
A Convolutional Autoencoder-Based Pipeline for Anomaly Detection and Classification of Periodic Variables

Ho-Sang, Chan

Department of Physics
The Chinese University of Hong Kong
Sha Tin, NT, Hong Kong
Center of Computational Astrophysics
Flatiron Institute
New York City, NY, USA
hschan@phy.cuhk.edu.hk

Siu-Hei, Cheung

Department of Physics
The Chinese University of Hong Kong
Sha Tin, NT, Hong Kong
Center of Computational Astrophysics
Flatiron Institute
New York City, NY, USA
shcheungpeter@link.cuhk.edu.hk

V. Ashley Villar

Department of Astronomy & Astrophysics
Institute for Computational & Data Sciences
Institute for Gravitation and the Cosmos
The Pennsylvania State University
University Park, PA, USA
vav5084@psu.edu

Shirley Ho

Center of Computational Astrophysics
Flatiron Institute
New York City, NY, USA
shirleyho@flatironinstitute.org

Abstract

The periodic pulsations of stars teach us about their underlying physical process. We present a convolutional autoencoder-based pipeline as an automatic approach to search for out-of-distribution anomalous periodic variables within The Zwicky Transient Facility Catalog of Periodic Variable Stars (ZTF CPVS). We use an isolation forest to rank each periodic variable by its anomaly score. Our overall most anomalous events have a unique physical origin: they are mostly highly variable and irregular evolved stars. Multiwavelength data suggest that they are most likely Red Giant or Asymptotic Giant Branch stars concentrated in the Milky Way galactic disk. Furthermore, we show how the learned latent features can be used for the classification of periodic variables through a hierarchical random forest. This novel semi-supervised approach allows astronomers to identify the most anomalous events within a given physical class, significantly increasing the potential for scientific discovery.

1 Introduction

Anomaly detection is an essential aspect of discovery in astronomy, as historically highlighted in the discovery of dark matter [1], Type Ia supernovae [2], and galaxies lacking dark matter [3]. As deep-sky surveys via modern telescopes continue to exponentially increase our discovery rates of variable cosmic phenomena, researchers are turning towards automated methods of anomaly detection [4–6].

The time for advanced techniques to search for anomalous astrophysical events is ripe. The Legacy Survey of Space and Time (LSST) conducted by the Vera Rubin Observatory is expected to commence in 2024 [7] and over 40 billion objects are to be observed within the 10 years of operation [8]. It is

reasonable to expect *anomalous* periodic variables stars (PVSs) which will defy expectations. Indeed, their discoveries have already been challenging our understanding of the chemical composition of stars [9, 10] and stellar evolution models [11]. Previous data-driven studies of PVSs that utilize machine learning have focused on classifications [12–14] and parameter estimation [15]. Here, we provide an anomaly detection pipeline to effectively search for peculiar PVSs detected with The Zwicky Transient Facility, a pathfinder survey for LSST. The Zwicky Transient Facility contains numerous publicly available data which serve as a good training set for big data problems in astronomy which will arise with LSST. We anticipate that our pipeline can be applied to these future surveys.

