ – 20% level. Additionally, we implement an Anderson–Darling test (Scholz & Stephens 1987) but are unable to reject the null hypothesis that any pair of parameter samples are drawn from the same distribution with a p value < 0.25 .

For completeness, we determine the corresponding K-S metrics for the parameters obtained using the rapid classifier. We find K-S metrics of 0.12 (0.97) for apparent magnitude, 0.18 (0.75) for Δm , 0.21 (0.52) for $(g-r)_p$, 0.07 (0.99) for W_g ,

0.11 (0.99) for W_r , and 0.14 (0.93) for $(g-r)_L$, showing that these metrics are comparable to the results obtained from the late-time classifier. We do not include the host-dependent features, since they do not depend on the classifier used.

Most of the parameters shown in Figure 8 show little difference between the full population of TDEs and the population of recovered TDEs with $P(\text{TDE}) > 0.5$. The biggest difference is for R_n , where TDEs with $R_n \gtrsim 0.3$ do not tend to be selected by FLEET. This is not to say that transients with $R_n > 0.3$ are not nuclear but simply that the uncertainty in their coordinates is high but still consistent with being nuclear. This uncertainty is dominated by the typical scatter in ZTF coordinates, which can be up to ~ 0.5 . We find that TDEs with $\Delta m \gtrsim 0$ or $W_g \lesssim -0.1$ are also less likely to be selected as TDEs.

Similarly, we explore the difference in physical parameters for the recovered TDEs compared to the full TDE sample. For the exploration of physical parameters, we use the TDE models from Nicholl et al. (2022), who fit the light curves of 32 TDEs using MOSFIT (Guillochon et al. 2018). We explore the differences in the impact parameter b , black hole mass M_{BH} ,

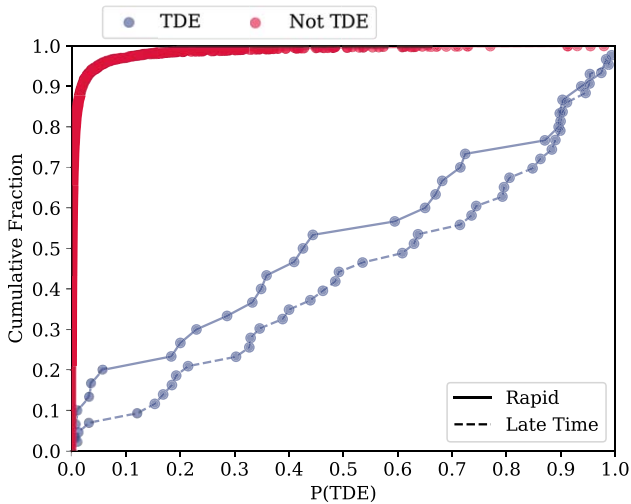


Figure 7. Cumulative distribution as a function of classification confidence P (TDE) for transients that are true TDEs (blue) or non-TDEs (red). We find that most non-TDEs are correctly identified as such, while half of the true TDEs have a $P(\text{TDE}) \gtrsim 0.6$. The rapid and late-time classifier lines for non-TDEs appear indistinguishable from each other in the plot.

stellar mass M_* , and viscous timescale T_V and find no obvious selection effects. We find a K-S metric (and p value) of 0.12 (0.94) for b , 0.12 (0.99) for M_{BH} , 0.18 (0.67) for M_* , and 0.12 (0.97) for T_V , indicating that the parameter distributions for the full TDE sample and the sample of recovered TDEs with $P(\text{TDE}) > 0.5$ are very similar and only differ at the $\sim 10\%$ level.

In conclusion, we find that even if FLEET selects events with certain observational properties, these do not translate to selection effects in physical parameters. In other words, FLEET does not appear to produce biases against physical or observational parameters not included as features in the classifier. Of course, this is in relation to the known sample of TDEs; finding anomalous TDEs not found in the current known sample, such as off-nuclear or white dwarf TDEs, will likely require larger samples from systematic surveys such as Rubin or Roman.

5. Implementation of FLEET for Time-domain Surveys

5.1. ZTF

The ZTF is the survey that currently reports the most transients to the TNS, but $\sim 90\%$ of them remain spectroscopically unclassified. Here we use the late-time FLEET classifier to obtain a $P(\text{TDE})$ estimate for all ZTF transients reported to the TNS with the aim of recovering TDEs that were previously missed.

At the time of writing, there were a total of 95,729 transients reported to the TNS. Since we are using the late-time classifier, we exclude transients that are too young, discovered after 2022 March 1. Additionally, FLEET only works for transients within the PS1/ 3π footprint and with at least two g -band and two r -band points. After applying these cuts, we run FLEET on a final list of 31,892 transients and list 39 likely TDE candidates with $P(\text{TDE}) \geq 0.5$ obtained from this experiment in Table 1. Follow-up and analysis of these targets will be presented in future work.

Given their individual probabilities of being TDEs, if we were to spectroscopically classify these 39 TDE candidates, we would expect ~ 28 of them to be confirmed as TDEs, which would represent a factor of $\sim 50\%$ increase in the current population of TDEs. These tools will be more powerful for

larger data sets like the ones Roman or Rubin will produce, allowing us to generate large photometrically selected samples of TDEs and analyze their statistical demographics without the need for spectra.

One limitation that ZTF has compared to future surveys is its relatively shallow depth of $r \sim 20.5$ mag. In Figure 9, we show the fraction of spectroscopically classified TDEs correctly identified as TDEs with $P(\text{TDE}) \geq 0.5$ as a function of redshift. We find that the increasingly poor data quality from ZTF makes it challenging to accurately classify TDEs at $z \gtrsim 0.15$. Future deeper surveys will be able to discover TDEs at much further redshifts.

5.2. Roman

Roman is an infrared wide-field survey telescope planned to commence observations in 2027 (Spergel et al. 2015). Major components of the Roman mission are the community surveys, large programs that will occupy a large fraction of the mission and are meant to provide data useful for the community at large. One such survey is the High Latitude Time Domain Survey (HLTDS), which focuses on the study of extragalactic transients (Rose et al. 2021). Here we estimate the number of TDEs that could be found by the Roman HLTDS and explore the possibility of using FLEET to find them.

