
BTSbot: A Multi-input Convolutional Neural Network to Automate and Expedite Bright Transient Identification for the Zwicky Transient Facility

Nabeel Rehemtulla^{1,2} Adam A. Miller^{1,2} Michael W. Coughlin³ Theophile Jegou du Laz⁴ on behalf of ZTF

Abstract

The Bright Transient Survey (BTS) relies on visual inspection (“scanning”) to select sources for accomplishing its mission of spectroscopically classifying all bright extragalactic transients found by the Zwicky Transient Facility (ZTF). We present **BTSbot**, a multi-input convolutional neural network, which provides a bright transient score to individual ZTF detections using their image data and 14 extracted features. **BTSbot** eliminates the need for scanning by automatically identifying and requesting follow-up observations of new bright ($m < 18.5$ mag) transient candidates. **BTSbot** outperforms BTS scanners in terms of completeness (99% vs. 95%) and identification speed (on average, 7.4 hours quicker).

1. Introduction

Large, wide-field surveys have recently revolutionized time-domain astronomy by repeatedly imaging the entire night sky. These surveys produce staggering amounts of data every night, an influx that demands the adoption of machine learning (ML) techniques. ML models have been applied to a variety of tasks in astronomy including real/bogus classification (e.g., Bloom et al., 2012; Brink et al., 2013), photometric transient classification (e.g., Villar et al., 2020), photometric redshift estimation (e.g., Carrasco Kind & Brunner, 2013), and many others. For the most part, ML models in astronomy perform their task using a small number of extracted numeric features. Limiting these models to

¹Department of Physics and Astronomy, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA ²Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), 1800 Sherman Ave., Evanston, IL 60201, USA ³School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA ⁴Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA. Correspondence to: Nabeel Rehemtulla <nabeelr@u.northwestern.edu>.

ICML 2023 Workshop on Machine Learning for Astrophysics, Honolulu, Hawaii, USA. PMLR 202, 2023. Copyright 2023 by the author(s).

extracted features ignores potentially valuable information present in the associated images. A comparatively small number of convolutional neural networks (CNNs) have been built that make use of the information embedded in astronomical images, and they have generally had great success (e.g., Dieleman et al., 2015; Lanusse et al., 2018; Duev et al., 2019). CNNs are particularly well suited to astronomy because they can capture properties, like galaxy morphology, which often remain largely obscured to other image processing techniques. Only a very small subset of these CNNs are *multi-input*, meaning they take in images and input of another type, like extracted features (Carrasco-Davis et al., 2021). Multi-input CNNs (MI-CNNs) have more and varied information to draw from compared to analogous single-input models. Here, we present a new MI-CNN for bright transient identification.

The Bright Transient Survey (BTS; Fremling et al., 2020; Perley et al., 2020) aims to classify all extragalactic transients with $m < 18.5$ mag in g or r band from the public alert stream of the Zwicky Transient Facility (ZTF; Bellm et al., 2019). Every night, experts inspect candidate BTS sources, a process called “scanning,” and select the bright extragalactic transients¹ for spectroscopic observation and classification. There are typically ~ 50 BTS candidates per night, of which ~ 10 are real bright supernovae (SNe), with the others being mostly dim SNe, active galactic nuclei (AGN), and cataclysmic variables (CVs). Since its origin, BTS has maintained near-perfect completeness of relevant sources and has rapidly and publicly released their classifications: a monumental service to the community. The BTS sample enables an enormous amount of science including some of the largest SN population studies ever conducted. The MI-CNN introduced here, dubbed **BTSbot**, is built to automate scanning for BTS by performing binary classification (bright SN / not bright SN) on ZTF alert packets.

2. Model scope, architecture, and training

Our MI-CNN, **BTSbot**, automates BTS scanning by assigning each individual ZTF alert packet a bright transient

¹We use “supernova” interchangeably with “extragalactic transient.” These are not equivalent, but it simplifies the prose.