2 Method

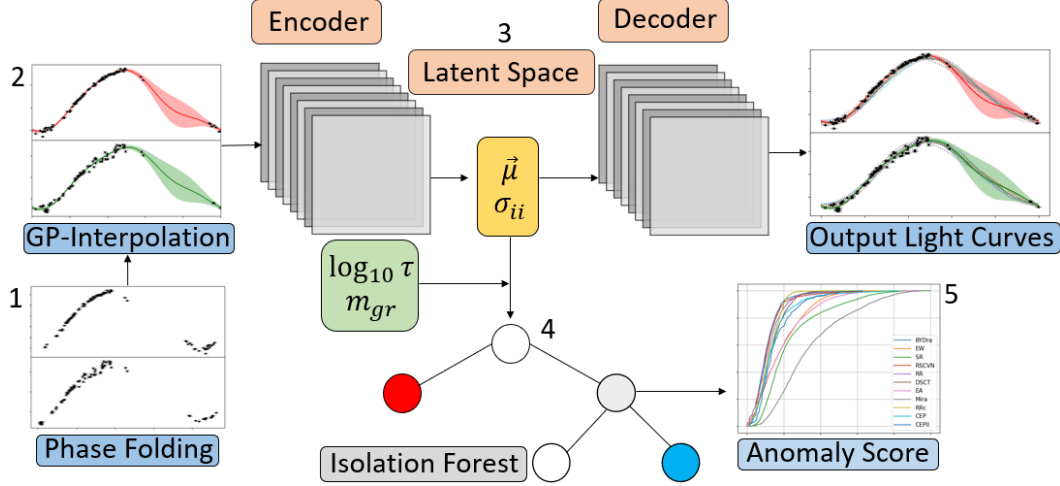


Figure 1: The anomaly detection pipeline: 1. Phase-folding the raw detection data. 2. Interpolation using the Multivariate Gaussian Process. 3. Encoding to learn latent vectors $\vec{\mu}$ and matrix elements σ_{ii} . 4. Append $\log_{10} \tau$ and m_{gr} to $\vec{\mu}$. 5. Isolation forest and ranking the anomalies.

Here we describe our training set and methodology. Our training set consists of the ZTF CPVS presented in Chen *et al.* [16]. The ZTF CPVS utilizes the Data Release 2 archive of the Zwicky Transient Facility [17] to search for and classify new PVSs down to a r -band magnitude, a measurement of brightness, of ~ 20.6 . They find a total of 781,602 PVSs, of which 621,702 are newly classified. The data are given as a multivariate, irregular time series of magnitudes, a measurement of brightness, in two filters (g -band and r -band) which are observed asynchronously. This time series is referred to as a light curve. The ZTF CPVS provides periods, which we use to phase-fold the light curves. To phase-fold a light curve, we cut the g - and r -band light curves into multiple sub-series with a time duration equal to the joint period, and we stack these sub-series on top of each other. Because the data are irregularly sampled by the telescope and taken in both bands, we choose to interpolate the phase-folded light curves using the multivariate Gaussian process [MGPR, 18, 5, 19] with periodic boundary conditions. The MGPR has a mean function $\eta(\phi, \lambda)$ and a covariance function K . We set $\eta(\phi, \lambda) = 0$, where ϕ is the temporal phase and λ is a wavelength in scaled units. We choose the following covariance function:

$$K(\vec{r}, \vec{r}') = C \exp\left(-\frac{|\vec{\phi} - \vec{\phi}'|^2}{l_\phi^2}\right) \exp\left(-\frac{|\vec{\lambda} - \vec{\lambda}'|^2}{l_\lambda^2}\right) + \begin{cases} \delta & \text{if } \vec{r} = \vec{r}', \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Here, $\vec{r} = (\vec{\phi}, \vec{\lambda})$ is a high dimensional position vector, l_ϕ and l_λ measure the correlation along the phase and wavelength direction respectively, C is a constant, and δ measures the white noise level of the raw data. After fitting for the kernel function, we generate 160 evenly spaced data points along with the phase direction for both g - and r -bands. Following the approach of Villar *et al.* [20] and Villar *et al.* [6], we stack both of them horizontally to form an ‘image’ of size 2×160 . The ‘images’ will be encoded through a convolutional variational autoencoder [C-VAE, 21] with a LeNet structure [22]. Our data size is 730,184 and we split them into a train to validation to test ratio of 7 : 2 : 1. The architecture of the encoder is described as follows:

1. **Input layer** of size 2×160
2. **3 Convolutional layers** with the *ReLU activation with Dropout*
3. **Dense layer** with 256 neurons, *Linear Activation*
4. **Latent space** of size 2×10

We adopt a kernel size of $(3, 4)$ and a stride of $(1, 2)$ for the convolutional layer. The filter size of the convolutional layer increases from $32 \rightarrow 64 \rightarrow 128$, and the dropout fraction is set to 0.1. The use of the ReLU activation is inspired by its robustness in various applications. We optimized the above-mentioned hyper-parameters through some rough grid searches, except for the kernel size and stride which are fixed. Given the variational nature of the autoencoder, the bottleneck latent space is described by a mean vector $\vec{\mu}$ and the diagonal covariance matrix σ_{ii} . Inspired by the work of Zhang and Bloom [13], we apply periodic padding to the ‘images’ during the convolution to enforce periodic boundary conditions. Our decoder is the symmetric counterpart of the encoder; however, no periodic padding and no dropout are applied. We train with the ADAM optimizer with default learning rates and a batch size of 1024. We train for 895 epochs for roughly ~ 6 hours by operating on the NVIDIA RTX2080-Ti GPU.

