Multi-Class Deep SVDD: Anomaly Detection Approach in Astronomy with Distinct Inlier Categories

M. Pérez-Carrasco [1 2 3] **G. Cabrera-Vives** [2 1 3] **L. Hernández-García** [4 3] **F. Förster** [3 5 6] **P. Sánchez-Sáez** [7 3]
A. Muñoz Arancibia [3 5] **N. Astorga** [8 3] **F. E. Bauer** [3 9 10] **A. Bayo** [4 7] **M. Cádiz-Leyton** [2 1] **M. Catelan** [3 9 10]
P. A. Estévez [3 8]

Abstract

With the increasing volume of astronomical data generated by modern survey telescopes, automated pipelines and machine learning techniques have become crucial for analyzing and extracting knowledge from these datasets. Anomaly detection, i.e. the task of identifying irregular or unexpected patterns in the data, is a complex challenge in astronomy. In this paper, we propose Multi-Class Deep Support Vector Data Description (MCDSVDD), an extension of the state-of-the-art anomaly detection algorithm One-Class Deep SVDD, specifically designed to handle different inlier categories with distinct data distributions. MCDSVDD uses a neural network to map the data into hyperspheres, where each hypersphere represents a specific inlier category. The distance of each sample from the centers of these hyperspheres determines the anomaly score. We evaluate the effectiveness of MCDSVDD by comparing its performance with several anomaly detection algorithms on a large dataset of astronomical light-curves obtained from the Zwicky Transient Facility. Our results demonstrate the efficacy of MCDSVDD in

detecting anomalous sources while leveraging the presence of different inlier categories. The code and the data needed to reproduce our results are publicly available at https://github.com/mperezcarrasco/AnomalyALeRCE.

1. Introduction

With modern survey telescopes producing unprecedented volumes of data, it has become unfeasible to analyze astronomical data massively by human inspection. Consequently, the necessity for automated pipelines arises, enabling the extraction of knowledge from these vast datasets in a data-driven manner. Among the intriguing and complex astronomical challenges lies the task of anomaly detection (AD), which refers to the identification of irregular or unexpected patterns that deviate from our existing understanding of the data (Chandola et al., 2009).

In recent years, the application of machine learning techniques in the field of astronomy has led to remarkable advancements in detecting anomalous sources in a systematic manner. For instance, Xiong et al. 2010 employed a hierarchical probabilistic model, while Baron & Poznanski 2016 utilized an unsupervised Random Forest approach to identify outliers within galaxy spectra from the Sloan Digital Sky Survey (SDSS). More recently, Villar et al. 2020 applied an IForest algorithm to a latent space derived from a Variational Recurrent Autoencoder (VRAE), specifically targeting a simulated dataset of supernovae (SNe). Similarly, Sánchez-Sáez et al. 2021a employed a VRAE architecture to analyze the light curves of active galactic nuclei (AGN) within the Zwicky Transient Facility Data Release 5 (ZTF DR5). Furthermore, the work of Ishida et al. 2021 employed Gaussian Processes to extract features and train Isolation Forest algorithm in an active learning manner using the Open Supernova Catalog dataset and the Photometric LSST Astronomical Time-series Classification Challenge (PLAsTiCC; Kessler et al. 2019; Hložek et al. 2020). In a similar vein, Muthukrishna et al. 2022 used a bayesian parametric model to find anomalous sources using simulated light curves from SNANA (Kessler et al., 2009). Previous works

---

[1]Data Science Unit, Universidad de Concepción, Concepción, Chile [2]Department of Computer Science, Universidad de Concepción, Concepción, Chile [3]Millennium Institute of Astrophysics (MAS), Santiago, Chile [4]Instituto de Física y Astronomía, Universidad of Valparaíso, Valparaíso, Chile [5]Center for Mathematical Modelling, Universidad de Chile, Santiago, Chile [6]Data and Artificial Intelligence Initiative, Universidad de Chile, Santiago, Chile [7]European Southern Observatory, Garching bei München, Germany [8]Department of Electrical Engineering, Universidad de Chile, Santiago, Chile [9]Instituto de Astrofísica, Facultad de Física, Pontificia Universidad Católica de Chile, Santiago, Chile [10]Centro de Astroingeniería, Pontificia Universidad Católica de Chile, Santiago, Chile. Correspondence to: Manuel Pérez-Carrasco <maperezc@inf.udec.cl>.

share a common training procedure: they utilize a sample of normal and well-identified objects categorized as inliers to train an algorithm that can detect samples deviating from these inliers. This approach, commonly known as one-class anomaly detection (Schölkopf et al., 1999), represents the current state-of-the-art for various anomaly detection problems. However, when different inlier categories with distinct data distributions are present in the training set, it remains unclear how to effectively utilize and leverage these methods.

