
Image-Based Multi-Survey Classification of Light Curves with a Pre-Trained Vision Transformer

Daniel Moreno-Cartagena¹ Guillermo Cabrera-Vives^{1 2 3 4} Alejandra M. Muñoz Arancibia^{3 5}
Pavlos Protopapas⁶ Francisco Förster^{7 3 5} Márcio Catelan^{8 3 9} A. Bayo¹⁰ Pablo A. Estévez^{11 3}
P. Sánchez-Sáez¹⁰ Franz E. Bauer¹² M. Pavez-Herrera⁸ L. Hernández-García^{13 3 14} Gonzalo Rojas^{1 2}

Abstract

We explore the use of SWIN TRANSFORMER V2, a pre-trained vision Transformer, for photometric classification in a multi-survey setting by leveraging light curves from the Zwicky Transient Facility (ZTF) and the Asteroid Terrestrial-impact Last Alert System (ATLAS). We evaluate different strategies for integrating data from these surveys and find that a multi-survey architecture which processes them jointly achieves the best performance. These results highlight the importance of modeling survey-specific characteristics and cross-survey interactions, and provide guidance for building scalable classifiers for future time-domain astronomy.

1. Introduction

The field of time-domain astronomy is entering a new era of data abundance and interoperability. Over the past two decades, large-scale sky surveys with diverse cadences, spectral coverage, and sky footprints have expanded the observational landscape (Catelan, 2023). Surveys such as the Zwicky Transient Facility (ZTF; Bellm et al., 2018), Asteroid Terrestrial-impact Last Alert System (ATLAS; Tonry et al., 2018), Panoramic Survey Telescope and Rapid Response System (Pan-STARRS1; Kaiser et al., 2010), *Gaia* (Gaia Collaboration et al., 2023), and Optical Gravitational Lensing Experiment (OGLE; Udalski et al., 2015), among others (Drake et al., 2012; Shappee et al., 2014; Catelan, 2023), have independently cataloged millions of celestial sources, providing a rich yet heterogeneous view of variability across the universe. In this context, photometric classification, the task of identifying and categorizing variable and transient objects based on their light curves, has become a central challenge in modern time-domain astronomy.

Historically, most machine learning approaches to photometric classification have relied on training models with light curves from a single survey (Muthukrishna et al., 2019; Villar et al., 2019; Jamal & Bloom, 2020; Sánchez-Sáez et al., 2021; Allam Jr & McEwen, 2024; Becker et al., 2025), often using data augmentation and feature extraction to improve predictions within this setting (Lochner et al., 2016; Boone, 2019; Gómez et al., 2020; Villar et al., 2020; García-Jara et al., 2022; Monsalves et al., 2024). More recently, transfer learning has been employed to adapt models trained on one survey to another (Mahabal et al., 2017; Carrasco-Davis et al., 2019; Pimentel et al., 2022; Gupta & Muthukrishna, 2025), and self-supervised learning has been used to produce survey-agnostic representations that can generalize across different surveys (Pan et al., 2022; Donoso-Oliva et al., 2023; Moreno-Cartagena et al., 2023; Cádiz-Leyton et al., 2024a;b). However, these methods still rely on models trained on light curves from individual surveys, without jointly processing data from multiple telescopes. In parallel, multimodal approaches that combine light curves with other data modalities (e.g., spectra, metadata, derived features)

¹Department of Computer Science, Universidad de Concepción, Edmundo Larenas 219, Concepción, Chile ²Center for Data and Artificial Intelligence, Universidad de Concepción, Edmundo Larenas 310, Concepción, Chile ³Millennium Institute of Astrophysics (MAS), Santiago, Chile ⁴Heidelberg Institute for Theoretical Studies, Heidelberg, Germany ⁵Center for Mathematical Modeling (CMM), Universidad de Chile, Santiago, Chile ⁶John A. Paulson School of Engineering and Applied Sciences, Harvard University, Boston, MA, USA ⁷Data and Artificial Intelligence Initiative (IDIA), Faculty of Physical and Mathematical Sciences, Universidad de Chile ⁸Instituto de Astrofísica, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, 7820436 Macul, Santiago, Chile ⁹Centro de Astro-Ingeniería, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, 7820436 Macul, Santiago, Chile ¹⁰European Southern Observatory (ESO), Karl-Schwarzschild-Strasse 2, 85748 Garching bei München, Germany ¹¹Dept. of Electrical Engineering, University of Chile, Santiago, Chile ¹²Instituto de Alta Investigación, Universidad de Tarapacá, Casilla 7D, Arica, Chile ¹³Millennium Nucleus on Transversal Research and Technology to Explore Supermassive Black Holes (TITANS), Gran Bretaña 1111, Playa Ancha, Valparaíso, Chile ¹⁴Instituto de Física y Astronomía, Facultad de Ciencias, Universidad de Valparaíso, Gran Bretaña 1111, Playa Ancha, Valparaíso, Chile. Correspondence to: Daniel Moreno-Cartagena <dmoreno2016@inf.udec.cl>.

(Cabrera-Vives et al., 2024; Rizhko & Bloom, 2024a;b; Zhang et al., 2024) have shown promise, highlighting their potential to improve predictive performance, but they still operate within single-survey settings. While several works have begun exploring multi-survey frameworks (Starr et al., 2012; Aguirre et al., 2019; Murray et al.; Engel et al., 2024), explicit methods for jointly exploiting complementary information across surveys remain an open area of research.