Once trained, we next use an isolation forest to rank each variable by a calculated anomaly score. The isolation forest works by building an ensemble of binary decision trees, which work to isolate samples of the population. The anomalies require few trees to isolate the event [23]; in this work, we focus on events with the top 100 most anomalous scores. We use the latent vectors $\vec{\mu}$, the log of period $\log_{10}\tau$, and the difference between the average g - and r -band magnitude ($m_{gr} := \langle m_g \rangle - \langle m_r \rangle$) as input features of the isolation forest [24]. We note that the period and magnitudes are explicitly included because this information is lost in the pre-processing used to train our autoencoder; however, we believe that they will be valuable in filtering out anomalies. We use the isolation forest implemented in the `scikit-learn` package with based estimators of 100,000, which we found sufficient to yield converged scores for the top anomalies.

We additionally use our learned latent space (and some hand-engineered features, including the joint period, and the amplitude and mean magnitudes of both the g - and r - band light curves) to *classify* the ZTF CPVS by their physical origin. We note that the ZTF CPVS provides class labels for PVSs, but the classifications are based on hand-engineered features only. Here we provide a new classification method that utilizes a hierarchical random forest classifier. We extracted event labels from the SIMBAD catalog [25] by cross-matching (using `Astroquery` [26]) their sky-coordinates with those listed in the ZTF CPVS. The SIMBAD catalog contains class labels obtained, typically, through spectroscopic analysis, which is more reliable but often expensive. We found 31,541 successfully cross-matched objects. We then construct 13 classes in 2 levels. The first level includes Active Galactic Nuclei-like (AGNL), Cepheid (CEP), Eclipsing Binaries (EB), Long-Period Variables (LPV), Mira variables (Mira), other Pulsating Variables (Pul_{oth}), RR Lyrae (RR), and Peculiar (Pec). The second level is further classifications of the Pec type, and it includes Carbon stars (C-Type), Horizontal Branch stars (HB), Red Giant Branch stars (RGB), S-Type stars (S-Type), Young Stellar Object-like (YSOL), and Other Variables (V_{oth}). They will serve as the data set for our classification model. We split the data set into a training-to-test set ratio of 7 : 3 by using the python package `scikit-learn` [27]. We note that our training set is highly imbalanced, with the largest set containing 10,745 events and the smallest containing just 41 events. We balanced the training set using the python package `imbalanced-learn` [28] with default learning parameters, which performs synthetic minority resampling [29, 28]. Finally, we train the hierarchical random forest classifier provided by `imbalanced-learn`, with no hyper-parameter optimization conducted.

3 Results and Discussion

We find that our learned latent space exhibits an annular structure, which inspires us to transform the latent space using spherical coordinates. We run our isolation forest on both sets of coordinates and compare results, selecting the top 100 most anomalous events in both cases. Our anomaly detection algorithm is seemingly sensitive to irregular oscillations which consist of several dominant Fourier modes. In general, these anomalous events are both multi-modal and highly irregular, with some examples exhibiting larger fluctuations that span over several magnitudes. Some of the anomalies also show weak or even anti-correlation between the g - and r - band light curves, suggesting significant temperature variations within the pulsations.