To address this challenge, we present Multi-Class Deep Support Vector Data Description (MCDSVDD), an extension of the state-of-the-art anomaly detection algorithm One-Class Deep SVDD (Ruff et al., 2018). Our proposed method employs a neural network to map the data into hyperspheres, where each hypersphere encapsulates the data representation of a specific inlier category. Consequently, the "weirdness" of each sample is determined by its distance from the center of the closest hypersphere. To evaluate the effectiveness of our approach, we compare the performance of several one-class anomaly detection algorithms on a large dataset of astronomical light-curves obtained from the Zwicky Transient Facility (ZTF; Bellm et al. 2018). Through our comprehensive evaluation, we demonstrate the efficacy of MCDSVDD in detecting anomalous sources while effectively leveraging the presence of different inlier categories.

2. Data

In this work we use data alerts from the from the ZTF data stream (Bellm et al., 2018). In this data stream, an alert is triggered by an object in the sky whose current (science) image has a significant difference with respect to a template (reference) image (Masci et al., 2018). For alerts to be streamed by ZTF, they need to pass the cut-off criteria defined by the real/bogus detection system designed by the ZTF Collaboration. These criteria include signal-to-noise ratios, near-edge image positioning, negative and bad pixels, and morphological and photometric features (Mahabal et al., 2019; Duev et al., 2019).

Light-curves are constructed using the same procedure of Section 4.4 of Förster et al. (2021). We perform cross-match with the AllWISE[1] public source catalog (Wright et al., 2010; Mainzer et al., 2011), using a matching radius of 2 arcseconds, obtaining W1, W2, and W3 photometry. Then, 152 features are calculated for each light-curve using the ALeRCE broker feature extractor (Förster et al., 2021). These features include detection and non-detection features (see Section 3 of Sánchez-Sáez et al. 2021b for details). We only use objects with at least six detections in either g or r

---

[1]The AllWISE Data Release can be found at `http://wise2.ipac.caltech.edu/docs/release/allwise/`

bands.

Our training set consisted in labeled light-curves that follow a taxonomy with three main classes (hereafter, the top level classes): transient, stochastic, and periodic. Each of these categories is then subdivided into the following subclasses:

- Transient: Type Ia supernova (SNIa), Type Ibc supernova (SNIbc), Type II supernova (SNII), and superluminous supernova (SLSN);

- Stochastic: Type 1 Seyfert galaxy (AGN; i.e., host-dominated active galactic nuclei), Type 1 quasar (QSO; i.e., nucleus-dominated AGN), blazar (blazar; i.e, beamed jet dominated AGN), young stellar object (YSO; including bursters, dippers and purely rotation modulated lightcurves of pre-main sequence stars), and cataclysmic variable/Nova (CV/Nova);

- Periodic: long-period variable (LPV; includes regular, semi-regular, and irregular variable stars), RR Lyrae (RRL), Cepheid (CEP), eclipsing binary (E), and δ Scuti (DSCT).

In Section 4 we explain in detail how the above-mentioned categories are used to train and evaluate our anomaly detection algorithms by hiding the light curves from a given subclass and considering all the others as inliers. This allows for a rigorous comparison of algorithms that helps to select the most promising method to be applied in a real-world scenario.

3. Method

In this work we present Multi-Class Deep SVDD (MCSVDD), an extension of the One-Class Deep SVDD method (Ruff et al., 2018).

3.1. One-Class Deep SVDD

Deep Support Vector Data Description (Deep SVDD; Ruff et al., 2018) is a neural network-based approach related to OCSVM (Schölkopf et al., 1999). The general idea is to map the data from an input space $\mathcal{X} \subseteq \mathbb{R}^d$ into a new feature space $\mathcal{F} \subseteq \mathbb{R}^p$ using a neural network $\phi(\cdot; \Theta) : \mathcal{X} \rightarrow \mathcal{F}$ with L hidden layers and parameters $\Theta = \{\theta^1, ..., \theta^L\}$. Specifically, we want the neural network ϕ to learn a representation that minimizes the volume of a hypersphere of radius $R > 0$ and center $\mathbf{c}$ that enclose the normal data in the output space $\mathcal{F}$.

The objective function for this task is defined by:

$$\min_{\Theta,R} R^2 + \frac{1}{N}\sum_{i=1}^{N}\max\{0, ||\phi(\mathbf{x}_i;\Theta) - \mathbf{c}||^2 - R^2\} + \frac{\lambda}{2}\boldsymbol{\Theta}^\top\boldsymbol{\Theta},$$
$$(1)$$

where λ is a hyperparameter that controls the weight decay on the network parameters θ. In practice, the encoding part of an autoencoder can be used to map the data into the latent space and define the center $\mathbf{c}$. The parameters θ and R are alternately and iteratively optimized via gradient descent and line-search respectively.