In this work, we build upon the straightforward approach proposed by Moreno-Cartagena et al. (2025), who fine-tuned a SWIN TRANSFORMER V2 (SWINV2), a hierarchical vision Transformer pre-trained on ImageNet-21K (Liu et al., 2022), for multi-band light curve classification. We extend this approach to the multi-survey setting by exploring fusion strategies to combine data from ZTF and ATLAS, two complementary surveys with overlapping sky coverage and distinct observational characteristics. We propose an architecture that jointly models light curves from different telescopes, using the vision Transformer as its core. To enable this, light curves are transformed into image-based representations that serve as input to the model. Our results suggest that incorporating multiple surveys can improve classification performance, that simple ensemble approaches already leverage complementary information, and that naive early fusion underperforms relative to single-survey baselines. Moreover, ATLAS contributes complementary information despite lower individual performance.

2. Methods

2.1. Multi-Band Light Curves to Images

In our approach, multi-band light curves are transformed into images used as input to a pre-trained vision Transformer. We adopt the Overlay visualization method proposed by Moreno-Cartagena et al. (2025), in which observations from different photometric bands are plotted simultaneously within a unified visual space, as illustrated in Figure 1. To ensure comparability across light curves, we apply Min-Max normalization independently to both time and flux values for each light curve, jointly considering all its bands. This normalization maps all values to the range $[0, 1]$, removing information about absolute time and brightness and encouraging the model to focus on the shape of the light curve rather than its amplitude or temporal scale to discriminate between classes. Global normalization across the dataset was explicitly avoided, as it disproportionately suppressed low-amplitude variability and negatively impacted classification performance. Additionally, flux errors are scaled using the same Min-Max factor applied to the flux, preserving relative uncertainty information. Figure 1 shows three example input images per dataset, including lines, points, and error bars (see Section 3.2). Each band is mapped to a location near a primary RGB axis, ensuring bal-

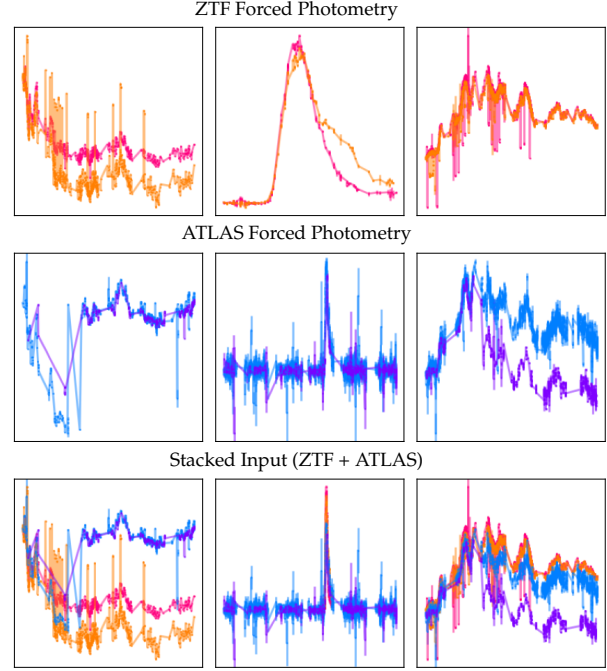


Figure 1. Model input images constructed from forced photometry light curves. Rows show ZTF (g and r bands), ATLAS (c and o bands), and the stacked input with all bands. Images represent flux difference over time, with axes omitted since time and flux are not explicitly used. From left to right, columns show ZTF18aajpcq (Blazar), ZTF21abfnfdw (SNIa), and ZTF18aampkcf (AGN).

anced color representation and avoiding visual bias across bands.

2.2. Vision Transformer Model

Following the methodology introduced by Moreno-Cartagena et al. (2025), we adopt the pre-trained SWINV2 architecture, adapted for the classification of light curves, as the core model. SWINV2 is a hierarchical vision Transformer originally pre-trained on ImageNet-21K for natural image recognition tasks (Liu et al., 2021; Carion et al., 2020; Wang et al., 2021; Chen et al., 2021). Basically, the model first partitions each input image into fixed-size patches, which are embedded and processed through a series of self-attention blocks with a hierarchical structure. This enables the model to capture both local patterns and global dependencies across the light curve image. After processing the image through the SWINV2 backbone, an adaptive average pooling layer computes the mean across all tokens, producing a single pooled embedding that is used for classification. Further architectural details and an illustration of the model are provided in Appendix A. To assess the effectiveness of this vision-based approach, we compare it with the Astronomical Transformer for time series And Tabular data (ATAT; Cabrera-Vives et al., 2024), a specialized

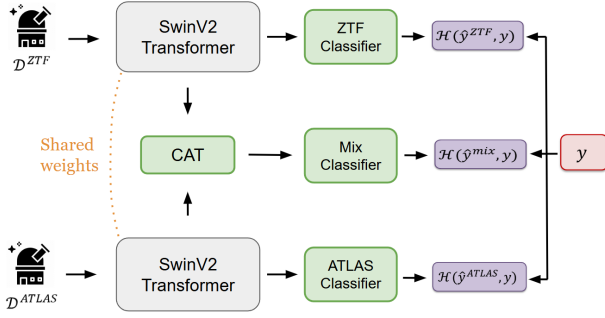


Figure 2. Multi-survey architecture.