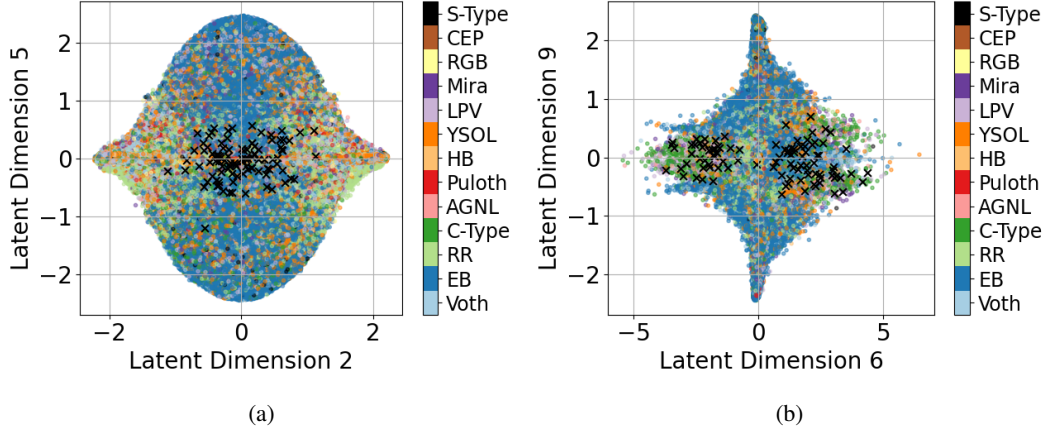


Figure 2: (a) and (b) Examples of the latent distributions for different PVSs labeled by distinct colors. Anomalies are marked as black crosses.

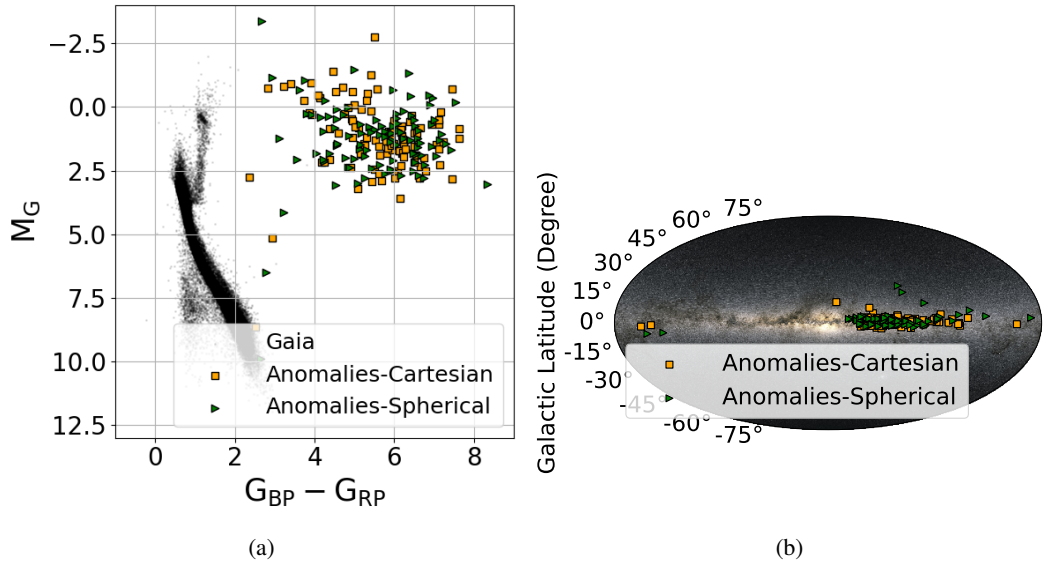


Figure 3: (a) Distribution of the top 100 anomalous events on the Gaia HR diagram (a common feature space in astronomy) against the main-sequence stars in black dots. Data are taken from Lindegren *et al.* [30], Bailer-Jones *et al.* [31], and Gaia Collaboration *et al.* [32], through the Vizier Catalogue [33] by Astroquery [26]. (b) The galactic distribution of the same set of anomalies. Note the tight distribution along the Milky Way galactic plane. Image credit: ESA/Gaia/DPAC.

To better understand the nature of the selected anomalies, we extracted Gaia [32] G-band absolute magnitudes M_G , and the difference between the Gaia B-band and R-band absolute magnitudes $G_{BP} - G_{RP}$ for the top 100 anomalies in both latent spaces. We plot their distribution in Figure 3 (a). We note that this is a common diagnostic phase space, which roughly correlates with the temperature and luminosity of the stars. The majority of the anomalies are *cool*, with $G_{BP} - G_{RP} > 4$ and *luminous*, with $0 < M_G < 2.5$. These properties correspond to evolved, luminous, and cold stars. Furthermore, we show the distribution of the top anomalies in the Milky Way galactic coordinates in Figure 3 (b). The majority of the anomalies concentrate in the Milky Way galactic disk, implying that (1) there is likely significant interstellar reddening due to the dust for these events and (2) the progenitor systems of these events are consistent with young and massive stars. Taken together, the observational evidence points to highly anomalous, young, cool, and massive Red Giant or Asymptotic Giant Branch stars. Spectroscopic follow-up observations and detailed light curve modeling are *essential* in fully understanding the anomalies detected in our data-driven pipeline; however, we note that these