In cases when training data contains only normal samples, one can simply penalize the squared distance of all the data points with respect to a center $\mathbf{c}$ as follows (Ruff et al., 2018; 2019):

$$\min_{\Theta} \frac{1}{N}\sum_{i=1}^{N}\max ||\phi(\mathbf{x}_i;\Theta) - \mathbf{c}||^2 + \frac{\lambda}{2}\boldsymbol{\Theta}^\top\boldsymbol{\Theta}. \quad (2)$$

This loss function will be used for comparisons in Section 5.

3.2. Multi-Class Deep SVDD

We extend the One-Class Deep SVDD to take into account the different categories presented in our taxonomy. Instead of modeling a single hypersphere to enclose the normal data, MCSVDD models multiple hyperspheres, where each one corresponds to a given class and enclose normal data for that class. The idea is to learn a neural network that maps objects from the same class close to each other and far from objects from different classes. As abnormal samples come from unseen classes, their distances to each hypersphere should be larger than those of normal datapoints. We named this method Multi-Class Deep SVDD (MCSVDD) as we now have multiple inlier classes.

Assuming normal data pairs coming from M different classes $y \in \{1, ..., M\}$, the objective function for a training set composed by different inlier classes is given by:

$$\min_{\boldsymbol{\Theta}} \sum_{j=1}^{M}\frac{1}{N_j}\sum_{i=1}^{N}\mathbb{1}_{(y_i=j)}||\phi(\mathbf{x}_i;\Theta) - \mathbf{c}_j||^2 + \frac{\lambda}{2}\boldsymbol{\Theta}^\top\boldsymbol{\Theta}, \quad (3)$$

where λ is a hyperparameter that controls the weight decay regularizer on the network parameters θ, $\mathbb{1}_{(y_i=j)}$ is an indicator function that becomes 1 if $y_i = j$ and 0 otherwise, and N_j is the number of data points that belong to class j.

By following this approach, it is possible to define an anomaly score $A(\cdot)$ based on the distance of the data points to the centers of the hyperspheres as follows:

$$A(x_i) = \min_{j}\{||\phi(\mathbf{x}_i;\boldsymbol{\Theta}^*) - \mathbf{c}_j||^2\}. \quad (4)$$

where Θ^* are the parameters of the trained neural network.

4. Methodology

For fair comparisons, we follow the same training and evaluation procedure for all our experiments. From the ZTF light-curves, we extracted 152 features as described in Sec. 2. All the features were normalized between -1 and 1 using the quantile normalization method. We randomly split the data into a training set (80%) and a test set (20%) in a stratified fashion in order to preserve the proportion of samples per class. The training set is divided into five stratified subsets in order to perform 5-fold cross-validation for model selection. Fig. 1 shows a scheme of our training and evaluation methodology.
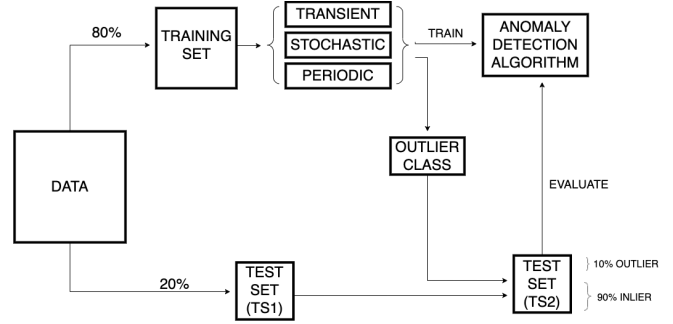

Figure 1. Methodology for training and evaluation of the anomaly detection algorithms. We split the data into a training set and test set, composed by 80% and 20% of the data, respectively. The training set is subdivided into transient, stochastic, and periodic data. For each of these classes, we choose each subclass as the outlier class. The outlier class is removed from the training set and added to the test set (TS2). Then, an anomaly detection algorithm is trained using the remaining objects of each of the classes, and is evaluated using TS2.

4.1. Training

To facilitate the problem at hand, our training dataset is divided into three primary classes: transient, stochastic, and periodic. These classes are further stratified into 14 subclasses, as described in Section 2. For each main class, we employ a dedicated anomaly detector, resulting in three separate detectors. Since anomalies are typically unknown in real-world scenarios, during cross-validation, we designate a subclass from each main class as the anomalous class and exclude it from the training set. Consequently, the model is trained on the remaining subclasses, with the excluded subclass reserved solely for evaluation purposes. This approach