The final survey parameters of the HLTDS have not been defined, but we simulate a Roman HLTDS observation sequence based on the parameters suggested in Rose et al. (2021). We exclude the other two community surveys from our simulation, since the Galactic Bulge Time Domain Survey will be heavily crowded by stars and affected by Galactic extinction, and the High Latitude Wide Area Survey has no time-domain component. We run a simulation that consists of 146 visits of a patch of sky located at R.A. = 25° , decl. = -53° with a cadence of 5 days. Each visit has a wide and a narrow component. The wide component covers an area of 15.6 deg^2 in 56 pointings with filters (and 5σ limits) F062 (26.4), F087 (25.6), F106 (25.5), and F129 (25.4), while the narrow component covers an overlapping area of 5.6 deg^2 in 20 pointings with filters (and 5σ limits) F106 (26.7), F129 (26.6), F158 (26.5), and F184 (26.7).¹³

We use the TDE MOSFIT models from Nicholl et al. (2022) to generate a sample of 240,000 TDEs distributed uniformly across space (R.A., decl., and comoving distance) and time (throughout the 2 yr duration of the HLTDS). Each TDE sample is generated from observations of a real TDE, which allows us to simply adopt the properties of existing TDEs without introducing additional parameter assumptions. MOSFIT generates a blackbody spectral energy distribution (SED) at each epoch of a TDE's evolution, allowing us to convolve this SED with the Roman passbands and measure their corresponding IR magnitudes. We then compare these magnitudes to the nominal 5σ limits of the HLTDS to quantify how many TDEs would be detected.

In Figure 10, we show the total number of TDEs per year expected to be found with at least four detections as part of the HLTDS. We adopt the volumetric TDE rate estimate from van Velzen (2018), which is defined as a function of TDE peak g -band luminosity as $\dot{N}_0 (L/L_0)^a$, where $\dot{N}_0 = (1.9 \pm 0.7) \times 10^{-7} \text{ Mpc}^{-3} \text{ yr}^{-1}$ and $a = -1.6 \pm 0.2$ for

¹³ For a definition of the Roman filters, see https://roman.gsfc.nasa.gov/science/WFI_technical.html.