Transformer architecture tailored for sequential modeling of multi-band astronomical time series.

2.3. Multi-Survey Methodology

We explore multiple strategies for combining ZTF and ATLAS data to improve classification performance. As a baseline, we first train separate single-survey SWIN V2 models for ZTF and ATLAS. We then perform prediction-level fusion by averaging the predicted probabilities of these two models to obtain a combined prediction. Additionally, we perform embedding-level fusion by extracting embeddings from the single-survey models and training a Balanced Random Forest (BRF) classifier on the concatenated embeddings. We also experiment with input-level fusion by creating a new dataset in which light curves from both surveys are represented as multi-band images in which bands g , r , o , and c are concatenated spatially into a single image, as shown in Figure 1 (stacked input). Finally, we propose a multi-survey architecture that simultaneously processes photometric data from both surveys. Figure 2 illustrates the overall architecture of this approach. Specifically, light curves from ZTF (\mathcal{D}^{ZTF}) and ATLAS ($\mathcal{D}^{\text{ATLAS}}$) are treated as independent input streams and processed by two parallel SWIN V2 Transformer branches with shared weights. The output of each branch is connected to a dedicated linear classifier, ZTF Classifier and ATLAS Classifier, that produces survey-specific predictions \hat{y}^{ZTF} and \hat{y}^{ATLAS} , each optimized using a cross-entropy loss $\mathcal{H}(\hat{y}, y)$. In parallel, the pooled embeddings from both streams are concatenated (CAT) and passed through a Mix Classifier, and we experiment with both a Linear Layer and a Multi-Layer Perceptron (MLP) implementation for this classifier, to generate a joint prediction \hat{y}^{mix} , also optimized via cross-entropy. The total loss is computed as $\mathcal{L}_{\text{total}} = \mathcal{H}(\hat{y}^{\text{ZTF}}, y) + \mathcal{H}(\hat{y}^{\text{ATLAS}}, y) + \mathcal{H}(\hat{y}^{\text{mix}}, y)$. This design enables the model to jointly learn survey-specific patterns and cross-survey interactions, while maintaining flexibility to output individual and joint predictions.

3. Experiments

3.1. Data Description

We use forced-photometry light curves from ZTF and ATLAS. Forced photometry retrieves fluxes at fixed sky positions, enabling the detection of faint signals. ZTF objects were selected from the public alert stream and processed using the Automatic Learning for the Rapid Classification of Events system (ALeRCE; Förster et al., 2021). Historical fluxes were obtained via the ZTF Forced Photometry Service (Masci et al., 2023), using each object’s RA and Dec, with standard quality filtering. Forced photometry complements the alert stream with more complete flux information. This type of data has been included in ZTF alerts only since October 2023. Labels were assigned through crossmatch with external catalogs and expert validation. Preprocessing removed epochs with negative flux uncertainties or bad-quality flags, selected forced photometry observations from up to 30 days before the first alert through to the last alert, discarded objects with fewer than eight detections, and filtered spurious events. We used calibrated differential fluxes in microjanskys (μJy). A forthcoming ALeRCE publication will describe the dataset in detail. After preprocessing, the ZTF dataset included 39,483 light curves in g and r bands spanning 21 classes, with a median cadence of 1.13 days and a standard deviation of 6.5 days, considering both bands. Data were split into training, validation, and test sets with class-stratified sampling: 20% for testing, and the remaining 80% split into five training folds.

For the multi-survey approach, we incorporated ATLAS forced-photometry data. Using the RA and Dec of each ZTF object, we queried the ATLAS service (Tonry et al., 2018; Heinze et al., 2018; Smith et al., 2020; Shingles et al., 2021) to retrieve historical fluxes at the same positions. ATLAS performs PSF-fit forced photometry on difference images using `tphot` (Tonry, 2010). To ensure quality, we retained only measurements with $0.5 < \chi^2/N < 3$ and sky magnitudes above 18 in the o band and 18.5 in the c band, following the criteria from Silva-Farfán et al. (2024). Fluxes are reported in μJy in the AB system, with the filter (c or o) indicated for each observation. The ATLAS dataset comprised 39,464 light curves, with the same classes and data split as the ZTF dataset. It has a median cadence of 1.52 days with a standard deviation of 11.03 days, considering both bands. The number of objects per class and the class taxonomy are shown in Table 1. The table also presents the division into transient, stochastic, and periodic sources.

3.2. Training Details

In our experiments, light curves were converted to images with `Matplotlib`, transformed into RGB tensors, and adjusted to minimize whitespace. We used the SWIN V2 model pre-trained on ImageNet-21K, available via Hugging

Table 1. Number of objects per class in ZTF, grouped as *Transient*, *Stochastic*, and *Periodic*, as defined by the ALerCE Team.