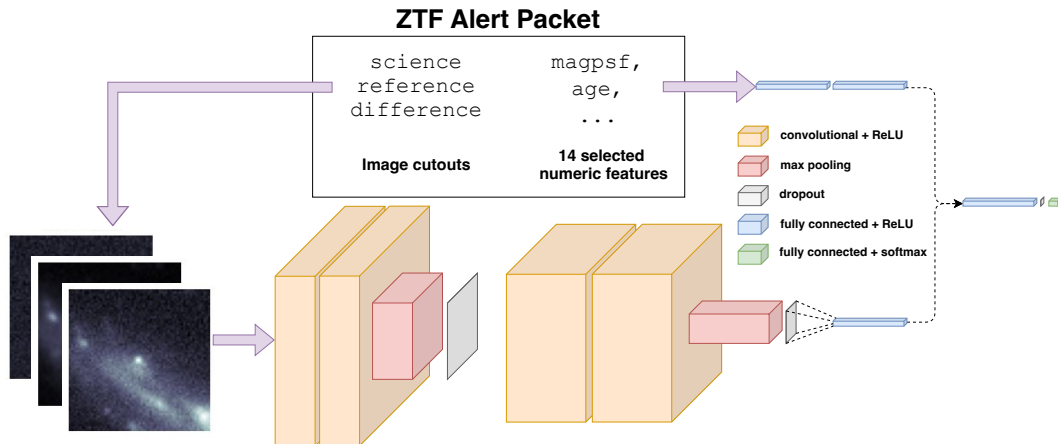


Figure 1. Diagram of BTSbot, a multi-input convolutional neural network, showing the layers that the input images and extracted features pass through before being rejoined and producing the score for bright extragalactic transient binary classification.

score. Alerts from real bright SNe will have high scores, allowing us to automatically select them from the pool of BTS candidates for follow-up and classification. Further, our model identifies sources and requests follow-up before a human can, helping to maintain the survey’s near-perfect completeness by obtaining a spectrum before the transient fades beyond detection.

Figure 1 shows how input is fed into BTSbot and how the information from the extracted features and images is combined to produce the output. BTSbot contains three main components: (1) The convolutional branch processes the science, reference, and difference images as a three-channel image through a VGG-like architecture (Simonyan & Zisserman, 2014) followed by flattening. (2) The fully-connected branch processes the extracted features through two dense layers. (3) The combined section concatenates the output of the two branches and passes it through two more fully-connected layers; the second of which produces the final output using a softmax activation function. The output is a unit-interval score where higher scores represent increased confidence that the source in the input alert packet is, or will become, a bright extragalactic transient.

The choice of a MI-CNN is motivated by the fact that the images and the extracted features provide different, valuable information for performing our task. For example, the features `distpsnr1` and `sgscore1` respectively represent the angular distance to and the star-galaxy score (Tachibana & Miller, 2018) of the Pan-STARRS1 (PS1; Kaiser et al., 2002) cataloged source nearest to this ZTF source. While the new SNe are not present in PS1, their host galaxies often are. Thus, alerts from bright SNe tend to have moderate `distpsnr1` and small `sgscore1` values, indicating a galaxy projected nearby to the source. In contrast, AGN and CVs will typically be cataloged in PS1 and thus have

`distpsnr1` very near to zero. The images also provide important information following a similar heuristic. Bright SNe tend to be associated with prominent off-center extended sources (their host galaxies); faint SNe tend to have less prominent host galaxies because they tend to be farther away; AGN will appear as exactly centered extended sources; CVs will often appear surrounded by many bright point sources because they tend to occur near dense star fields. A MI-CNN is able to pool information from all input types and consider them together when making a prediction. Where a single-input CNN might fail due to a lack of discriminating information, a MI-CNN leverages additional and distinct information when making its prediction. For example, a very faint alert that shows next to no features in its images can be identified as an AGN by a MI-CNN because AGN often have a large number of previous detections at their location (represented by `ndethist`), information a single-input CNN could not have.