Method	Transient				Stochastic					Periodic				
	SLSN	SNII	SNIa	SNIbc	AGN	Blazar	CV/Nova	QSO	YSO	CEP	DSCT	E	RRL	LPV
IForest	0.640	0.721	0.428	0.490	0.573	0.710	0.975	0.468	0.913	0.359	0.295	0.469	0.549	0.971
(Liu et al., 2012)	±0.014	±0.021	±0.032	±0.038	±0.017	±0.009	±0.001	±0.016	±0.003	±0.007	±0.012	±0.021	±0.033	±0.007
OCSVM	0.577	0.587	0.434	0.492	0.532	0.443	0.909	**0.517**	0.792	0.432	0.557	0.555	0.539	0.943
(Schölkopf et al., 1999)	±0.014	±0.014	±0.021	±0.011	±0.008	±0.002	±0.001	±0.003	±0.005	±0.004	±0.005	±0.003	±0.004	±0.001
AE	**0.736**	0.807	0.438	0.537	**0.701**	**0.762**	**0.980**	0.443	**0.990**	0.564	0.367	0.864	0.907	**0.996**
(Rumelhart & McClelland, 1987)	**±0.022**	±0.021	±0.015	±0.019	**±0.010**	**±0.006**	**±0.016**	±0.004	**±0.001**	±0.024	±0.015	±0.009	±0.015	**±0.000**
VAE	0.669	0.690	0.404	0.522	0.596	0.597	0.849	0.500	0.795	0.442	0.417	0.561	0.451	0.936
(Kingma & Welling, 2014)	±0.015	±0.023	±0.018	±0.025	±0.007	±0.010	±0.028	±0.009	±0.009	±0.010	±0.007	±0.007	±0.006	±0.007
Deep SVDD	0.644	0.731	0.475	0.507	0.496	0.607	0.932	0.411	0.901	0.707	0.482	0.636	0.774	0.785
(Ruff et al., 2018)	±0.043	±0.043	±0.040	±0.040	±0.025	±0.044	±0.015	±0.008	±0.022	±0.027	±0.054	±0.055	±0.068	±0.025
MCDSVDD	0.686	**0.828**	**0.624**	**0.584**	0.706	0.512	0.770	0.483	0.854	**0.858**	**0.819**	**0.945**	**0.953**	0.953
(This work)	±0.051	**±0.024**	**±0.039**	**±0.032**	±0.069	±0.113	±0.127	±0.080	±0.041	**±0.025**	**±0.015**	**±0.006**	**±0.003**	±0.008

Table 1: Evaluation of the performance of each model when applied to each of the top level taxonomy (transient, stochastic, periodic). Each row represents a different outlier detection algorithm, and each column represents the subclass considered as outlier. The performance is evaluated using the AUROC scores. Best metrics per class are marked in boldface

ensures that the model does not incorporate data from the chosen anomalous subclass during the training phase. The process is iterated for each subclass, thereby providing a comprehensive evaluation of the anomaly detectors' overall performance. The assessment primarily focuses on the detectors' ability to successfully identify the removed subclasses, aligning with established practices in the machine learning literature (Ruff et al., 2018).

4.2. Evaluation

Although real outlier events are typically unavailable for evaluation, we aim to create a realistic scenario to identify the most promising anomaly detection models. As previously mentioned, 20% of the dataset is designated as the test set (TS1) and is never utilized for training any model. We create a second test set (TS2), including all the objects from TS1 that belong to the inlier subclasses, and outliers from both TS1 and those who were removed from the training set. To ensure a realistic distribution, TS2 is composed of 10% outliers and 90% inliers. This setup enables a comprehensive assessment of the models' ability to detect the chosen outlier class within TS2. See Figure 1 for an illustrative diagram of training en evaluation methodology.

To evaluate model performance when a subclass is treated as an outlier, we employ the Area Under the Receiver Operating Characteristic curve (AUROC; Davis & Goadrich 2006).

5. Results

We conducted a comprehensive comparison of our proposed method, Multi-Class Deep SVDD (MCDSVDD), against five state-of-the-art algorithms for anomaly detection in the field of astronomy. The baseline algorithms we considered were Isolation Forest (IForest; Liu et al. 2012), One-class Support Vector Machine (OCSVM; Schölkopf et al. 1999), Autoencoder (AE; Rumelhart & McClelland 1987), Variational Autoencoders (VAE; Kingma & Welling 2014), and Deep Support Vector Data Description (Deep SVDD; Ruff

et al. 2018). See Appendix I for details and hyperparameters of all the algorithms used in this work. Table 1 presents the 5-fold AUROC values obtained by each anomaly detection model on the TS2 dataset. The table showcases the performance of each model when considering each of the 14 subclasses as an outlier in separate trials. The highest AUROC values for each subclass are highlighted in bold, indicating the best-performing model for that particular task.