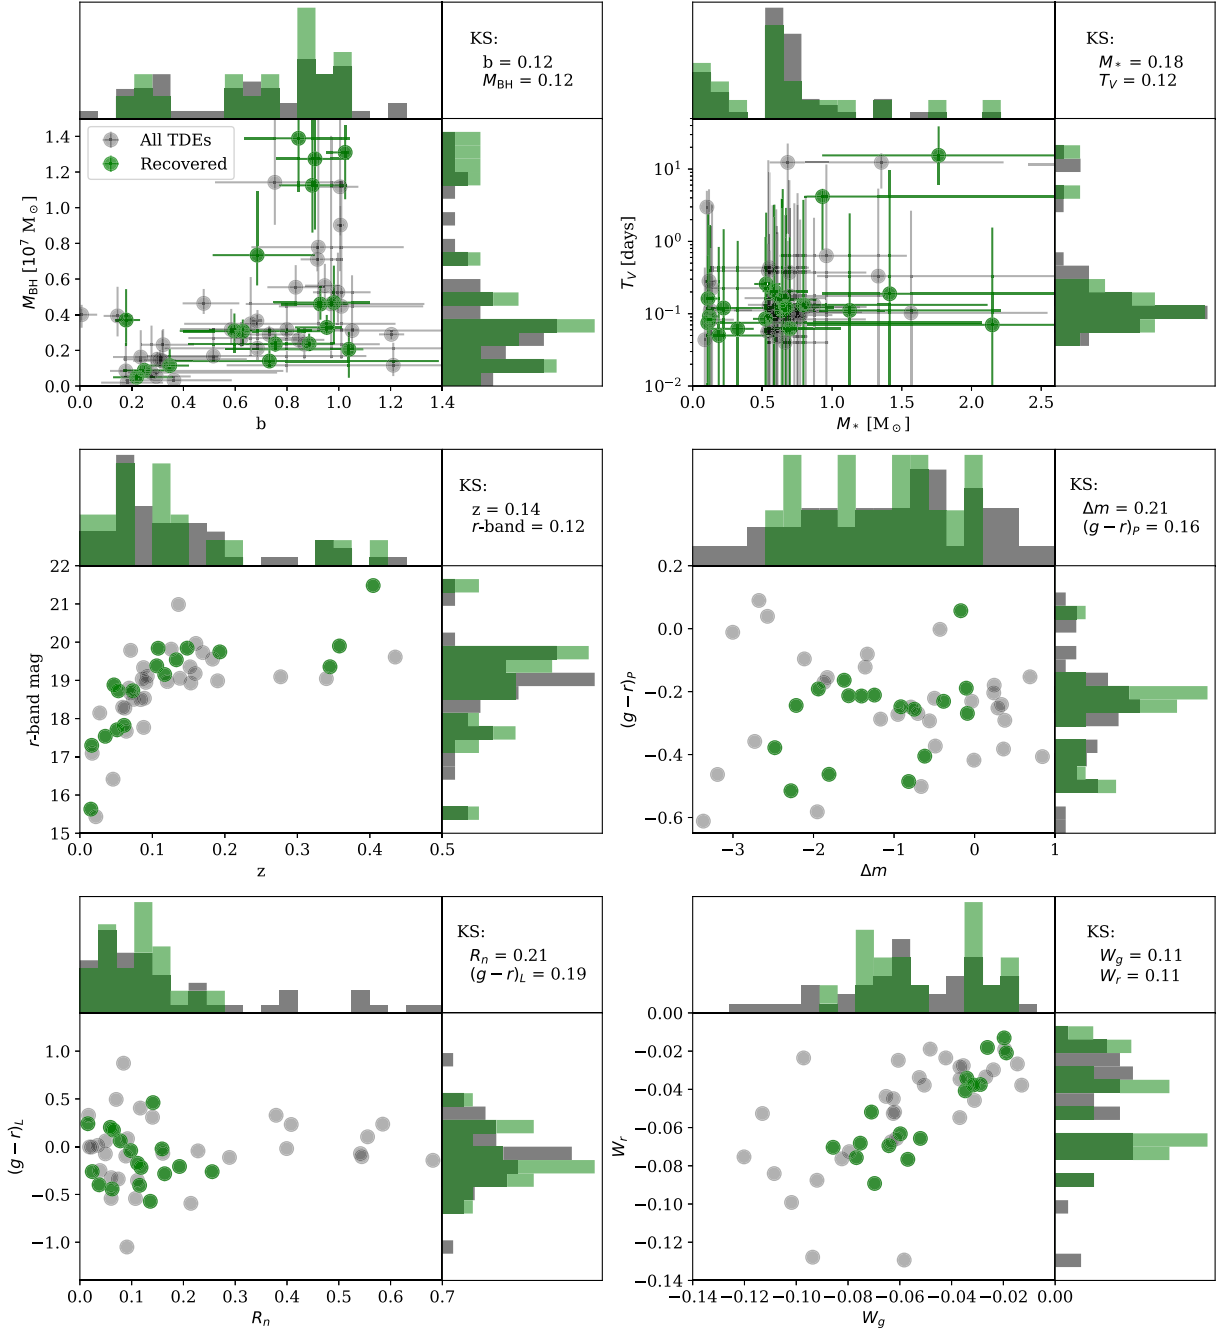


Figure 8. The gray data points show the physical and observational parameters from a sample of TDEs obtained from Nicholl et al. (2022). The green data points are TDEs that were recovered by FLEET with $P(\text{TDE}) > 0.5$. Top left: black hole mass and impact parameter. Top right: stellar mass and viscous timescale. Middle left: redshift and apparent peak r -band magnitude. Middle right: Δm compared to the $(g-r)_p$ color at peak. Bottom left: late-time color $(g-r)_L$ compared to the normalized host separation R_n . Bottom right: width of the light curve in the g and r bands. We include the K-S metric for each parameter in the upper right corner of each panel. We find no obvious bias against physical or observational parameters that are not already features of the classifier.

$L_0 = 10^{43} \text{ erg s}^{-1}$. Most TDEs in our sample have peak g -band luminosities of $\sim 10^{42} - 10^{43} \text{ erg s}^{-1}$, which translates to a rate of $\sim 10^{-6} - 10^{-7} \text{ TDEs Mpc}^{-3} \text{ yr}^{-1}$. Using these estimates, we predict that the Roman HLTDS could find $\sim 300 \text{ TDEs yr}^{-1}$ out to a redshift of $z \sim 7$ with at least four detections each, or ~ 200 “well-observed” TDEs with at least 20 detections spanning a minimum 70 day baseline between the first and last detection. It is also evident from Figure 10 that the distribution of TDEs found by the HLTDS peaks around a redshift of $z \sim 1-2$.

We find that out of all TDEs injected into our simulation, $\sim 35\%$ of the ones at $z < 0.2$ were detected with at least four data points in either the deep or wide components of the HLTDS. If we only consider the deep component of the HLTDS, we find that $\sim 81\%$ of TDEs at $z < 0.2$ are detected with at least four data points. In the top panels of Figure 11, we show some example light curves of TDEs detected by our simulation of the Roman HLTDS. It is clear that while Roman will find TDEs at high redshifts, the most distant ones at $z \sim 6$ will have poor coverage and are only detectable by the deep

Table 1
Likely TDEs from ZTF

Name	$P(\text{TDE})$	Name	$P(\text{TDE})$	Name	$P(\text{TDE})$	Name	$P(\text{TDE})$
2021crk	0.97	2021ldl	0.80	2020kri	0.67	2020aexc	0.58
2019gtm	0.97	2020ygl	0.80	2021pqg	0.66	2021qhb	0.57
2021aees	0.91	2019zbt	0.80	2018jil	0.66	2020afap	0.57
2021aeuf	0.90	2021zvy	0.72	2021uhs	0.64	2019aamf	0.57
2021lsi	0.86	2021ony	0.71	2021xbz	0.63	2021abkx	0.56
2020bgf	0.86	2020qfm	0.71	2021acqt	0.63	2019enj	0.56
2021kqp	0.85	2019phf	0.71	2019aami	0.62	2021aux	0.52
2021wdh	0.82	2021icq	0.70	2021nhq	0.61	2020qmx	0.51
2021iqs	0.82	2019qdh	0.70	2019saj	0.60	2020aawn	0.50
2020dxw	0.81	2019wlv	0.69	2020nmc	0.58		

Note. List of 39 unclassified ZTF transients with a high likelihood of being a TDE, ordered by $P(\text{TDE})$.

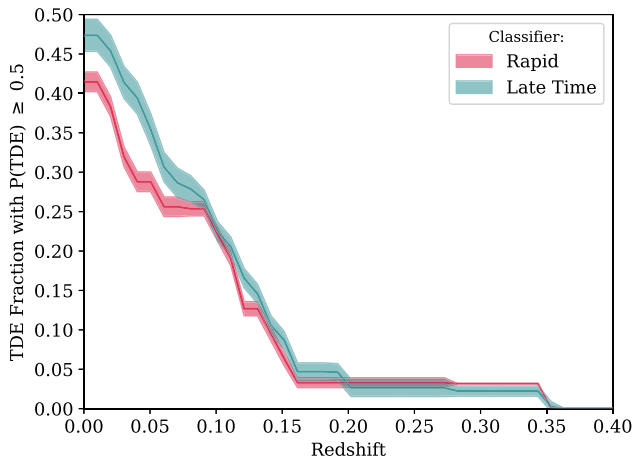


Figure 9. Fraction of TDEs with $P(\text{TDE}) \geq 0.5$ as a function of redshift with both the rapid and late-time classifiers. The decreasing quality of ZTF data for TDEs with $z \gtrsim 0.15$ makes it increasingly challenging to classify them.

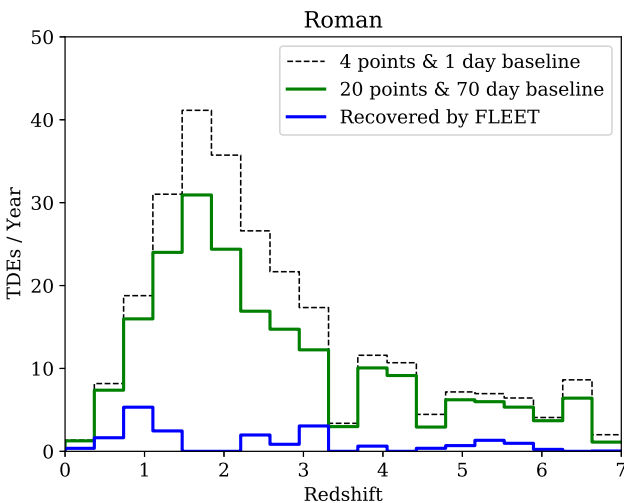


Figure 10. Number of TDEs expected to be found by the Roman HLTDS per year per redshift bin. The black dashed line shows the total number of TDEs with at least four detections spanning at least a 1 day baseline, while the green solid line shows only the “well-observed” TDEs with at least 20 detections over a 70 day baseline. The blue line shows the TDEs recovered by FLEET with $P(\text{TDE}) > 0.5$.

component of the HLTDS, which would make their characterization challenging. The wide component of the HLTDS can only detect TDEs out to a redshift of $z \sim 4$. Finding TDEs at

redshifts this high will allow us to test different theories regarding the evolution of the TDE rate. For example, Kochanek (2016) predicted that the rate of TDEs is expected to decrease at higher redshifts.