Transient	Stochastic	Periodic
SNIa (2486)	Microlensing (29)	LPV (2931)
SESN (565)	QSO (2923)	EA (2909)
SNII (1393)	AGN (2898)	EB/EW (2921)
SNIIIn (267)	Blazar (1393)	Periodic-Other (1269)
SLSN (185)	YSO (2773)	RSCVn (2452)
TDE (81)	CV/Nova (1640)	CEP (2891)
-		RRLab (2931)
-		RRLc (2936)
-		DSCT (1610)

Face¹. The model operates on $256 \times 256 \times 3$ images, divides them into 4×4 patches, uses an embedding dimension of $C = 96$, and performs self-attention within 16×16 windows. Image processing used `AutoImageProcessor`, including resizing to 256×256 , rescaling to $[0, 1]$, and normalization with `ImageNet-21K` statistics. Training was conducted on an NVIDIA RTX A5000 GPU.

We used a marker size of 1.0, a line width of 2.0, flux errors, and a learning rate of $5 \cdot 10^{-6}$ as hyperparameters for image generation, following [Moreno-Cartagena et al. \(2025\)](#), who found this configuration to be optimal for the two bands in the public Catalog of Variable Stars from the Massive Compact Halo Object (MACHO) survey. To mitigate class imbalance, we employed a weighted sampling strategy² based on class frequencies, promoting a more balanced representation and better generalization. The model was optimized with Adam and early stopping (patience: 5 epochs) using validation F1-score.

4. Results

Table 2 summarizes the classification performance on the test set, in terms of F1-score, for the different strategies explored for combining ZTF and ATLAS data. As single-survey baselines, we compare the ATAT and SWINV2 models trained independently on ZTF and ATLAS. ATAT is a specialized Transformer architecture designed to process multi-band light curves jointly with tabular data. In this work, however, we focus exclusively on light curves and therefore do not use metadata or derived features from the light curves, to ensure a fair comparison with the vision Transformer. This configuration is referred to as ATAT (LC). We prepared the input data for ATAT using sliding windows of 200 observations per light curve and per band, as the quadratic complexity of the attention mechanism limits the feasible input length. The median number of observations

Table 2. Comparison of ATAT and SWINV2-based strategies for photometric classification using ZTF and ATLAS data.

MODEL	DESCRIPTION	F1-SCORE
1	ATAT (LC) (ZTF ONLY)	57.0 ± 1.4
2	ATAT (LC) (ATLAS ONLY)	49.6 ± 1.9
3	SWINV2 (ZTF ONLY)	64.9 ± 1.6
4	SWINV2 (ATLAS ONLY)	54.0 ± 1.3
3 + 4	AVERAGED PROBABILITIES	68.9 ± 1.3
3 + 4	BRF ON EMBEDDINGS	68.7 ± 0.9
5	SWINV2 (STACKED INPUT)	59.0 ± 1.3
6	MS-SWINV2-LINEAR	69.9 ± 1.7
7	MS-SWINV2-MLP	69.8 ± 1.5

per light curve is 952.0 for ZTF and 1493.0 for ATLAS. Final class probabilities were obtained by averaging the predicted probability vectors across all windows for each light curve. Table 2 shows that ATAT achieved F1-scores of 57.0% (ZTF, Model 1) and 49.6% (ATLAS, Model 2), while SWINV2 reached higher F1-scores of 64.9% (ZTF, Model 3) and 54.0% (ATLAS, Model 4). These results indicate that distinguishing classes on ZTF is easier for both models than on ATLAS, likely due to differences in filter sets, cadence, and photometric noise. They also highlight the potential of vision Transformers for multi-band light curve classification compared to the specialized ATAT architecture. We adopt SWINV2 as the baseline for subsequent fusion experiments, as it achieved the best overall performance.

When combining the two single-survey models, both prediction-level fusion (averaged probabilities) and embedding-level fusion (BRF on embeddings) achieved higher F1-scores than the individual baselines, reaching 68.9% and 68.7%, respectively. A hyperparameter search is reported in Table 3 of Appendix B for the BRF classifier. These results indicate that even simple post-hoc fusion strategies can leverage complementary information between surveys. For input-level fusion, where light curves from both surveys are represented as stacked multi-band images (Figure 1, stacked input) and processed by SWINV2 (Model 5), reached only 59.0%, below the single-survey ZTF baseline. This suggests that naive early fusion may introduce additional complexity or noise when combining heterogeneous light curves into a single input.

Finally, the proposed multi-survey SWINV2 architectures, which process each survey as an independent input stream with shared weights, achieved the highest performance among the strategies evaluated. The Linear variant (Model 6) reached 69.9%, and the MLP variant (Model 7) 69.8% F1-score. Although ATLAS-only models perform worse than ZTF-only models, incorporating ATLAS data in the multi-survey framework provides additional signals that improve overall classification. In contrast to post-hoc fusion (e.g., 68.9% with averaged probabilities), the multi-survey

¹Pre-trained `Hugging Face` model used.

²`PyTorch`'s `WeightedRandomSampler`.

approach captures cross-survey interactions during training. This is reflected in the confusion matrices shown in Figure 4 (Appendix C). Moreover, it is well-suited for scaling to additional surveys, simplifies deployment with a unified model, and offers a flexible foundation for incorporating other modalities, such as metadata, contextual features, or spectra.

These results suggest the benefit of modeling survey-specific representations while jointly capturing cross-survey interactions within a unified architecture. Using a pre-trained vision Transformer also avoids the complexity of designing specialized representations for multi-band light curve data, offering a simple alternative to classical sequence models.