As can be seen, our proposed method MCDSVDD consistently outperforms all other methods we considered for the transient and periodic objects. Our method is able to detect SNII, SNIa, and SNIb as outliers with higher AUROC values than the rest. For the SLSN subclass, AE's performance is statistically indistinguishable from MCDSVDD (i.e., not statistically significant; p-value = 0.0789). In terms of the periodic classes, MCDSVDD shows superior performance in detecting CEP, DSCT, E, and RRL as outliers, while AE demonstrates better performance in detecting LPVs. Regarding stochastic sources, AE outperforms the other methods for four out of the five subclasses (AGN, Blazar, CV/Nova, and YSO), while OCSVM achieves the highest performance for QSOs. It is worth mentioning that QSO light curves often exhibit slow and smooth temporal variations, making them more challenging to detect as anomalies. Therefore, AE emerges as the best outlier detection algorithm for stochastic sources, whereas MCDSVDD excels in detecting transient and periodic sources. We hypothesize that the high variability of stochastic light curves may impact MCDSVDD's capability to detect different types of sources, unlike AE, which benefits from its reconstruction loss function.

6. Conclusion

In this paper, we addressed the challenge of effectively utilizing different inlier categories with distinct data distributions in anomaly detection for astronomical data. We proposed Multi-Class Deep SVDD (MCDSVDD), an extension of the state-of-the-art One-Class Deep SVDD algorithm. MCDSVDD employs a neural network to

map the data into hyperspheres, where each hypersphere represents a specific inlier category. By measuring the distance of each sample from the centers of these hyperspheres, MCDSVDD determines the anomaly scores of the samples. We evaluated the performance of MCDSVDD along with several one-class anomaly detection algorithms on a large dataset of astronomical light-curves from the Zwicky Transient Facility (ZTF). The results demonstrated that MCDSVDD outperforms other algorithms in detecting anomalous sources while effectively leveraging the presence of different inlier categories. Our findings highlight the importance of considering distinct inlier categories in anomaly detection tasks and showcase the potential of MCDSVDD for identifying anomalies in astronomical data. Finally, the code and data to reproduce the results presented in this work are publicly available at https://github.com/mperezcarrasco/AnomalyALeRCE.

7. Acknowledgements

References

Baron, D. and Poznanski, D. The weirdest SDSS galaxies: results from an outlier detection algorithm. *Monthly Notices of the Royal Astronomical Society*, 465(4): 4530–4555, 11 2016. ISSN 0035-8711. doi: 10.1093/mnras/stw3021. URL https://doi.org/10.1093/mnras/stw3021.

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., Brinnel, 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 Roestel, 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., Sistine, A. V., 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. *Publications of the Astronomical Society of the Pacific*, 131(995):018002, dec 2018. doi: 10.1088/1538-3873/aaecbe. URL https://doi.org/10.1088%2F1538-3873%2Faaecbe.

Chandola, V., Banerjee, A., and Kumar, V. Anomaly detection: A survey. *ACM Comput. Surv.*, 41(3), July 2009. ISSN 0360-0300. doi: 10.1145/1541880.1541882. URL https://doi.org/10.1145/1541880.1541882.

Davis, J. and Goadrich, M. The relationship between precision-recall and roc curves. In *Proceedings of the 23rd International Conference on Machine Learning*, ICML '06, pp. 233–240, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 1595933832. doi: 10.1145/1143844.1143874. URL https://doi.org/10.1145/1143844.1143874.

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. *Astrophysical Journal*, 489(3):3582–3590, November 2019. doi: 10.1093/mnras/stz2357.

Förster, F., Cabrera-Vives, G., Castillo-Navarrete, E., Estévez, P. A., Sánchez-Sáez, P., Arredondo, J., Bauer, F. E., Carrasco-Davis, R., Catelan, M., Elorrieta, F., Eyheramendy, S., Huijse, P., Pignata, G., Reyes, E.,

Reyes, I., Rodríguez-Mancini, D., Ruz-Mieres, D., Valenzuela, C., Álvarez-Maldonado, I., Astorga, N., Borissova, J., Clocchiatti, A., De Cicco, D., Donoso-Oliva, C., Hernández-García, L., Graham, M. J., Jordán, A., Kurtev, R., Mahabal, A., Maureira, J. C., Muñoz-Arancibia, A., Molina-Ferreiro, R., Moya, A., Palma, W., Pérez-Carrasco, M., Protopapas, P., Romero, M., Sabatini-Gacitua, L., Sánchez, A., San Martín, J., Sepúlveda-Cobo, C., Vera, E., and Vergara, J. R. The Automatic Learning for the Rapid Classification of Events (ALeRCE) Alert Broker. *The Astronomical Journal*, 161(5):242, May 2021. doi: 10.3847/1538-3881/abe9bc.

Hložek, R., Ponder, K. A., Malz, A. I., Dai, M., Narayan, G., Ishida, E. E. O., Allam, T., Bahmanyar, A., Biswas, R., Galbany, L., Jha, S. W., Jones, D. O., Kessler, R., Lochner, M., Mahabal, A. A., Mandel, K. S., Martínez-Galarza, J. R., McEwen, J. D., Muthukrishna, D., Peiris, H. V., Peters, C. M., and Setzer, C. N. Results of the photometric lsst astronomical time-series classification challenge (plasticc), 2020. URL https://arxiv.org/abs/2012.12392.