Given the scope and architecture of BTSbot we encounter a number of challenges. First, we are requiring the model to learn multiple non-trivial separations. BTSbot must learn to separate SNe from other sources without using distance information because it is not known *a priori*. It must also learn to identify bright SNe from a single alert irrespective of the SN’s current phase. Early in its rise or late in its fade, a bright SN can appear very similar to a near-peak dim SN. This is related to the second complication, which is that BTSbot uses no time-series information. Although light curves contain crucial information for evaluating a source as a bright SN or not, we make the choice to omit time-series information, in part, because it would introduce complexity and, likely, noise. There is no established method for representing partial light curves of the wide variety our model encounters in a way that is fit for input into a neural network. There has been a great deal of work to accomplish




Figure 2. Stacked bar chart showing the distribution of classification outcomes as a function of the input alerts’ PSF magnitude. The highlighted region to the left is most important because it marks the alerts which BTS scanners use to trigger spectroscopic observations. Performance of alerts with $m < 18.5$ mag is excellent and misclassifications are mostly in the dimmest bin.

this for SNe alone (Villar et al., 2020), but these methods are not applicable to all the types of sources that our model encounters. Further, omitting time-series information offers BTSbot an advantage over light curve-based models. With its alert-based architecture, BTSbot can identify a bright SN from its very first detection, rather than preferring multiple detections to begin constraining a light curve. We expect that a similar model with light curve information would identify bright SNe as or less quickly than BTSbot does, potentially hurting BTS’s completeness. Section 4.1 illustrates how prompt and automatic identification of SNe with BTSbot can aid in uncovering poorly understood SN physics. To our knowledge, no other CNNs exist that simultaneously filter out non-SNe and predict a SN’s future behavior given a single snapshot in time.

2.1. Training

Solving this complex classification task requires a significant training set. Having run since 2018, BTS has now amassed the largest set of public SNe classifications ever. The size of this labeled data set enables the construction of BTSbot. Bright SNe classified by BTS populate the positive class, while AGN, CVs, and dim SNe rejected by BTS constitute the negative class. In total, we have 561,000 alerts ($\sim 44\%$ of which belong to the positive class) from 14,348 sources ($\sim 33\%$ of which belong to the positive class). This difference in class distribution between alerts and sources stems from some sources having many more alerts than others: AGN can have thousands, bright SNe typically have

dozens, and dim SNe can have as little as a few. Training on every alert would result in some types of sources being over-represented. To remedy this we define a hyperparameter called N_{\max} , the maximum number of alerts included per source. We find $N_{\max} = 60$ to be optimal; it balances between thinning extra alerts from sources like AGN and maximizing the training set size. We also weight contributions to the loss function by the relative size of each class to mitigate the effects of class imbalance.

BTSbot is trained with mostly standard hyperparameters. We adopt the Adam optimizer (Kingma & Ba, 2014) and the binary cross-entropy loss function. We perform a number of Bayesian hyperparameter searches with the Weights and Biases platform (Biewald, 2020) to optimize our choice of batch size, Adam parameters, N_{\max} , which extracted features to include, and more. We also employ data augmentation to the image cutouts like random rotations of 0° , 90° , 180° , and 270° and random horizontal and vertical flipping. These help prevent overfitting and ensure that BTSbot is invariant to these transformations.

We employ a standard 81/9/10% train/validation/test split. We also prevent sources from having alerts in multiple splits. This prevents different splits from containing extremely similar data, which would introduce bias.

3. Model performance

The trained model is run on the full validation split (i.e. no N_{\max} thinning) to create unbiased performance diagnostics.

Figure 2 illustrates classification outcomes as a function of PSF magnitude. The highlight marks $m < 18.5$ mag, the BTS magnitude threshold. The alerts within the highlight are especially important to classify correctly because they influence whether or not a spectrum of the source is to be collected. In that region, misclassifications are very limited and mostly in the dimmest bin. For $m > 18.5$ mag, there are more misclassifications and many more alerts overall.

Ultimately, alert-based metrics are not perfectly representative of the model’s real-world performance; source-based classification is more relevant. The chosen metrics must consider that the model has multiple chances to correctly (or incorrectly) classify a source. To this end, we define a “policy” that maps the real-time history of a source’s scores to a source-based classification. For simplicity, we put aside policy optimization for now and only consider one naive policy with the impression that we could improve performance with a more sophisticated choice of policy.