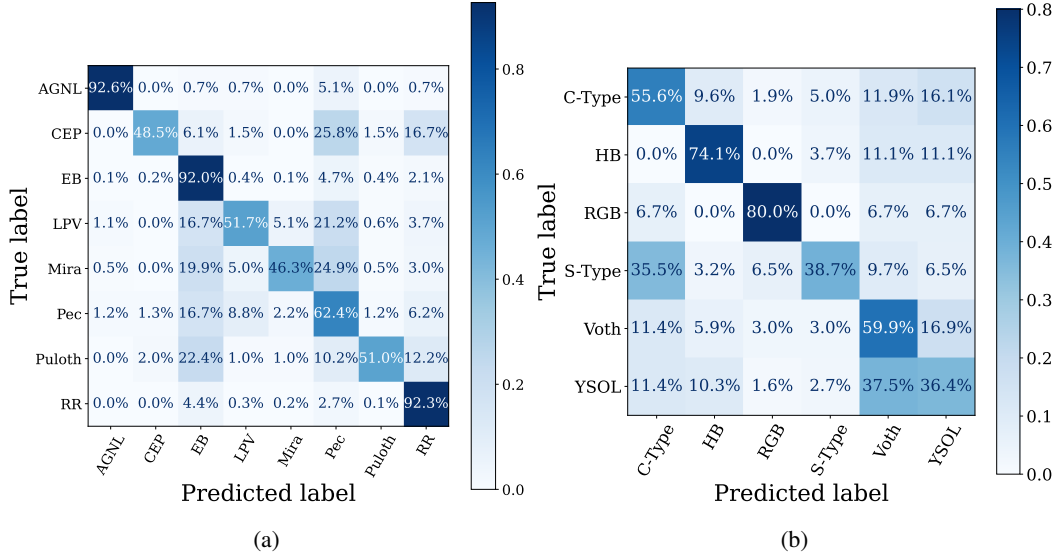


Figure 4: (a) Confusion matrix of the first level classification labels of the test set. We show in each row the completeness of each class. (b) Same as (a), but for the second level labels (of Peculiar type).

anomalies were discovered with *limited* survey observations. Similar techniques will be invaluable for future missions.

Finally, we highlight the classification results from our hierarchical classifier in Figure 4 (a) and (b). We find the ZTF CPVS likely contains variable, non-stellar objects. For example, we find that 11837 out of 97137 (12.2 %) semi-regular *galactic* variable stars in the ZTF CPVS are classified as AGNL objects (objects associated with supermassive black holes in other galaxies) in our new classification model. We note that we do not use the same class breakdown as the original ZTF CPVS, making it difficult to directly use the ZTF CPVS as a baseline. A detailed comparison between classifiers will be left to future work. However, we find that our learned latent space is sufficient to classify the set of ZTF CPVS light curves explored here with reasonable levels of accuracy. We plot the latent distribution using our new labels in Figures 2 (a) and (b), to show the robustness of our autoencoder in distinguishing objects that belongs to different categories.

4 Conclusion And Broader Impact

We present a convolutional autoencoder-based pipeline as an automatic yet robust approach to search for and classify anomalous PVSs. We note that Malanchev *et al.* [34] recently presented an anomaly search within the Zwicky Transient Facility transient data using a deep-generative learning approach, but without a specific focus on PVSs. Our pipeline is specifically for PVSs, and could be applied to PVS data obtained from deep-sky surveys in the future. Detailed spectroscopic follow-up is essential to reveal their true identity of the anomalous PVSs identified here and how these anomalies fit into our current understanding of late-stage stellar processes.

Astrophysical science has entered the big-data era. We anticipate our method can contribute to the community, such as detecting other anomalous periodic/non-periodic transient or better classifying astronomical objects. Additionally, our unsupervised learning pipeline can make use of supervised object labels to look for categorical-wise anomalies. Techniques similar to those presented here can be used in broader applications to identify anomalies in periodic/non-periodic, multi-variate time series.