Ishida, E. E. O., Kornilov, M. V., Malanchev, K. L., Pruzhinskaya, M. V., Volnova, A. A., Korolev, V. S., Mondon, F., Sreejith, S., Malancheva, A. A., and Das, S. Active anomaly detection for time-domain discoveries. *Astronomy and Astrophysics*, 650:A195, June 2021. doi: 10.1051/0004-6361/202037709.

Kessler, R., Bernstein, J. P., Cinabro, D., Dilday, B., Frieman, J. A., Jha, S., Kuhlmann, S., Miknaitis, G., Sako, M., Taylor, M., and Vanderplas, J. SNANA: A Public Software Package for Supernova Analysis. *Publications of the Astronomical Society of the Pacific*, 121(883):1028, September 2009. doi: 10.1086/605984.

Kessler, R., Narayan, G., Avelino, A., Bachelet, E., Biswas, R., Brown, P. J., Chernoff, D. F., Connolly, A. J., Dai, M., Daniel, S., Stefano, R. D., Drout, M. R., Galbany, L., González-Gaitán, S., Graham, M. L., Hložek, R., Ishida, E. E. O., Guillochon, J., Jha, S. W., Jones, D. O., Mandel, K. S., Muthukrishna, D., O'Grady, A., Peters, C. M., Pierel, J. R., Ponder, K. A., Prša, A., Rodney, S., Villar, V. A., Collaboration, T. L. D. E. S., the Transient, and Collaboration), V. S. S. Models and simulations for the photometric lsst astronomical time series classification challenge (plasticc). *Publications of the Astronomical Society of the Pacific*, 131(1003):094501, jul 2019. doi: 10.1088/1538-3873/ab26f1. URL https://dx.doi.org/10.1088/1538-3873/ab26f1.

Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.

Liu, F. T., Ting, K. M., and Zhou, Z.-H. Isolation-based anomaly detection. *ACM Trans. Knowl. Discov. Data*, 6(1), March 2012. ISSN 1556-4681. doi: 10.1145/2133360.2133363. URL https://doi.org/10.1145/2133360.2133363.

Mahabal, A., Rebbapragada, U., Walters, R., Masci, F. J., Blagorodnova, N., van Roestel, J., Ye, Q.-Z., Biswas, R., Burdge, K., Chang, C.-K., Duev, D. A., Golkhou, V. Z., Miller, A. A., Nordin, J., Ward, C., Adams, S., Bellm, E. C., Branton, D., Bue, B., Cannella, C., Connolly, A., Dekany, R., Feindt, U., Hung, T., Fortson, L., Frederick, S., Fremling, C., Gezari, S., Graham, M., Groom, S., Kasliwal, M. M., Kulkarni, S., Kupfer, T., Lin, H. W., Lintott, C., Lunnan, R., Parejko, J., Prince, T. A., Riddle, R., Rusholme, B., Saunders, N., Sedaghat, N., Shupe, D. L., Singer, L. P., Soumagnac, M. T., Szkody, P., Tachibana, Y., Tirumala, K., van Velzen, S., and Wright, D. Machine learning for the zwicky transient facility. *Publications of the Astronomical Society of the Pacific*, 131(997):038002, jan 2019. doi: 10.1088/1538-3873/aaf3fa. URL https://doi.org/10.1088%2F1538-3873%2Faaf3fa.

Mainzer, A., Bauer, J., Grav, T., Masiero, J., Cutri, R. M., Dailey, J., Eisenhardt, P., McMillan, R. S., Wright, E., Walker, R., Jedicke, R., Spahr, T., Tholen, D., Alles, R., Beck, R., Brandenburg, H., Conrow, T., Evans, T., Fowler, J., Jarrett, T., Marsh, K., Masci, F., McCallon, H., Wheelock, S., Wittman, M., Wyatt, P., DeBaun, E., Elliott, G., Elsbury, D., Gautier, T., Gomillion, S., Leisawitz, D., Maleszewski, C., Micheli, M., and Wilkins, A. PRELIMINARY RESULTS FROM NEOWISE: AN ENHANCEMENT TO THEWIDE-FIELD INFRARED SURVEY EXPLORERFOR SOLAR SYSTEM SCIENCE. *The Astrophysical Journal*, 731(1):53, mar 2011. doi: 10.1088/0004-637x/731/1/53. URL https://doi.org/10.1088%2F0004-637x%2F731%2F1%2F53.