Our naive policy is named `gt1`; it classifies a source as a bright SN once it has *one* alert with score ≥ 0.5 . The left panel of Figure 3 shows the purity and completeness of the sample produced when following this policy. We




Figure 3. *BTSbot* performance when classifying sources as bright supernovae given at least one alert with high score. Left panel: completeness (navy) and purity (orange) as a function of peak magnitude and respective averages (dashed lines). *BTSbot* produces a sample that is 99.1% complete and 90.7% pure overall. The bright limit is noisy due to the small number of sources in those bins. Right panel: distribution of delay between *BTSbot* and BTS scanner identification times with median at -0.31 days, indicating that, on average, *BTSbot* identifies sources about 7 hours quicker than BTS scanners.

observe that *gt1* yields a sample with 99.1% completeness and 90.7% purity overall. BTS scanners produce a sample that is roughly 95% in both completeness and purity (Perley et al., 2020). In tests with other more conservative policies like *gt3* (defined analogously to *gt1*), we observe that purity is increased at the cost of completeness. We favor *gt1* because it maximizes completeness, BTS’s highest priority. The right panel of Figure 3 compares the time between when *BTSbot* and the BTS scanners identified some source. Here, negative numbers represent that *BTSbot* identified the bright SN before the scanners and positive numbers represent the opposite. The median is -0.31 days, indicating that *BTSbot* expedites the identification of bright SNe by, on average, 7.44 hours. Expediting identification by several hours will frequently mean observing the source an entire night earlier, thus yielding an improvement of ~ 24 hours in practice. With *gt1*, *BTSbot* outperforms human scanners in completeness and speed, while making only a small compromise with slightly lower purity.

4. Real-time, real-world usage

BTSbot has been integrated into Fritz, the first-party ZTF alert-broker and a SkyPortal instance (Coughlin et al., 2023), and is currently posting scores to all new ZTF alert packets. Alert packets are created and augmented with a bright transient scores just minutes after the observations are taken with ZTF. During the night, a tool we call *autoscan* checks for new BTS candidates every half hour and immediately saves those that pass *gt1*. *autoscan* will soon also simultaneously request follow-up for passing sources. With this, we can monitor the model’s real-time, real-world performance and note any frequent misclassifications. We have been awarded a large observing program for the 2023B

semester with the SEDMv2 spectrograph on the Kitt Peak 2.1 m telescope. With this time, we will build a spectroscopic sample that represents the model’s unbiased performance. The resulting performance metrics will be invaluable when considering adaptations of this MI-CNN, e.g., to specific SN sub-types or exotic transients like kilonovae.

4.1. SN 2023ixf & Rapid follow-up with *BTSbot*

We take SN 2023ixf, the recent Type-II supernova in M101, as an example illustrative of the additional discovery potential of an alert-based model like *BTSbot*. SN 2023ixf was reported to TNS by Koichi Itagaki at 14:42 PDT on May 19th (Itagaki, 2023). The earliest published spectrum was collected by Daniel Perley less than an hour later with the SPRAT spectrograph (Perley et al., 2023). About 14 hours before the first TNS report, SN 2023ixf was detected by ZTF, and, just minutes later, this alert packet was assigned a bright transient score of 0.840 by *BTSbot*. If it was in-place at the time, *autoscan* would have identified this new source passing *gt1* at about 01:00 PDT and requested a spectrum from one of the numerous robotic spectrographs associated with ZTF and BTS, e.g., SEDM. At this point, there is about half of the night remaining at SEDM’s location, plenty of time for a spectrum to be collected. Even if the spectrum is collected at the end of the night ($\sim 05:00$ PDT), this represents a ~ 10 hour speed-up over the otherwise earliest spectrum. In this example, *BTSbot* and *autoscan* probe the early, rapidly-evolving explosion physics mostly unavailable to traditional triggering methods. This would not be possible with light curve-based models, which typically require using the source’s evolution over multiple days to identify it as a transient. This extremely rapid follow-up is enabled by the alert-based architecture of *BTSbot*.

Software and Data

Code relating to BTSbot and autoscan are made public at <https://github.com/nabeelre/BTSbot>. In their development, we use Astropy (Astropy Collaboration et al., 2013; 2018), cron (Reznick, 1993), Keras (Chollet et al., 2015), Matplotlib (Hunter, 2007), NumPy (Harris et al., 2020), pandas (pandas development team, 2020; Wes McKinney, 2010), penquins (Duvet et al., 2021), scikit-learn (Pedregosa et al., 2011), and Tensorflow (Abadi et al., 2015).