Given that Roman does not have g - or r -band observations, we cannot directly run FLEET on the simulated Roman light curves. Nevertheless, we can adopt the most frequent Roman bands (F106 and F129) as proxies for g and r and run the late-time FLEET classifier on these. We show the fraction of TDEs recovered with $P(\text{TDE}) > 0.5$ as a function of redshift in Figure 13. We determine that FLEET can recover with $P(\text{TDE}) > 0.5$ about 10% of Roman TDEs at $z < 3.0$ or about 20% of TDEs at $z < 0.5$. This translates to ~ 20 TDEs that FLEET could uncover every year out to a redshift of $z = 0.5$ or ~ 40 TDEs out to a redshift of $z = 3.0$. We caution that these estimates assume that FLEET would perform equally well on Roman and ZTF data. Eventually, we will be able to retrain FLEET using Roman data, including the host galaxy images from the Wide Area Survey, the light curves from the HLTDS, and the spectroscopic classifications for the $\sim 10\%$ of detected transients that are expected to be produced by the Roman slitless spectrograph (Rose et al. 2021). But currently, our approximation depends on a number of factors. First, the difference in optical versus IR filters is likely to have a negative effect only for the closest TDEs at $z < 0.3$; since TDEs appear blue at early times, they will be harder to distinguish from nuclear SNe in IR wavelengths. Nevertheless, this will likely not be a problem for the majority of TDEs discovered by Roman, since most TDEs are expected to be detected at $z \sim 1-2$, which, for the F129 and F158 filters, approximately corresponds to the rest-frame g and r bands. The fact that we extrapolate a blackbody SED to calculate the magnitudes of the TDEs in the IR bands is possibly a conservative estimate given that Lu & Bonnerot (2020) predicted TDEs to be brighter in the IR than what a simple blackbody predicts. At least three of the TDEs in our sample (AT 2018iih, AT 2018zr, and AT 2018hyz) have been shown to have late-time IR emission from dust formation. Additionally, Roman will observe these transients with at least four filters, likely improving our selection criteria thanks to the doubling in bands available. Moreover, we do not account for the higher angular resolution of Roman, which will likely increase our purity, as it will allow us to better discern whether a transient is nuclear compared to what we can currently do with ZTF. Lastly, we do not account for intrinsic host galaxy extinction of the transients from Roman, which would only be different for the nearby TDEs observed in the rest-frame IR. In conclusion, our algorithms are

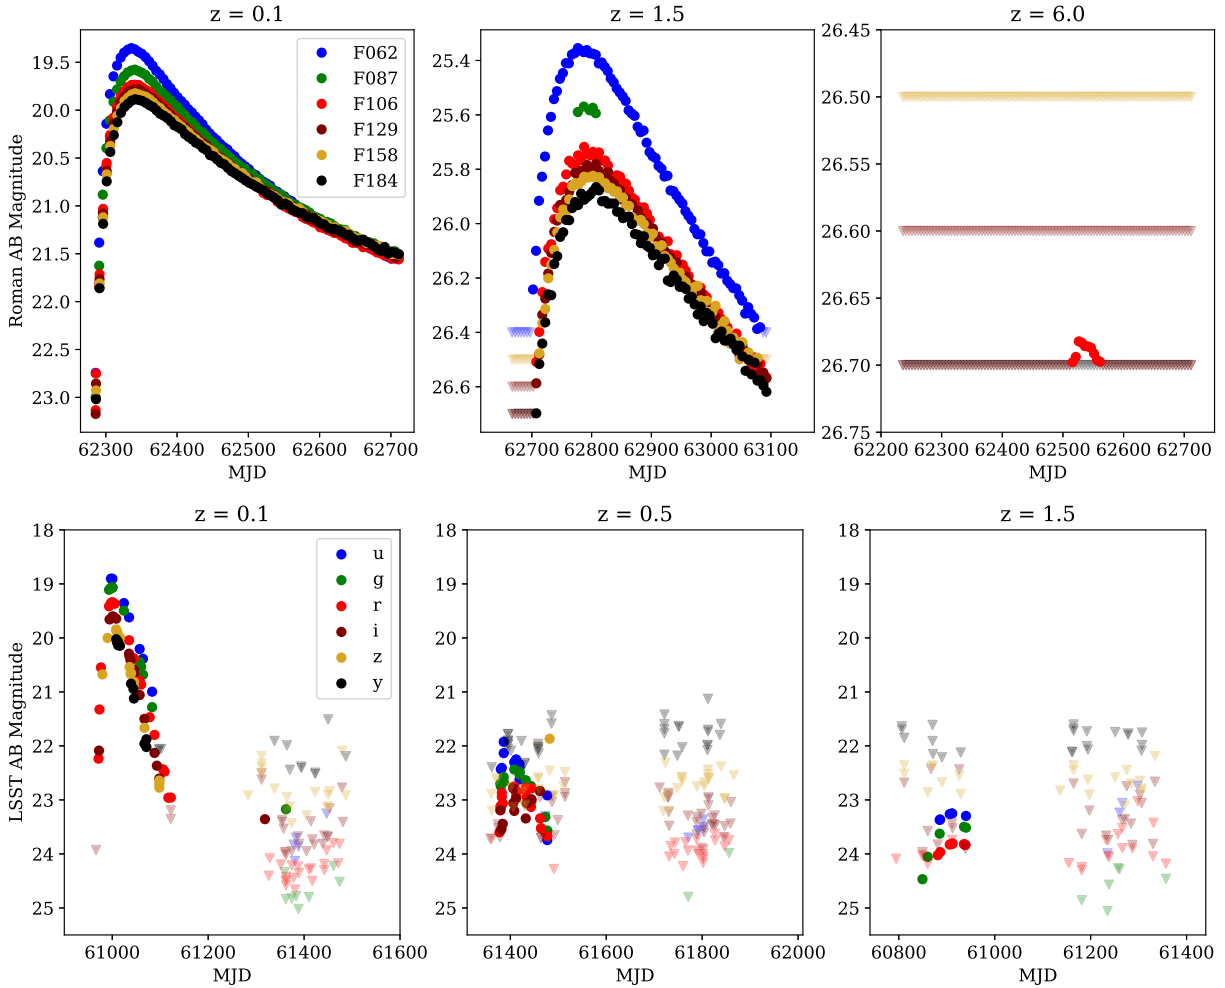


Figure 11. Representative model light curves of TDEs from the Roman HLTDS (top) and Rubin (bottom) simulated surveys. We show three light curves for each survey at three different redshifts.