5. Conclusion

In this work, we presented a systematic exploration of fusion strategies for photometric classification using multi-survey light curves from ZTF and ATLAS. Our results indicate that multi-survey architectures, which process each survey as an independent input, obtained the highest classification performance among the approaches evaluated, compared to single-survey baselines and simpler fusion methods. These observations suggest that incorporating information from additional telescopes can be beneficial for classification, even when the data are noisier or of lower individual quality. The results also point to the importance of explicitly modeling survey-specific characteristics when integrating heterogeneous time-domain data. As multi-survey and multi-instrument data become increasingly common, such architectures may provide a useful foundation for scalable and adaptable photometric classifiers. Future work will extend this approach to incorporate metadata and additional features (Cabrera-Vives et al., 2024), and to explore its use in real-time alert processing.

Acknowledgements

The authors acknowledge support from the National Agency for Research and Development (ANID) grants: FONDECYT Regular 1231877 (DMC, GCV); FONDECYT Regular 1241005 (FEB); FONDECYT Regular 1231637 (MC); FONDECYT Iniciación 11241477 (LHG); FONDECYT 1220829 (PAE); Millennium Science Initiative Program ICN12.009 (GCV, AMMA); Millennium Science Initiative Program AIM23.0001, awarded to the Millennium Institute of Astrophysics (FEB, FF, LHG, MC, PAE, AMMA); Millennium Science Initiative Program NCN2023.002 (LHG). This work was also supported by the BASAL Center of Mathematical Modeling Grant FB210005 (AMMA, FF), ANID-Chile BASAL CATA FB210003 (FEB), ANID’s BASAL project FB210003 (MC), and the Quimal infrastructure funds QUIMAL190012 and QUIMAL240008 (FF).

References

- Aguirre, C., Pichara, K., and Becker, I. Deep multi-survey classification of variable stars. *Monthly Notices of the Royal Astronomical Society*, 482(4):5078–5092, 2019.
- Allam Jr, T. and McEwen, J. D. Paying attention to astronomical transients: introducing the time-series transformer for photometric classification. *RAS Techniques and Instruments*, 3(1):209–223, 2024.
- Becker, I., Protopapas, P., Catelan, M., and Pichara, K. Multiband embeddings of light curves. *A&A*, 694:A183, 2025. doi: 10.1051/0004-6361/202347461. URL <https://doi.org/10.1051/0004-6361/202347461>.
- 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., et al. The zwicky transient facility: system overview, performance, and first results. *Publications of the Astronomical Society of the Pacific*, 131(995): 018002, 2018.
- Boone, K. Avocado: Photometric classification of astronomical transients with gaussian process augmentation. *The Astronomical Journal*, 158(6):257, 2019.
- Cabrera-Vives, G., Moreno-Cartagena, D., Astorga, N., Reyes-Jainaga, I., Förster, F., Huijse, P., Arredondo, J., Arancibia, A. M., Bayo, A., Catelan, M., et al. Atat: Astronomical transformer for time series and tabular data. *Astronomy & Astrophysics*, 689:A289, 2024.
- Cádiz-Leyton, M., Cabrera-Vives, G., Protopapas, P., Moreno-Cartagena, D., and Donoso-Oliva, C. Transformer-based astronomical time series model with uncertainty estimation for detecting misclassified instances. In *LatinX in AI workshop, 41 st International Conference on Machine Learning (ICML), PMLR 235*, Vienna, Austria, 2024a. URL <https://arxiv.org/abs/2411.01363>.
- Cádiz-Leyton, M., Cabrera-Vives, G., Protopapas, P., Moreno-Cartagena, D., Donoso-Oliva, C., and Becker, I. Uncertainty estimation for time series classification: Exploring predictive uncertainty in transformer-based models for variable stars. *Astronomy & Astrophysics*, 2024b. Accepted for publication.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., and Zagoruyko, S. End-to-end object detection with transformers. In *European conference on computer vision*, pp. 213–229. Springer, 2020.
- Carrasco-Davis, R., Cabrera-Vives, G., Förster, F., Estévez, P. A., Huijse, P., Protopapas, P., Reyes, I., Martínez-Palomera, J., and Donoso, C. Deep learning for image