Acknowledgements

The material contained in this document is based upon work supported by a National Aeronautics and Space Administration (NASA) grant or cooperative agreement. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author and do not necessarily reflect the views of NASA. This work was supported through a NASA grant awarded to the Illinois/NASA Space Grant Consortium.

M. W. Coughlin acknowledges support from the National Science Foundation with grant numbers PHY-2010970 and OAC-2117997.

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. Major funding has been provided by the U.S National Science Foundation under Grant No. AST-1440341 and by the ZTF partner institutions: the California Institute of Technology, the Oskar Klein Centre, the Weizmann Institute of Science, the University of Maryland, the University of Washington, Deutsches Elektronen-Synchrotron, the University of Wisconsin-Milwaukee, and the TANGO Program of the University System of Taiwan.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., Greenfield, P., Droettboom, M., Bray, E., Aldcroft, T., Davis, M., Ginsburg, A., Price-Whelan, A. M., Kerzendorf, W. E., Conley, A., Crighton, N., Barbary, K., Muna, D., Ferguson, H., Grollier, F., Parikh, M. M., Nair, P. H., Unther, H. M., Deil, C., Woillez, J., Conseil, S., Kramer, R., Turner, J. E. H., Singer, L., Fox, R., Weaver, B. A., Zabalza, V., Edwards, Z. I., Azalee Bostroem, K., Burke, D. J., Casey, A. R., Crawford, S. M., Dencheva, N., Ely, J., Jenness, T., Labrie, K., Lim, P. L., Pierfederici, F., Pontzen, A., Ptak, A., Refsdal, B., Servillat, M., and Streicher, O. Astropy: A community Python package for astronomy. , 558:A33, October 2013. doi: 10.1051/0004-6361/201322068.
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., Lim, P. L., Crawford, S. M., Conseil, S., Shupe, D. L., Craig, M. W., Dencheva, N., Ginsburg, A., VanderPlas, J. T., Bradley, L. D., Pérez-Suárez, D., de Val-Borro, M., Aldcroft, T. L., Cruz, K. L., Robitaille, T. P., Tollerud, E. J., Ardelean, C., Babej, T., Bach, Y. P., Bachetti, M., Bakanov, A. V., Bamford, S. P., Barentsen, G., Barmby, P., Baumbach, A., Berry, K. L., Biscani, F., Boquien, M., Bostroem, K. A., Bouma, L. G., Brammer, G. B., Bray, E. M., Breytenbach, H., Buddelmeijer, H., Burke, D. J., Calderone, G., Cano Rodríguez, J. L., Cara, M., Cardoso, J. V. M., Cheedella, S., Copin, Y., Corrales, L., Crichton, D., D’Avella, D., Deil, C., Depagne, É., Dietrich, J. P., Donath, A., Droettboom, M., Earl, N., Erben, T., Fabbro, S., Ferreira, L. A., Finethy, T., Fox, R. T., Garrison, L. H., Gibbons, S. L. J., Goldstein, D. A., Gommers, R., Greco, J. P., Greenfield, P., Groener, A. M., Grollier, F., Hagen, A., Hirst, P., Homeier, D., Horton, A. J., Hosseinzadeh, G., Hu, L., Hunkeler, J. S., Ivezić, Ž., Jain, A., Jenness, T., Kanarek, G., Kendrew, S., Kern, N. S., Kerzendorf, W. E., Khvalko, A., King, J., Kirkby, D., Kulkarni, A. M., Kumar, A., Lee, A., Lenz, D., Littlefair, S. P., Ma, Z., Macleod, D. M., Mastrogiuseppe, M., McCully, C., Montagnac, S., Morris, B. M., Mueller, M., Mumford, S. J., Muna, D., Murphy, N. A., Nelson, S., Nguyen, G. H., Ninan, J. P., Nöthe, M., Ogaz, S., Oh, S., Parejko, J. K., Parley, N., Pascual, S., Patil, R., Patil, A. A., Plunkett, A. L., Prochaska, J. X., Rastogi, T., Reddy Janga, V., Sabater, J., Sakurikar, P., Seifert, M., Sherbert, L. E., Sherwood-Taylor, H., Shih, A. Y., Sick, J., Silbiger, M. T., Singanamalla, S., Singer, L. P., Sladen, P. H., Sooley, K. A., Sornarajah, S., Streicher, O., Teuben, P., Thomas, S. W., Tremblay, G. R., Turner, J. E. H., Terrón, V., van Kerkwijk, M. H., de la Vega, A., Watkins, L. L., Weaver, B. A., Whitmore, J. B., Woillez, J., Zabalza, V., and Astropy Contributors. The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package. , 156(3):123, September 2018. doi: 10.3847/1538-3881/aabc4f.
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., Dekany, R., Smith, R. M., Riddle, R., Masci, F. J., Helou, G., Prince, T. A., Adams, S. M., Barbarino, C., Barlow, T., Bauer,