likely to perform similarly well for high-redshift TDEs, but it is uncertain how well they will perform for the closest TDEs with rest-frame IR observations. Obtaining a more realistic estimate of the number of TDEs that will be found by Roman will require simulations of the Roman survey, noise properties, images, template subtraction capabilities and artifacts, and eventually real data from Roman on which to train the FLEET algorithm, including the spectroscopic data for $\sim 10\%$ of detected transients that are expected to be produced by the Roman slitless spectrograph (Rose et al. 2021).

Finally, we note that real-time follow-up of TDEs from Roman will completely rely on the availability of real-time transient alerts. Therefore, we urge the implementation of a real-time transient alert system to be part of the Roman survey.

5.3. Rubin

We estimate the number of TDEs that could be found by Rubin (LSST Science Collaboration et al. 2017) following the methods outlined in Villar et al. (2018). Here we inject 30,000 TDE light curves into the 10 yr Rubin baseline simulation (baseline_v2.2_10yrs.db) using the same methods for generating and injecting light curves outlined in Section 5.2. Fewer random draws are used for Rubin due to the computational limitations imposed by the more complex Rubin survey simulation.

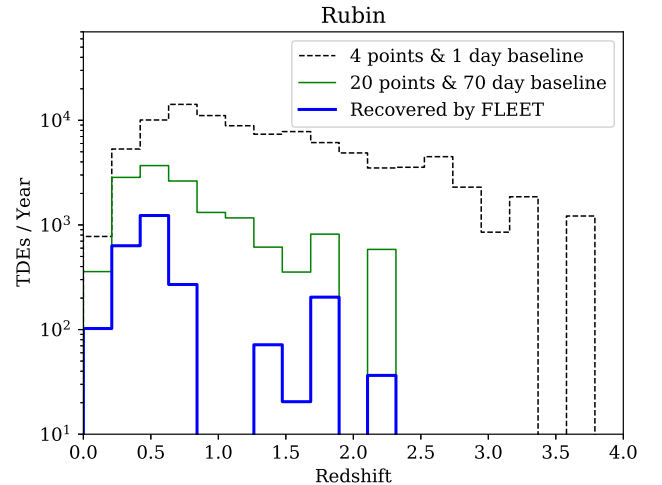


Figure 12. Number of TDEs expected to be found by Rubin per year per redshift bin. The black dashed line shows the total number of TDEs with at least four detections spanning at least a 1 day baseline, while the green solid line shows only the “well-observed” TDEs with at least 20 detections over a 70 day baseline. The blue line shows the TDEs recovered by FLEET with $P(\text{TDE}) > 0.5$.

In Figure 12, we show the number of TDEs expected to be found by Rubin assuming the same volumetric rate from van Velzen (2018) described in Section 5.2. We find that the




Figure 13. Cumulative number of TDEs with $P(\text{TDE}) > 0.5$ as a fraction of all TDEs lower than some redshift z detected by Rubin or Roman. Out of all of the injected TDEs detected in our Rubin survey simulation, FLEET can recover $\sim 30\%$ of those at redshift $z < 0.5$ or $\sim 10\%$ of the ones at $z < 1.5$.

number of TDEs detected by Rubin does not extend past $z \sim 4$. We determine that $\sim 10^5$ TDEs could be detected by Rubin every year with at least four data points each or $\sim 10^4$ “well-observed” TDEs per year with at least 20 detections spanning a minimum 70 day baseline. This is comparable to the highest value of ~ 8000 TDEs yr^{-1} that Bricman & Gomboc (2020) estimated would be discovered by Rubin, which the authors calculated based on theoretical models of TDEs drawn from uniformly sampled impact factors and black hole masses using a fixed rate of 10^{-5} TDEs per galaxy per year. Alternatively, French & Zabludoff (2018) proposed targeting TDEs based on the colors of their hosts. The authors generated a sample of $\sim 68,000$ galaxies that are likely to host TDEs and predicted that 119–248 TDEs yr^{-1} can be detected with LSST with a purity of 73%.

Since Rubin includes g - and r -band observations, we can directly run FLEET on the light curves of the simulated TDEs that were detected by the survey. We run the late-time FLEET classifier on all of the TDEs detected by the Rubin simulation and show the fraction of TDEs with $P(\text{TDE}) > 0.5$ as a function of redshift in Figure 13. We find that with $P(\text{TDE}) > 0.5$, FLEET can recover about 10% of the TDEs detected by Rubin at redshift $z < 1.5$ or about 30% of TDEs at redshift $z < 0.5$. This translates to ~ 2000 TDEs that FLEET could uncover every year out to a redshift of $z = 0.5$ or ~ 3000 TDEs out to a redshift of $z = 1.5$. We show sample light curves of well-observed TDEs from the Rubin simulation located at redshifts of $z = 0.1$, 0.5 , and 1.5 in the bottom row of Figure 11. We note that these estimates come from running the current version of FLEET, trained on ZTF data, on Rubin data. Optimizing the algorithm with Rubin data is only expected to improve these estimates. Training the algorithm on the u band will be particularly useful, since TDEs are known to separate well from other transients in this band (van Velzen et al. 2020).

Despite the fact that Rubin is expected to find a larger total number of TDEs than Roman, the fraction of TDEs recovered by Rubin with 20 data points is an order of magnitude lower than the fraction of TDEs recovered with four data points. This means that Rubin is likely to find many TDEs that will go unclassified due to their poor light-curve coverage. This is not the case for the Roman HLTDS, for which the fraction of TDEs recovered with four data points is almost the same as the

fraction of TDEs recovered with 20 data points, as a consequence of the small footprint, cadence, and depth of the Roman HLTDS.