5 Acknowledgement

We thank the Flatiron institute for providing computer cluster access. We thank Prof. Ming-Chung, Chu in the Chinese University of Hong Kong for providing the NVIDIA RTX2080-Ti GPU as the computational resources for the neural network training. We also thank Prof. Maria Drout and Anna O’Grady at the University of Toronto; Matteo Cantiello, Mathieu Renzo, and Adam Jermyn at the Center of Computational Astrophysics, Flatiron Institute for their valuable discussion on the anomalous periodic variables stars. VAV acknowledges support by the Simons Foundation through a Simons Junior Fellowship (#718240) during the early phases of this project, as well as support in part by the NSF through grant AST-2108676.

6 Checklist

For all authors...

- (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? Yes
- (b) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes
- (c) Did you discuss any potential negative societal impacts of your work?. Yes
- (d) Did you describe the limitations of your work? Yes
- (a) If your work uses existing assets, did you cite the creators? Yes
- (b) Did you mention the license of the assets? Yes
- (c) Did you include any new assets either in the supplemental material or as a URL? Yes
- (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? Yes
- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? Yes

References

- [1] V. C. Rubin and J. Ford, W. Kent, *ApJ* **159**, 379 (1970).
- [2] W. Li, A. V. Filippenko, E. Gates, R. Chornock, A. Gal-Yam, E. O. Ofek, D. C. Leonard, M. Modjaz, R. M. Rich, A. G. Riess, and R. R. Treffers, *Publications of the Astronomical Society of the Pacific* **113**, 1178 (2001).
- [3] P. van Dokkum, S. Danieli, Y. Cohen, A. Merritt, A. J. Romanowsky, R. Abraham, J. Brodie, C. Conroy, D. Lokhorst, L. Mowla, E. O’Sullivan, and J. Zhang, *Nature* **555**, 629 (2018).
- [4] M. Henrion, D. J. Mortlock, D. J. Hand, and A. Gandy, Classification and anomaly detection for astronomical survey data, in *Astrostatistical Challenges for the New Astronomy*, edited by J. M. Hilbe (Springer New York, New York, NY, 2013) pp. 149–184.
- [5] M. V. Pruzhinskaya, K. L. Malanchev, M. V. Kornilov, E. E. O. Ishida, F. Mondon, A. A. Volnova, and V. S. Korolev, *Monthly Notices of the Royal Astronomical Society* **489**, 3591 (2019), <https://academic.oup.com/mnras/article-pdf/489/3/3591/30029513/stz2362.pdf> .
- [6] V. A. Villar, M. Cranmer, E. Berger, G. Contardo, S. Ho, G. Hosseinzadeh, and J. Yao-Yu Lin, *arXiv e-prints* , arXiv:2103.12102 (2021), arXiv:2103.12102 [astro-ph.HE] .
- [7] M. Graham, in *The Extragalactic Explosive Universe: the New Era of Transient Surveys and Data-Driven Discovery* (2019) p. 23.
- [8] Ž. Ivezić *et al.*, *The Astrophysical Journal* **873**, 111 (2019).
- [9] K. Bernhard, S. Hümmerich, E. Paunzen, and J. Šupíková, *Monthly Notices of the Royal Astronomical Society* 10.1093/mnras/stab2065