- J., Beck, R., Belicki, J., Biswas, R., Blagorodnova, N., Bodewits, D., Bolin, B., Brinell, V., Brooke, T., Bue, B., Bulla, M., Burruss, R., Cenko, S. B., Chang, C.-K., Connolly, A., Coughlin, M., Cromer, J., Cunningham, V., De, K., Delacroix, A., Desai, V., Duev, D. A., Eadie, G., Farnham, T. L., Feeney, M., Feindt, U., Flynn, D., Franckowiak, A., Frederick, S., Fremling, C., Gal-Yam, A., Gezari, S., Giomi, M., Goldstein, D. A., Golkhou, V. Z., Goobar, A., Groom, S., Hacopians, E., Hale, D., Henning, J., Ho, A. Y. Q., Hover, D., Howell, J., Hung, T., Huppenkothen, D., Imel, D., Ip, W.-H., Ivezić, Ž., Jackson, E., Jones, L., Juric, M., Kasliwal, M. M., Kaspi, S., Kaye, S., Kelley, M. S. P., Kowalski, M., Kramer, E., Kupfer, T., Landry, W., Laher, R. R., Lee, C.-D., Lin, H. W., Lin, Z.-Y., Lunnan, R., Giomi, M., Mahabal, A., Mao, P., Miller, A. A., Monkewitz, S., Murphy, P., Ngeow, C.-C., Nordin, J., Nugent, P., Ofek, E., Patterson, M. T., Penprase, B., Porter, M., Rauch, L., Rebbapragada, U., Reiley, D., Rigault, M., Rodriguez, H., van Roesstel, J., Rusholme, B., van Santen, J., Schulze, S., Shupe, D. L., Singer, L. P., Soumagnac, M. T., Stein, R., Surace, J., Sollerman, J., Szkody, P., Taddia, F., Terek, S., Van Sistine, A., van Velzen, S., Vestrand, W. T., Walters, R., Ward, C., Ye, Q.-Z., Yu, P.-C., Yan, L., and Zolkower, J. The Zwicky Transient Facility: System Overview, Performance, and First Results. , 131(995):018002, January 2019. doi: 10.1088/1538-3873/aaecbe.
- Biewald, L. Experiment tracking with weights and biases, 2020. URL <https://www.wandb.com/>. Software available from wandb.com.
- Bloom, J. S., Richards, J. W., Nugent, P. E., Quimby, R. M., Kasliwal, M. M., Starr, D. L., Poznanski, D., Ofek, E. O., Cenko, S. B., Butler, N. R., Kulkarni, S. R., Gal-Yam, A., and Law, N. Automating Discovery and Classification of Transients and Variable Stars in the Synoptic Survey Era. , 124(921):1175, November 2012. doi: 10.1086/668468.
- Brink, H., Richards, J. W., Poznanski, D., Bloom, J. S., Rice, J., Negahban, S., and Wainwright, M. Using machine learning for discovery in synoptic survey imaging data. , 435(2):1047–1060, October 2013. doi: 10.1093/mnras/stt1306.
- Carrasco-Davis, R., Reyes, E., Valenzuela, C., Förster, F., Estévez, P. A., Pignata, G., Bauer, F. E., Reyes, I., Sánchez-Sáez, P., Cabrera-Vives, G., Eyheramendy, S., Catelan, M., Arredondo, J., Castillo-Navarrete, E., Rodríguez-Mancini, D., Ruz-Mieres, D., Moya, A., Sabatini-Gacitúa, L., Sepúlveda-Cobo, C., Mahabal, A. A., Silva-Farfán, J., Camacho-Iñiguez, E., and Galbany, L. Alert Classification for the ALeRCE Broker System: The Real-time Stamp Classifier. , 162(6):231, December 2021. doi: 10.3847/1538-3881/ac0ef1.