Masci, F. J., Laher, R. R., Rusholme, B., Shupe, D. L., Groom, S., Surace, J., Jackson, E., Monkewitz, S., Beck, R., Flynn, D., Terek, S., Landry, W., Hacopians, E., Desai, V., Howell, J., Brooke, T., Imel, D., Wachter, S., Ye, Q.-Z., Lin, H.-W., Cenko, S. B., Cunningham, V., Rebbapragada, U., Bue, B., Miller, A. A., Mahabal, A., Bellm, E. C., Patterson, M. T., Jurić, M., Golkhou, V. Z., Ofek, E. O., Walters, R., Graham, M., Kasliwal, M. M., Dekany, R. G., Kupfer, T., Burdge, K., Cannella, C. B., Barlow, T., Sistine, A. V., Giomi, M., Fremling, C., Blagorodnova, N., Levitan, D., Riddle, R., Smith, R. M., Helou, G., Prince, T. A., and Kulkarni, S. R. The zwicky transient facility: Data processing, products, and archive. *Publications of the Astronomical Society of the Pacific*, 131(995):018003, dec 2018.

doi: 10.1088/1538-3873/aae8ac. URL https://doi.org/10.1088%2F1538-3873%2Faae8ac.

Muthukrishna, D., Mandel, K. S., Lochner, M., Webb, S., and Narayan, G. Real-time detection of anomalies in large-scale transient surveys. *Astrophysical Journal*, 517(1):393–419, November 2022. doi: 10.1093/mnras/stac2582.

Ruff, L., Vandermeulen, R., Goernitz, N., Deecke, L., Siddiqui, S. A., Binder, A., Müller, E., and Kloft, M. Deep one-class classification. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4393–4402, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL http://proceedings.mlr.press/v80/ruff18a.html.

Ruff, L., Vandermeulen, R., Görnitz, N., Binder, A., Müller, E., and Kloft, M. Deep support vector data description for unsupervised and semi-supervised anomaly detection. 06 2019.

Rumelhart, D. E. and McClelland, J. L. *Learning Internal Representations by Error Propagation*, pp. 318–362. 1987.

Sánchez-Sáez, P., Lira, H., Martí, L., Sánchez-Pi, N., Arredondo, J., Bauer, F. E., Bayo, A., Cabrera-Vives, G., Donoso-Oliva, C., Estévez, P. A., Eyheramendy, S., Förster, F., Hernández-García, L., Arancibia, A. M. M., Pérez-Carrasco, M., Sepúlveda, M., and Vergara, J. R. Searching for Changing-state AGNs in Massive Data Sets. I. Applying Deep Learning and Anomaly-detection Techniques to Find AGNs with Anomalous Variability Behaviors. *The Astronomical Journal*, 162(5):206, November 2021a. doi: 10.3847/1538-3881/ac1426.

Sánchez-Sáez, P., Reyes, I., Valenzuela, C., Förster, F., Eyheramendy, S., Elorrieta, F., Bauer, F. E., Cabrera-Vives, G., Estévez, P. A., Catelan, M., Pignata, G., Huijse, P., De Cicco, D., Arévalo, P., Carrasco-Davis, R., Abril, J., Kurtev, R., Borissova, J., Arredondo, J., Castillo-Navarrete, E., Rodriguez, D., Ruz-Mieres, D., Moya, A., Sabatini-Gacitúa, L., Sepúlveda-Cobo, C., and Camacho-Iñiguez, E. Alert Classification for the ALeRCE Broker System: The Light Curve Classifier. *The Astronomical Journal*, 161(3):141, March 2021b. doi: 10.3847/1538-3881/abd5c1.

Schölkopf, B., Williamson, R., Smola, A., Shawe-Taylor, J., and Platt, J. Support vector method for novelty detection. In *Proceedings of the 12th International Conference on Neural Information Processing Systems*, NIPS'99, pp. 582–588, Cambridge, MA, USA, 1999. MIT Press.

Villar, V. A., Cranmer, M., Contardo, G., Ho, S., and Lin, J. Y.-Y. Anomaly Detection for Multivariate Time Series of Exotic Supernovae. In *34th Conference on Neural Information Processing Systems*, 10 2020.

Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., Ressler, M. E., Cutri, R. M., Jarrett, T., Kirkpatrick, J. D., Padgett, D., McMillan, R. S., Skrutskie, M., Stanford, S. A., Cohen, M., Walker, R. G., Mather, J. C., Leisawitz, D., Gautier, T. N., McLean, I., Benford, D., Lonsdale, C. J., Blain, A., Mendez, B., Irace, W. R., Duval, V., Liu, F., Royer, D., Heinrichsen, I., Howard, J., Shannon, M., Kendall, M., Walsh, A. L., Larsen, M., Cardon, J. G., Schick, S., Schwalm, M., Abid, M., Fabinsky, B., Naes, L., and Tsai, C.-W. THE WIDE-FIELD INFRARED SURVEY EXPLORER (WISE): MISSION DESCRIPTION AND INITIAL ON-ORBIT PERFORMANCE. *The Astronomical Journal*, 140(6):1868–1881, nov 2010. doi: 10.1088/0004-6256/140/6/1868. URL https://doi.org/10.1088%2F0004-6256%2F140%2F6%2F1868.