6. Conclusions

We have presented a new version of FLEET, a machine-learning classifier designed specifically to rapidly identify TDEs with a high purity and without the need for redshift information. We trained this classifier on a sample of 4779 spectroscopically classified transients, including 45 TDEs. We provide two classifiers, a rapid one trained on 20 days of photometry meant to be used for real-time classification and a late-time classifier trained on 40 days of data meant to be used for more robust estimates even if a transient has begun to fade. Both classifiers use light-curve and contextual host galaxy information to calculate the probability of a transient being a TDE, $P(\text{TDE})$. Our key findings are as follows.

1. The most important features for distinguishing TDEs from other transients are the normalized host separation R_n and the light-curve color during peak $(g - r)_p$.
2. We can recover TDEs with a purity of $\approx 30\%$ using the rapid classifier for events with $P(\text{TDE}) > 0.5$. This is a factor of ~ 60 improvement compared to random selection. The corresponding completeness for this threshold is $\approx 40\%$, or about five times better than using simple selection cuts.
3. We find that a peak purity of $\approx 50\%$ can be achieved with the rapid classifier for transients with $P(\text{TDE}) > 0.8$, corresponding to a completeness of $\approx 20\%$.
4. The late-time classifier trained on 40 days of data performs similarly, with a completeness of $\approx 30\%$ for transients with $P(\text{TDE}) > 0.8$ but a higher peak purity of $\approx 90\%$.

Additionally, we explore the application of FLEET to current and future time-domain surveys. We present a list of 39 TDE candidates with $P(\text{TDE}) > 0.5$ that were found by ZTF but remain currently unclassified. We inject TDE light curves into simulations of the Roman and Rubin time-domain surveys to estimate how many TDEs can be recovered given their current survey designs. We find that, using FLEET, we can recover ~ 30 TDEs yr^{-1} with the Roman HLTDS out to a redshift of $z \sim 7$ and ~ 3000 TDEs yr^{-1} with Rubin out to a redshift of $z \sim 3$. Tools like FLEET will be not only useful but necessary in the era of big data and large time-domain surveys. Given the successful performance of FLEET to date, the predictions for detecting TDEs with Rubin, and the lack of strong biases against physical parameters, we are confident we can begin running FLEET on Rubin data and discovering interesting transients as soon as the survey begins.

S.G. is supported by an STScI Postdoctoral Fellowship. The Berger Time Domain group at Harvard is supported in part by NSF and NASA grants, including support by the NSF under grant AST-2108531, as well as the NSF under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, <http://iafi.org/>). This project was supported in part by the Transients Science @ Space Telescope group. M.N. is supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement

No. 948381) and by a Fellowship from the Alan Turing Institute. Operation of the Pan-STARRS1 telescope is supported by the National Aeronautics and Space Administration under grant Nos. NNX12AR65G and NNX14AM74G issued through the NEO Observation Program. This work has made use of data from the European Space Agency (ESA) mission Gaia (<https://www.cosmos.esa.int/gaia>), processed by the Gaia Data Processing and Analysis Consortium (DPAC; <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the Gaia Multilateral Agreement. This research has made use of NASA’s Astrophysics Data System. This research has made use of the SIMBAD database, operated at CDS, Strasbourg, France. Based on observations obtained with MegaPrime/MegaCam, a joint project of CFHT and CEA/IRFU, 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. This work is based in part on data products produced at Terapix, available at the Canadian Astronomy Data Centre as part of the Canada–France–Hawaii Telescope Legacy Survey, a collaborative project of NRC and CNRS. This research has made use of the NASA/IPAC

Extragalactic Database, which is funded by the National Aeronautics and Space Administration and operated by the California Institute of Technology.

Facilities: ADS, TNS.

Software: Astropy (Astropy Collaboration et al. 2018), MOSFiT (Guillochon et al. 2017), extinction (Barbary 2016), Matplotlib (Hunter 2007), emcee (Foreman-Mackey et al. 2013), NumPy (van der Walt et al. 2011), scikit-learn (Pedregosa et al. 2011), SMOTE (Chawla et al. 2002), FLEET (Gomez et al. 2020b).

Appendix









In Table 2, we list the TDEs used to train the FLEET classifier. Only TDEs with at least two g -band and two r -band points are included. We exclude AT 2019eve from our classifier because even though it was originally classified as a TDE (van Velzen et al. 2021), the classification has since been retracted due to the uncertainty in its spectral features (Hammerstein et al. 2023). Similarly, we exclude AT 2018dyk, since it was reclassified from a TDE to a changing-look AGN by Frederick et al. (2019).

Table 2
TDEs Used for Training FLEET

Name	Redshift	Reference	Name	Redshift	Reference	Name	Redshift	Reference
AT 2018bsi	0.0510	[1, 2]	AT 2019meg	0.1520	[1, 3]	AT 2020vwl	0.0350	[7]
AT 2018hco	0.0900	[1, 3]	AT 2019mha	0.1480	[1, 3]	AT 2020wey	0.0273	[1]
AT 2018hyz	0.0458	[1, 3, 19, 20]	AT 2019qiz	0.0151	[1, 3, 23, 25]	AT 2020ysg	0.2770	[1]
AT 2018iih	0.2120	[1, 3, 26]	AT 2019teq	0.0878	[1]	AT 2020zso	0.0610	[1, 24]
AT 2018jbv	0.3400	[1]	AT 2019vcv	0.0890	[1]	AT 2021ack	0.1330	[8]
AT 2018lna	0.0910	[1, 3]	AT 2020ddv	0.1600	[1]	AT 2021axu	0.1900	[9]
AT 2018lni	0.1380	[1]	AT 2020mbq	0.0930	[1]	AT 2021ehb	0.0170	[10]
AT 2018zr	0.0710	[1, 3, 26]	AT 2020mot	0.0700	[1]	AT 2021gje	0.3580	[11]
AT 2019azh	0.0223	[1, 3, 21]	AT 2020neh	0.0620	[5]	AT 2021jjm	0.1530	[12]
AT 2019bbf	0.1206	[1, 3]	AT 2020nov	0.0840	[6]	AT 2021jsg	0.1260	[13]
AT 2019cho	0.1930	[1, 3]	AT 2020ocn	0.0700	[1]	AT 2021mhg	0.0730	[14]
AT 2019dsg	0.0512	[1, 3, 22]	AT 2020opy	0.1590	[1]	AT 2021nwa	0.0470	[15]
AT 2019ehz	0.0740	[1, 3]	AT 2020pj	0.0680	[1]	AT 2021sdu	0.0590	[16]
AT 2019gte	0.0860	[4]	AT 2020qhs	0.3450	[1]	AT 2021uqv	0.1060	[17]
AT 2019lwu	0.1170	[1, 3]	AT 2020riz	0.4350	[1]	AT 2021yte	0.0530	[18]