- sequence classification of astronomical events. *Publications of the Astronomical Society of the Pacific*, 131 (1004):108006, 2019.
- Catelan, M. Stellar Variability in Ground-Based Photometric Surveys: AN Overview. In *Memorie della Societa Astronomica Italiana*, volume 94, pp. 56, December 2023. doi: 10.36116/MEMSAIT_94N4.2023.56.
- Chen, X., Yan, B., Zhu, J., Wang, D., Yang, X., and Lu, H. Transformer tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8126–8135, 2021.
- Donoso-Oliva, C., Becker, I., Protopapas, P., Cabrera-Vives, G., Vishnu, M., and Vardhan, H. Astromer-a transformer-based embedding for the representation of light curves. *Astronomy & Astrophysics*, 670:A54, 2023.
- Drake, A. J., Djorgovski, S. G., Mahabal, A., Prieto, J. L., Beshore, E., Graham, M. J., Catalan, M., Larson, S., Christensen, E., and Donalek, C. The catalina real-time transient survey. In *Proceedings of the International Astronomical Union, Symposium S285: New Horizons in Time-Domain Astronomy*, volume 7, pp. 306–308. Cambridge University Press, 2012. doi: 10.1017/S1743921312000889.
- Engel, A., Narayan, G., and Byler, N. Preliminary report on mantis shrimp: a multi-survey computer vision photometric redshift model. *arXiv preprint arXiv:2402.03535*, 2024. URL <https://arxiv.org/abs/2402.03535>. Submitted to AI4Differential Equations in Science Workshop, ICLR 2024.
- Förster, F., Cabrera-Vives, G., Castillo-Navarrete, E., Estévez, P., Sánchez-Sáez, P., Arredondo, J., Bauer, F., Carrasco-Davis, R., Catelan, M., Elorrieta, F., et al. The automatic learning for the rapid classification of events (alerce) alert broker. *The Astronomical Journal*, 161(5): 242, 2021.
- Gaia Collaboration, Vallenari, A., Brown, A. G. A., Prusti, T., de Bruijne, J. H. J., and et al. Gaia data release 3: Summary of the content and survey properties. *Astronomy & Astrophysics*, 674:A1, 2023. doi: 10.1051/0004-6361/202243940.
- García-Jara, G., Protopapas, P., and Estévez, P. A. Improving astronomical time-series classification via data augmentation with generative adversarial networks. *The Astrophysical Journal*, 935(1):23, 2022.
- Gómez, C., Neira, M., Hernández Hoyos, M., Arbeláez, P., and Forero-Romero, J. E. Classifying image sequences of astronomical transients with deep neural networks. *Monthly Notices of the Royal Astronomical Society*, 499 (3):3130–3138, 2020.
- Gupta, R. and Muthukrishna, D. Transfer learning for transient classification: From simulations to real data and ztf to lsst. *arXiv preprint arXiv:2502.18558*, 2025. Submitted to MNRAS.
- Heinze, A., Tonry, J. L., Denneau, L., Flewelling, H., Stalder, B., Rest, A., Smith, K. W., Smartt, S. J., and Weiland, H. A first catalog of variable stars measured by the asteroid terrestrial-impact last alert system (atlas). *The Astronomical Journal*, 156(5):241, 2018.
- Jamal, S. and Bloom, J. S. On neural architectures for astronomical time-series classification with application to variable stars. *The Astrophysical Journal Supplement Series*, 250(2):30, 2020.
- Kaiser, N., Burgett, W., Chambers, K., Denneau, L., Heasley, J., Jedicke, R., Magnier, E., Morgan, J., Onaka, P., and Tonry, J. The pan-starrs wide-field optical/nir imaging survey. In *Ground-based and airborne telescopes III*, volume 7733, pp. 159–172. SPIE, 2010.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 10012–10022, 2021.
- Liu, Z., Hu, H., Lin, Y., Yao, Z., Xie, Z., Wei, Y., Ning, J., Cao, Y., Zhang, Z., Dong, L., et al. Swin transformer v2: Scaling up capacity and resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12009–12019, 2022.
- Lochner, M., McEwen, J. D., Peiris, H. V., Lahav, O., and Winter, M. K. Photometric supernova classification with machine learning. *The Astrophysical Journal Supplement Series*, 225(2):31, 2016.
- Mahabal, A., Sheth, K., Gieseke, F., Pai, A., Djorgovski, S. G., Drake, A. J., and Graham, M. J. Deep-learned classification of light curves. In *2017 IEEE symposium series on computational intelligence (SSCI)*, pp. 1–8. IEEE, 2017.
- Masci, F. J., Laher, R. R., Rusholme, B., Shupe, D., Paladini, R., Groom, S., Wold, A., Miller, A. A., and Drake, A. A new forced photometry service for the zwicky transient facility. *arXiv preprint arXiv:2305.16279*, 2023.
- Monsalves, N., Jaque Arancibia, M., Bayo, A., Sánchez-Sáez, P., Angeloni, R., Damke, G., and Segura Van de Perre, J. Application of convolutional neural networks to time domain astrophysics. 2d image analysis of ogle light curves. *Astronomy & Astrophysics*, 691:A106, 2024.