- (2021), [stab2065](https://academic.oup.com/mnras/stab2065), <https://academic.oup.com/mnras/advance-article-pdf/doi/10.1093/mnras/stab2065/39170168/stab2065.pdf> .
- [10] E. Niemczura, S. Hümmerich, F. Castelli, E. Paunzen, K. Bernhard, F.-J. Hamsch, and K. Hełminiak, *Scientific Reports* **7**, 5906 (2017).
 - [11] M. A. T. Groenewegen and M. I. Jurkovic, *A&A* **604**, A29 (2017), arXiv:1705.04487 [astro-ph.SR] .
 - [12] S. Jamal and J. S. Bloom, *The Astrophysical Journal Supplement Series* **250**, 30 (2020).
 - [13] K. Zhang and J. S. Bloom, *Monthly Notices of the Royal Astronomical Society* **505**, 515 (2021), <https://academic.oup.com/mnras/article-pdf/505/1/515/38334334/stab1248.pdf> .
 - [14] B. Naul, J. S. Bloom, F. Pérez, and S. van der Walt, *Nature Astronomy* **2**, 151 (2018), arXiv:1711.10609 [astro-ph.IM] .
 - [15] J. Martínez-Palomera, J. S. Bloom, and E. S. Abrahams, arXiv e-prints , arXiv:2005.07773 (2020), arXiv:2005.07773 [astro-ph.IM] .
 - [16] X. Chen, S. Wang, L. Deng, R. de Grijs, M. Yang, and H. Tian, *The Astrophysical Journal Supplement Series* **249**, 18 (2020).
 - [17] E. C. Bellm *et al.*, *PASP* **131**, 018002 (2019), arXiv:1902.01932 [astro-ph.IM] .
 - [18] D. Foreman-Mackey, S. Hoyer, J. Bernhard, and R. Angus, *george: George (v0.2.0)* (2014).
 - [19] H. Qu, M. Sako, A. Möller, and C. Doux, arXiv e-prints , arXiv:2106.04370 (2021), arXiv:2106.04370 [astro-ph.IM] .
 - [20] V. A. Villar, G. Hosseinzadeh, E. Berger, M. Ntampaka, D. O. Jones, P. Challis, R. Chornock, M. R. Drout, R. J. Foley, R. P. Kirshner, R. Lunnan, R. Margutti, D. Milisavljevic, N. Sanders, Y.-C. Pan, A. Rest, D. M. Scolnic, E. Magnier, N. Metcalfe, R. Wainscoat, and C. Waters, *The Astrophysical Journal* **905**, 94 (2020).
 - [21] D. P. Kingma and M. Welling, arXiv e-prints , arXiv:1312.6114 (2013), arXiv:1312.6114 [stat.ML] .
 - [22] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, *Neural Computation* **1**, 541 (1989), <https://direct.mit.edu/neco/article-pdf/1/4/541/811941/neco.1989.1.4.541.pdf> .
 - [23] F. T. Liu, K. M. Ting, and Z.-H. Zhou, in *2008 Eighth IEEE International Conference on Data Mining* (2008) pp. 413–422.
 - [24] F. T. Liu, K. M. Ting, and Z.-H. Zhou, *ACM Trans. Knowl. Discov. Data* **6**, 10.1145/2133360.2133363 (2012).
 - [25] M. Wenger, F. Ochsenbein, D. Egret, P. Dubois, F. Bonnarel, S. Borde, F. Genova, G. Jasniewicz, S. Laloë, S. Lesteven, and R. Monier, *A&AS* **143**, 9 (2000), arXiv:astro-ph/0002110 [astro-ph] .
 - [26] A. Ginsburg *et al.*, *AJ* **157**, 98 (2019), arXiv:1901.04520 [astro-ph.IM] .
 - [27] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, *Journal of Machine Learning Research* **12**, 2825 (2011).
 - [28] G. Lemaître, F. Nogueira, and C. K. Aridas, *Journal of Machine Learning Research* **18**, 1 (2017).
 - [29] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, *Journal of artificial intelligence research* **16**, 321 (2002).
 - [30] L. Lindgren *et al.*, *A&A* **616**, A2 (2018), arXiv:1804.09366 [astro-ph.IM] .
 - [31] C. A. L. Bailer-Jones, J. Rybizki, M. Fouesneau, M. Demleitner, and R. Andrae, *AJ* **161**, 147 (2021), arXiv:2012.05220 [astro-ph.SR] .

- [32] Gaia Collaboration *et al.*, *A&A* **649**, A1 (2021), arXiv:2012.01533 [astro-ph.GA] .
- [33] B. A. Skiff, *VizieR Online Data Catalog* , B/mk (2014).
- [34] K. L. Malanchev, M. V. Pruzhinskaya, V. S. Korolev, P. D. Aleo, M. V. Kornilov, E. E. O. Ishida, V. V. Krushinsky, F. Mondon, S. Sreejith, A. A. Volnova, A. A. Belinski, A. V. Dodin, A. M. Tatarnikov, S. G. Zheltoukhov, and (The SNAD Team), *MNRAS* **502**, 5147 (2021), arXiv:2012.01419 [astro-ph.IM] .