Note. All TDEs used to train our classifiers, listed alphabetically. 1: Hammerstein et al. (2023); 2: van Velzen et al. (2021); 3: Nicholl et al. (2022); 4: Swann et al. (2019); 5: Dahiwalé & Fremling (2020a); 6: Dahiwalé & Fremling (2020b); 7: Hammerstein et al. (2021d); 8: Hammerstein et al. (2021a); 9: Hammerstein et al. (2021e); 10: Yao (2021a); 11: Hammerstein et al. (2021b); 12: Yao et al. (2021d); 13: Yao et al. (2021c); 14: Chu et al. (2021a); 15: Yao et al. (2021b); 16: Chu et al. (2021b); 17: Yao (2021b); 18: Yao et al. (2021a); 19: Short et al. (2020); 20: Gomez et al. (2020c); 21: Liu et al. (2022); 22: Cannizzaro et al. (2021); 23: Hung et al. (2021); 24: Wevers et al. (2022); 25: Nicholl et al. (2020); 26: Cao et al. (2022).

ORCID iDs

Sebastian Gomez  <https://orcid.org/0000-0001-6395-6702>
 V. Ashley Villar  <https://orcid.org/0000-0002-5814-4061>
 Edo Berger  <https://orcid.org/0000-0002-9392-9681>
 Suvi Gezari  <https://orcid.org/0000-0003-3703-5154>
 Sjoert van Velzen  <https://orcid.org/0000-0002-3859-8074>
 Matt Nicholl  <https://orcid.org/0000-0002-2555-3192>
 Peter K. Blanchard  <https://orcid.org/0000-0003-0526-2248>
 Kate. D. Alexander  <https://orcid.org/0000-0002-8297-2473>