- Moreno-Cartagena, D., Cabrera-Vives, G., Protopapas, P., Donoso-Oliva, C., Pérez-Carrasco, M., and Cádiz-Leyton, M. Positional encodings for light curve transformers: Playing with positions and attention. In *Machine Learning for Astrophysics Workshop, 40th International Conference on Machine Learning (ICML), PMLR 202*, Honolulu, Hawaii, USA, 2023. URL <https://ml4astro.github.io/icml2023/assets/55.pdf>.
- Moreno-Cartagena, D., Protopapas, P., Cabrera-Vives, G., Cádiz-Leyton, M., Becker, I., and Donoso-Oliva, C. Leveraging pre-trained visual transformers for multi-band photometric light curve classification. *arXiv preprint arXiv:2502.20479*, 2025. Submitted to Astronomy & Astrophysics.
- Murray, C., Fabbro, S., Herzberg, N., and Yi, K. M. Unsupervised learning of deep features through best-fits for observational cosmology.
- Muthukrishna, D., Narayan, G., Mandel, K. S., Biswas, R., and Hložek, R. Rapid: early classification of explosive transients using deep learning. *Publications of the Astronomical Society of the Pacific*, 131(1005):118002, 2019.
- Pan, J., Ting, Y.-S., and Yu, J. Astroconformer: Inferring surface gravity of stars from stellar light curves with transformer. In *Machine Learning for Astrophysics Workshop, 39th International Conference on Machine Learning (ICML), PMLR 162*, Baltimore, Maryland, USA, 2022.
- Pimentel, Ó., Estévez, P. A., and Förster, F. Deep attention-based supernovae classification of multiband light curves. *The Astronomical Journal*, 165(1):18, 2022.
- Rizhko, M. and Bloom, J. S. Self-supervised multimodal model for astronomy. In *Neurips 2024 Workshop Foundation Models for Science: Progress, Opportunities, and Challenges*, 2024a.
- Rizhko, M. and Bloom, J. S. Astrom³: A self-supervised multimodal model for astronomy. *arXiv preprint arXiv:2411.08842*, 2024b. Submitted to MNRAS.
- Sánchez-Sáez, P., Reyes, I., Valenzuela, C., Förster, F., Eyheramendy, S., Elorrieta, F., Bauer, F., Cabrera-Vives, G., Estévez, P., Catelan, M., et al. Alert classification for the alerce broker system: The light curve classifier. *The Astronomical Journal*, 161(3):141, 2021.
- Shappee, B. J., Prieto, J., Grupe, D., Kochanek, C., Stanek, K., De Rosa, G., Mathur, S., Zu, Y., Peterson, B., Pogge, R., et al. The man behind the curtain: X-rays drive the uv through nir variability in the 2013 active galactic nucleus outburst in ngc 2617. *The Astrophysical Journal*, 788(1):48, 2014.
- Shingles, L., Smith, K., Young, D., Smartt, S., Tonry, J., Denneau, L., Heinze, A., Weiland, H., Flewelling, H., Stalder, B., et al. Release of the atlas forced photometry server for public use. *Transient Name Server AstroNote 2021-7*, 7:1–7, 2021.
- Silva-Farfán, J., Förster, F., Moriya, T. J., Hernández-García, L., Arancibia, A. M., Sánchez-Sáez, P., Anderson, J. P., Tonry, J. L., and Clocchiatti, A. Physical properties of type ii supernovae inferred from ztf and atlas photometric data. *The Astrophysical Journal*, 969(1):57, 2024.
- Smith, K., Smartt, S., Young, D., Tonry, J., Denneau, L., Flewelling, H., Heinze, A., Weiland, H., Stalder, B., Rest, A., et al. Design and operation of the atlas transient science server. *Publications of the Astronomical Society of the Pacific*, 132(1014):085002, 2020.
- Starr, D. L., Richards, J. W., Brink, H., Miller, A. A., Bloom, J. S., Butler, N. R., James, J. B., and Long, J. P. ALLStars: Overcoming Multi-Survey Selection Bias using Crowd-Sourced Active Learning. In Ballester, P., Egret, D., and Lorente, N. P. F. (eds.), *Astronomical Data Analysis Software and Systems XXI*, volume 461 of *Astronomical Society of the Pacific Conference Series*, pp. 581, September 2012.
- Tonry, J., Denneau, L., Heinze, A., Stalder, B., Smith, K., Smartt, S., Stubbs, C., Weiland, H., and Rest, A. Atlas: a high-cadence all-sky survey system. *Publications of the Astronomical Society of the Pacific*, 130(988):064505, 2018.
- Tonry, J. L. An early warning system for asteroid impact. *Publications of the Astronomical Society of the Pacific*, 123(899):58, 2010.
- Udalski, A., Szymański, M. K., and Szymański, G. Ogleiv: Fourth phase of the optical gravitational lensing experiment. *Acta Astronomica*, 65(1):1–38, 2015. URL https://acta.astro.uw.edu.pl/Vol65/n1/pdf/pap_65_1_1.pdf.
- Villar, V., Berger, E., Miller, G., Chornock, R., Rest, A., Jones, D., Drout, M., Foley, R., Kirshner, R., Lunnan, R., et al. Supernova photometric classification pipelines trained on spectroscopically classified supernovae from the pan-starrs1 medium-deep survey. *The Astrophysical Journal*, 884(1):83, 2019.
- 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., et al. Superraenn: a semisupervised supernova photometric classification pipeline trained on pan-starrs1 medium-deep survey supernovae. *The Astrophysical Journal*, 905(2):94, 2020.

Wang, H., Zhu, Y., Adam, H., Yuille, A., and Chen, L.-C. Max-deeplab: End-to-end panoptic segmentation with mask transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5463–5474, 2021.

Zhang, G., Helfer, T., Gagliano, A. T., Mishra-Sharma, S., and Villar, V. A. Maven: a multimodal foundation model for supernova science. *Machine Learning: Science and Technology*, 5(4):045069, 2024.