Xiong, L., Poczos, B., Connolly, A., and Schneider, J. Anomaly detection for astronomical data, Jun 2010. URL https://kilthub.cmu.edu/articles/journal_contribution/Anomaly_Detection_for_Astronomical_Data/6475475/1.

A. Appendix I

In this work, six anomaly detection algorithms are examined in order to compare their performances in finding outliers. The algorithms and their respective hyperparameters are:

Isolation Forest

The Isolation Forest algorithm is based on the concept of isolation trees. It randomly selects features and divides them into distinct non-overlapping regions using a randomly selected threshold criterion. The anomaly score is proportional to the number of splits required to isolate each object in the sample. Anomalous objects are expected to require fewer splits to be isolated. To address the issue of overfitting, Isolation Forest employs ensembles of isolation trees.

We utilized the implementation of Isolation Forest provided by scikit-learn. The hyperparameters were set as follows: number of trees = 100, number of samples to train each base estimator = 256, and contamination parameter = 0.1. These values were recommended in the original work by Liu et al. (2012). Since outliers were not used for training, we did not perform hyperparameter selection through cross-validation.

One-Class Support Vector Machine

The One-Class Support Vector Machine (OCSVM) is an anomaly detection method based on Support Vector Machines (SVM). OCSVM maps the data into a new feature space, where the inner product between two objects can be represented with a kernel function, such as the Gaussian kernel. It learns a hyperplane in this feature space that separates the region where most of the data lie. During testing, the anomaly score is computed by evaluating the distance of data points with respect to the learned hyperplane.

We used the One-Class SVM implementation provided by scikit-learn with the radial basis function (RBF) kernel. The hyperparameters were set to $\nu = 0.01$ and the contamination parameter $c = 0.1$, following the default values. As the anomalous samples were assumed to be unknown, we did not perform hyperparameter selection and used the default settings.

Autoencoder

Autoencoders (AE) are unsupervised neural network algorithms that aim to reconstruct the input data using a lower-dimensional representation called the latent space. They consist of an encoder function that maps the input data to the lower-dimensional representation and a decoder function that reconstructs the original data from the latent space. AE's effectiveness in anomaly detection is attributed to the assumption that anomalies are incompressible and cannot be effectively reconstructed from low-dimensional projections. The reconstruction error, computed as the mean squared error, serves as the anomaly score.

We implemented the Autoencoder using PyTorch 1.0.0. Our model architecture consisted in an encoder with 4 hidden layers with {512, 256, 128, 64} neurons, and a decoder with {64, 128, 256, 512} neurons. We use batch normalization and Leaky ReLU as activation function, except for the last layer that used a Tanh activation function. We selected the hyperparameters by considering the reconstruction error over a validation set consisting only of inliers.

Variational Autoencoder

Variational Autoencoders (VAE) extend autoencoders by incorporating a regularization term. This term enforces the latent space to follow a known distribution, typically a normal distribution. By generating multiple reconstructions for each sample and averaging their reconstruction errors, VAEs compute an anomaly score.

We implemented the Variational Autoencoder using PyTorch 1.0.0. We use the same model architecture as the autoencoder, but we add two extra layers at the end of the encoder to learn the parameters μ and σ of the normal distribution. We selected the hyperparameters using the unsupervised loss function over a validation set consisting only of inliers.

Deep Support Vector Data Description

To train the Deep Support Vector Data Description (Deep SVDD), an Autoencoder is trained until convergence. Then, the decoder is removed, and the center of the hypersphere is estimated as the average of the encoder's outputs on the training data. Subsequently, the parameters of the encoder are re-optimized using the pretrained parameters, with the objective of minimizing the distance between the encoder's output and the center of the hypersphere, along with a weight decay regularizer.

During testing, the anomaly score is defined as the distance between each data point and the center of the hypersphere.

We implemented Deep SVDD using PyTorch 1.0.0. We use the same model architecture as the encoder of the autoencoder. The hyperparameter $\lambda = 0.5 \times 10^{-6}$ was selected by measuring the unsupervised loss function over a validation set of inliers.

Multi-Class Deep SVDD

Similar to Deep SVDD, an Autoencoder is trained until convergence, and the decoder is removed. The centers of the hyperspheres are estimated as the averages of the encoder's outputs, considering only the data points of each class. The parameters of the encoder are re-optimized using a similar objective as Deep SVDD, but considering multiple hyperspheres.

During testing, the anomaly score is determined by measuring the distance between each data point and the center of its closest hypersphere.

We implemented MCDSVDD using PyTorch 1.0.0. We use the same model architecture as the encoder of the autoencoder. We set the hyperparameter $\lambda = 0.5 \times 10^{-6}$ through cross-validation of the unsupervised loss function over a validation set consisting only of inliers.