References

- Ahumada, R., Allende Prieto, C., Almeida, A., et al. 2020, *ApJS*, **249**, 3
 Alam, S., Albareti, F. D., Allende Prieto, C., et al. 2015, *ApJS*, **219**, 12
 Alexander, K. D., Berger, E., Guillochon, J., Zauderer, B. A., & Williams, P. K. G. 2016, *ApJL*, **819**, L25
 Alexander, K. D., van Velzen, S., Horesh, A., & Zauderer, B. A. 2020, *SSRv*, **216**, 81
 Arcavi, I., Gal-Yam, A., Sullivan, M., et al. 2014, *ApJ*, **793**, 38
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
 Auchettl, K., Guillochon, J., & Ramirez-Ruiz, E. 2017, *ApJ*, **838**, 149
 Barbary, K. 2016, Extinction v0.3.0, Zenodo, doi:10.5281/zenodo.804967
 Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, *PASP*, **131**, 018002
 Berger, E. 2010, *ApJ*, **722**, 1946
 Bertin, E., & Arnouts, S. 1996, *A&AS*, **117**, 393
 Blagorodnova, N., Cenko, S. B., Kulkarni, S. R., et al. 2019, *ApJ*, **873**, 92
 Bloom, J. S., Kulkarni, S. R., & Djorgovski, S. G. 2002, *AJ*, **123**, 1111
 Boone, K. 2019, *AJ*, **158**, 257
 Bower, G. C., Metzger, B. D., Cenko, S. B., Silverman, J. M., & Bloom, J. S. 2013, *ApJ*, **763**, 84
 Breiman, L. 2001, *Mach. Learn.*, **45**, 5
 Bricman, K., & Gomboc, A. 2020, *ApJ*, **890**, 73
 Cannizzaro, G., Wevers, T., Jonker, P. G., et al. 2021, *MNRAS*, **504**, 792
 Cao, X.-M., Wang, S.-Q., Gan, W.-P., & Li, J.-Y. 2022, arXiv:2210.04115
 Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, *ApJ*, **345**, 245
 Cendes, Y., Berger, E., Alexander, K., et al. 2022, *ApJ*, **938**, 28
 Chambers, K. & Pan-STARRS Team 2018, AAS Meeting Abstracts, **231**, 102.01
 Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. 2002, *J. Artif. Intell. Res.*, **16**, 321
 Chornock, R., Berger, E., Gezari, S., et al. 2014, *ApJ*, **780**, 44
 Chu, M., Dahiwalé, A., & Fremling, C. 2021a, TNSCR, **2021-2672**, 1
 Chu, M., Dahiwalé, A., & Fremling, C. 2021b, TNSCR, **2021-2712**, 1
 Dahiwalé, A., & Fremling, C. 2020a, TNSCR, **2020-2126**, 1
 Dahiwalé, A., & Fremling, C. 2020b, TNSCR, **2020-3800**, 1
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
 Förster, F., Cabrera-Vives, G., Castillo-Navarrete, E., et al. 2021, *AJ*, **161**, 242
 Förster, F., Muñoz Arancibia, A. M., Reyes-Jainaga, I., et al. 2022, *AJ*, **164**, 195
 Frederick, S., Gezari, S., Graham, M. J., et al. 2019, *ApJ*, **883**, 31
 Fremling, C., Miller, A. A., Sharma, Y., et al. 2020, *ApJ*, **895**, 32
 French, K. D., Arcavi, I., & Zabludoff, A. 2016, *ApJL*, **818**, L21
 French, K. D., & Zabludoff, A. I. 2018, *ApJ*, **868**, 99
 Gagliano, A., Narayan, G., Engel, A., Carrasco Kind, M., & LSST Dark Energy Science Collaboration 2021, *ApJ*, **908**, 170
 Gezari, S. 2021, *ARA&A*, **59**, 21
 Gezari, S., Cenko, S. B., & Arcavi, I. 2017, *ApJL*, **851**, L47
 Gezari, S., Chornock, R., Rest, A., et al. 2012, *Natur*, **485**, 217
 Gezari, S., Heckman, T., Cenko, S. B., et al. 2009, *ApJ*, **698**, 1367
 Gomez, S., Berger, E., Blanchard, P. K., et al. 2020a, *ApJ*, **904**, 74
 Gomez, S., Berger, E., Blanchard, P. K., et al. 2020b, FLEET Finding Luminous and Exotic Extragalactic Transients, v1.0.0, Zenodo, doi:10.5281/zenodo.4013965
 Gomez, S., Nicholl, M., Short, P., et al. 2020c, *MNRAS*, **497**, 1925
 Gomez, S., Berger, E., Blanchard, P. K., et al. 2022, arXiv:2210.10811
 Graur, O., French, K. D., Zahid, H. J., et al. 2018, *ApJ*, **853**, 39
 Guillochon, J., Nicholl, M., Villar, V. A., et al. 2017, MOSFIT: Modular Open-Source Fitter for Transients, Astrophysics Source Code Library, ascl:1710.006
 Guillochon, J., Nicholl, M., Villar, V. A., et al. 2018, *ApJS*, **236**, 6
 Guillochon, J., Ramirez-Ruiz, E., Rosswog, S., & Kasen, D. 2009, *ApJ*, **705**, 844
 Hammerstein, E., Gezari, S., Velzen, S. V., et al. 2021a, TNSCR, **2021-732**, 1
 Hammerstein, E., Yao, Y., Gezari, S., et al. 2021b, TNSCR, **2021-1723**, 1
 Hammerstein, E., Gezari, S., van Velzen, S., et al. 2021c, *ApJL*, **908**, L20
 Hammerstein, E., Gezari, S., Velzen, S. V., et al. 2021d, TNSCR, **2021-159**, 1
 Hammerstein, E., Gezari, S., Velzen, S. V., et al. 2021e, TNSCR, **2021-955**, 1
 Hammerstein, E., van Velzen, S., Gezari, S., et al. 2023, *ApJ*, **942**, 9
 Hills, J. G. 1975, *Natur*, **254**, 295
 Holoien, T. W. S., Kochanek, C. S., Prieto, J. L., et al. 2016, *MNRAS*, **455**, 2918
 Hosseinzadeh, G., Auphin, F., Villar, V. A., et al. 2020, *ApJ*, **905**, 93
 Hudelot, P., Cuillandre, J. C., Withington, K., et al. 2012, *yCat*, **2317**, 0
 Hung, T., Foley, R. J., Veilleux, S., et al. 2021, *ApJ*, **917**, 9
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111
 Kessler, R., Narayan, G., Avelino, A., et al. 2019, *PASP*, **131**, 094501
 Kochanek, C. S. 2016, *MNRAS*, **461**, 371
 Leloudas, G., Dai, L., Arcavi, I., et al. 2019, *ApJ*, **887**, 218
 Liu, X.-L., Dou, L.-M., Chen, J.-H., & Shen, R.-F. 2022, *ApJ*, **925**, 67
 Lu, W., & Bonnerot, C. 2020, *MNRAS*, **492**, 686
 LSST Science Collaboration, Marshall, P., Anguita, T., et al. 2017, arXiv:1708.04058
 Mockler, B., Guillochon, J., & Ramirez-Ruiz, E. 2019, *ApJ*, **872**, 151
 Muthukrishna, D., Narayan, G., Mandel, K. S., Biswas, R., & Hložek, R. 2019, *PASP*, **131**, 118002
 Nicholl, M., Blanchard, P. K., Berger, E., et al. 2019, *MNRAS*, **488**, 1878
 Nicholl, M., Lanning, D., Ramsden, P., et al. 2022, *MNRAS*, **515**, 5604
 Nicholl, M., Wevers, T., Oates, S. R., et al. 2020, *MNRAS*, **499**, 482
 Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *J. Mach. Learn. Res.*, **12**, 2825
 Rees, M. J. 1988, *Natur*, **333**, 523
 Rose, B. M., Baltay, C., Hounsell, R., et al. 2021, arXiv:2111.03081
 Sánchez-Sáez, P., Reyes, I., Valenzuela, C., et al. 2021, *AJ*, **161**, 141
 Schlafly, E. F., & Finkbeiner, D. P. 2011, *ApJ*, **737**, 103
 Scholtz, F. W., & Stephens, M. A. 1987, *J. Am. Stat. Assoc.*, **82**, 918
 Short, P., Nicholl, M., Lawrence, A., et al. 2020, *MNRAS*, **498**, 4119
 Spergel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv:1503.03757
 Swann, E., Frohmaier, C., Nicholl, M., Short, P., & Yaron, O. 2019, TNSCR, **2019-975**, 1
 van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
 van Velzen, S. 2018, *ApJ*, **852**, 72
 van Velzen, S., Anderson, G. E., Stone, N. C., et al. 2016, *Sci*, **351**, 62
 van Velzen, S., Frail, D. A., Körding, E., & Falcke, H. 2013, *A&A*, **552**, A5
 van Velzen, S., Gezari, S., Hammerstein, E., et al. 2021, *ApJ*, **908**, 4
 van Velzen, S., Holoien, T. W. S., Onori, F., Hung, T., & Arcavi, I. 2020, *SSRv*, **216**, 124
 Villar, V. A., Berger, E., Miller, G., et al. 2019, *ApJ*, **884**, 83
 Villar, V. A., Hosseinzadeh, G., Berger, E., et al. 2020, *ApJ*, **905**, 94
 Villar, V. A., Nicholl, M., & Berger, E. 2018, *ApJ*, **869**, 166
 Wevers, T., Nicholl, M., Guolo, M., et al. 2022, *A&A*, **666**, A6
 Wevers, T., Pasham, D. R., van Velzen, S., et al. 2019, *MNRAS*, **488**, 4816
 Yao, Y. 2021a, TNSCR, **2021-2295**, 1
 Yao, Y. 2021b, TNSCR, **2021-3411**, 1
 Yao, Y., Chu, M., Das, K. K., et al. 2021a, TNSCR, **2021-3611**, 1
 Yao, Y., Gezari, S., Velzen, S. V., Hammerstein, E., & Somalwar, J. 2021b, TNSCR, **2021-2155**, 1
 Yao, Y., Sharma, Y., Gezari, S., et al. 2021c, TNSCR, **2021-1221**, 1
 Yao, Y., Velzen, S. V., Perley, D., et al. 2021d, TNSCR, **2021-1632**, 1
 Zauderer, B. A., Berger, E., Soderberg, A. M., et al. 2011, *Natur*, **476**, 425