- Carrasco Kind, M. and Brunner, R. J. TPZ: photometric redshift PDFs and ancillary information by using prediction trees and random forests. , 432(2):1483–1501, June 2013. doi: 10.1093/mnras/stt574.
- Chollet, F. et al. Keras. <https://keras.io>, 2015.
- Coughlin, M. W., Bloom, J. S., Nir, G., Antier, S., Jegou du Laz, T., van der Walt, S., Crellin-Quick, A., Culino, T., Duev, D. A., Goldstein, D. A., Healy, B. F., Karambelkar, V., Lilleboe, J., Shin, K. M., Singer, L. P., Ahumada, T., Anand, S., Bellm, E. C., Dekany, R., Graham, M. J., Kasliwal, M. M., Kostadinova, I., Weizmann Kien-drebeogo, R., Kulkarni, S. R., Jenkins, S., LeBaron, N., Mahabal, A. A., Neill, J. D., Parazin, B., Peloton, J., Perley, D. A., Riddle, R., Rusholme, B., van Santen, J., Sollerman, J., Stein, R., Turpin, D., Wold, A., Amat, C., Bonnefon, A., Bonnefoy, A., Flament, M., Kerkow, F., Kishore, S., Jani, S., Mahanty, S. K., Liu, C., Llinares, L., Makarison, J., Olliéric, A., Perez, I., Pont, L., and Sharma, V. A data science platform to enable time-domain astronomy. *arXiv e-prints*, art. arXiv:2305.00108, April 2023. doi: 10.48550/arXiv.2305.00108.
- Dieleman, S., Willett, K. W., and Dambre, J. Rotation-invariant convolutional neural networks for galaxy morphology prediction. , 450(2):1441–1459, June 2015. doi: 10.1093/mnras/stv632.
- Duev, D., Shin, K. M., and Singer, L. dmitry-duev/penquins: a python client for dmitryduev/kowalski, November 2021. URL <https://doi.org/10.5281/zenodo.5651471>.
- Duev, D. A., Mahabal, A., Masci, F. J., Graham, M. J., Rusholme, B., Walters, R., Karmarkar, I., Frederick, S., Kasliwal, M. M., Rebbapragada, U., and Ward, C. Real-bogus classification for the Zwicky Transient Facility using deep learning. , 489(3):3582–3590, November 2019. doi: 10.1093/mnras/stz2357.
- Fremling, C., Miller, A. A., Sharma, Y., Dugas, A., Perley, D. A., Taggart, K., Sollerman, J., Goobar, A., Graham, M. L., Neill, J. D., Nordin, J., Rigault, M., Walters, R., Andreoni, I., Bagdasaryan, A., Belicki, J., Cannella, C., Bellm, E. C., Cenko, S. B., De, K., Dekany, R., Frederick, S., Golkhou, V. Z., Graham, M. J., Helou, G., Ho, A. Y. Q., Kasliwal, M. M., Kupfer, T., Laher, R. R., Mahabal, A., Masci, F. J., Riddle, R., Rusholme, B., Schulze, S., Shupe, D. L., Smith, R. M., van Velzen, S., Yan, L., Yao, Y., Zhuang, Z., and Kulkarni, S. R. The Zwicky Transient Facility Bright Transient Survey. I. Spectroscopic Classification and the Redshift Completeness of Local Galaxy Catalogs. , 895(1):32, May 2020. doi: 10.3847/1538-4357/ab8943.

- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., Fernández del Río, J., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. Array programming with NumPy. *Nature*, 585:357–362, 2020. doi: 10.1038/s41586-020-2649-2.
- Hunter, J. D. Matplotlib: A 2D Graphics Environment. *Computing in Science and Engineering*, 9:90–95, May 2007. doi: 10.1109/MCSE.2007.55.
- Itagaki, K. Transient Discovery Report for 2023-05-19. *Transient Name Server Discovery Report*, 2023-1158:1, May 2023.
- Kaiser, N., Aussel, H., Burke, B. E., Boesgaard, H., Chambers, K., Chun, M. R., Heasley, J. N., Hodapp, K.-W., Hunt, B., Jedicke, R., Jewitt, D., Kudritzki, R., Luppino, G. A., Maberry, M., Magnier, E., Monet, D. G., Onaka, P. M., Pickles, A. J., Rhoads, P. H. H., Simon, T., Szalay, A., Szapudi, I., Tholen, D. J., Tonry, J. L., Waterson, M., and Wick, J. Pan-STARRS: A Large Synoptic Survey Telescope Array. In Tyson, J. A. and Wolff, S. (eds.), *Survey and Other Telescope Technologies and Discoveries*, volume 4836 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pp. 154–164, December 2002. doi: 10.1117/12.457365.
- Kingma, D. P. and Ba, J. Adam: A Method for Stochastic Optimization. *arXiv e-prints*, art. arXiv:1412.6980, December 2014. doi: 10.48550/arXiv.1412.6980.
- Lanusse, F., Ma, Q., Li, N., Collett, T. E., Li, C.-L., Ravanbakhsh, S., Mandelbaum, R., and Póczos, B. CMU DeepLens: deep learning for automatic image-based galaxy-galaxy strong lens finding. , 473(3):3895–3906, January 2018. doi: 10.1093/mnras/stx1665.
- pandas development team, T. pandas-dev/pandas: Pandas, February 2020. URL <https://doi.org/10.5281/zenodo.3509134>.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Perley, D. A., Fremling, C., Sollerman, J., Miller, A. A., Dahiwal, A. S., Sharma, Y., Bellm, E. C., Biswas, R., Brink, T. G., Bruch, R. J., De, K., Dekany, R., Drake, A. J., Duev, D. A., Filippenko, A. V., Gal-Yam, A., Gobar, A., Graham, M. J., Graham, M. L., Ho, A. Y. Q., Irani, I., Kasliwal, M. M., Kim, Y.-L., Kulkarni, S. R., Mahabal, A., Masci, F. J., Modak, S., Neill, J. D., Nordin, J., Riddle, R. L., Soumagnac, M. T., Strotjohann, N. L., Schulze, S., Taggart, K., Tzanidakis, A., Walters, R. S., and Yan, L. The Zwicky Transient Facility Bright Transient Survey. II. A Public Statistical Sample for Exploring Supernova Demographics. , 904(1):35, November 2020. doi: 10.3847/1538-4357/abbd98.
- Perley, D. A., Gal-Yam, A., Irani, I., and Zimmerman, E. LT Classification of SN 2023ixf as a Type II Supernova in M101. *Transient Name Server AstroNote*, 119:1, May 2023.
- Reznick, L. Using cron and crontab. *Sys Admin*, 2(4):29–32, 1993.
- Simonyan, K. and Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv e-prints*, art. arXiv:1409.1556, September 2014. doi: 10.48550/arXiv.1409.1556.
- Tachibana, Y. and Miller, A. A. A Morphological Classification Model to Identify Unresolved PanSTARRS1 Sources: Application in the ZTF Real-time Pipeline. , 130(994):128001, December 2018. doi: 10.1088/1538-3873/aae3d9.
- Villar, V. A., Hosseinzadeh, G., Berger, E., Ntampaka, M., Jones, D. O., Challis, P., Chornock, R., Drout, M. R., Foley, R. J., Kirshner, R. P., Lunnan, R., Margutti, R., Milisavljevic, D., Sanders, N., Pan, Y.-C., Rest, A., Scolnic, D. M., Magnier, E., Metcalfe, N., Wainscoat, R., and Waters, C. SuperRAENN: A Semisupervised Supernova Photometric Classification Pipeline Trained on Pan-STARRS1 Medium-Deep Survey Supernovae. , 905(2):94, December 2020. doi: 10.3847/1538-4357/abc6fd.
- Wes McKinney. Data Structures for Statistical Computing in Python. In Stéfan van der Walt and Jarrod Millman (eds.), *Proceedings of the 9th Python in Science Conference*, pp. 56 – 61, 2010. doi: 10.25080/Majora-92bf1922-00a.