A. SwinV2 Architecture Details

The processing pipeline, illustrated in Figure 3, proceeds as follows. The SWINV2 architecture first partitions each input image (Section 2.1) into fixed-size patches (Patch Partition). Each patch is then flattened and projected into a higher-dimensional space via a linear transformation (Linear Embedding), producing tokens that serve as input representations for the first SWIN TRANSFORMER Block, where Multi-head Self-Attention (MSA) is applied. These tokens first pass through Window-based MSA (W-MSA), where attention is restricted to non-overlapping local windows, reducing computational complexity from quadratic to linear with respect to the number of tokens, while capturing local dependencies. To enable cross-window interactions, the model alternates W-MSA with Shifted Window MSA (SW-MSA), which shifts the window partitions so that tokens at the edges of one window are included in neighboring windows, enabling the capture of global context. SWINV2 builds hierarchical feature representations by progressively applying patch merging layers between SWIN TRANSFORMER Blocks (Patch Merging), reducing spatial resolution while increasing embedding dimensionality and enabling multi-scale feature extraction. Within this framework, scaled cosine attention normalizes attention scores, improving training stability, while residual post-normalization mitigates the accumulation of large activations in deeper layers. Additionally, a log-spaced continuous position bias emphasizes local token interactions while preserving the ability to model long-range dependencies, enhancing generalization across varying image resolutions. Together, these components enable SWINV2 to efficiently and scalably process high-resolution light curve images. After the hierarchical stages, the output of the last SWIN TRANSFORMER Block is passed through an adaptive average pooling layer, which computes the mean across all tokens, producing a single pooled output vector that serves as the final embedding of the image. Finally, fine-tuning adapts the model to capture class-discriminative patterns for light curve classification. Further details of this approach can be found in Liu et al. (2021) and Moreno-Cartagena et al. (2025).

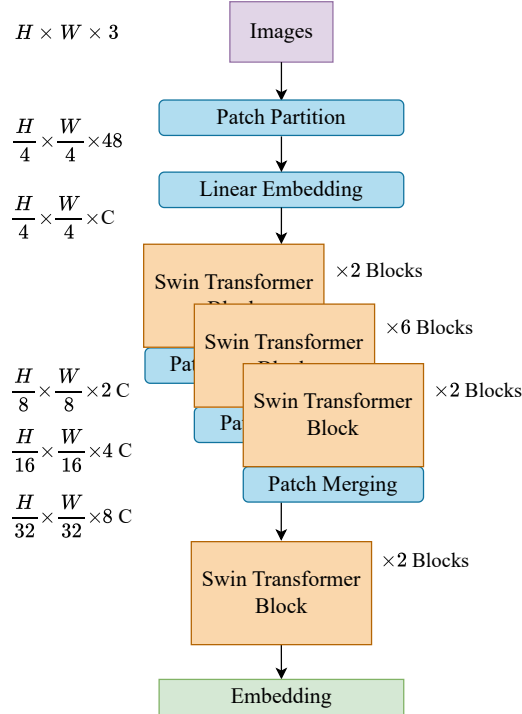


Figure 3. SWINV2 architecture. The rounded light blue rectangles highlight the components where the model changes the dimensionality of the information, while the regular rectangles indicate where it remains fixed. The figure is from Moreno-Cartagena et al. (2025).

B. Hyperparameter Tuning for the BRF Classifier

The BRF classifier hyperparameters were optimized using F1-score on the validation set across five stratified folds, matching those used to train, validate, and test the SWINV2 embeddings. A total of 120 combinations were evaluated via grid search over the following parameter grid: `n_estimators` = {100, 200, 300, 500}, `max_depth` = {None, 10, 20, 30, 50}, `max_features` = {'sqrt', 'log2'}, and `min_samples_split` = {2, 5, 10}. Table 3 reports the top 20 configurations,

with F1-score summarized as mean \pm standard deviation across the five validation folds.

Table 3. Top 20 hyperparameter configurations for the BRF classifier on the validation set.

N_ESTIMATORS	MAX FEATURES	MAX DEPTH	MIN SAMPLES SPLIT	F1-SCORE
500	LOG2	30	5	65.7 \pm 0.5
500	LOG2	NONE	5	65.7 \pm 0.5
500	LOG2	50	5	65.7 \pm 0.5
500	LOG2	10	2	65.7 \pm 1.2
500	LOG2	20	5	65.7 \pm 0.5
300	LOG2	20	2	65.7 \pm 1.3
300	LOG2	NONE	2	65.6 \pm 1.5
300	LOG2	30	2	65.6 \pm 1.5
300	LOG2	50	2	65.6 \pm 1.5
500	LOG2	NONE	2	65.6 \pm 1.2
500	LOG2	50	2	65.6 \pm 1.2
500	LOG2	30	2	65.6 \pm 1.2
500	SQRT	20	2	65.5 \pm 0.2
500	LOG2	20	2	65.5 \pm 1.2
300	SQRT	20	2	65.5 \pm 0.3
500	SQRT	10	5	65.5 \pm 0.6
500	SQRT	30	2	65.4 \pm 0.4
500	SQRT	50	2	65.4 \pm 0.4
500	SQRT	NONE	2	65.4 \pm 0.4
500	SQRT	20	5	65.4 \pm 0.5

C. Confusion Matrices of Selected Models

This appendix provides the confusion matrices for the key models reported in Section 4. Specifically, we include the ZTF-only and ATLAS-only single-survey baselines (Model 3 and Model 4), the post-hoc fusion with averaged probabilities (Model 3+4), and the MS-SWINV2-Linear (Model 6). These visualizations offer additional insight into class-wise performance across the different fusion strategies. The confusion matrices reveal differences between the single-survey baselines. The SWINV2 model trained on ZTF (Figure 4a) achieves higher recall than the ATLAS-only baseline (Figure 4b) in 19 out of 21 classes, with the ATLAS-only baseline showing higher recall only for RSCVn and DSCT. The ATLAS-only baseline exhibits particularly low recall for several rare or challenging classes, including SESN (which contains SNIbc and SNIib), SLSN, TDE, and Microlensing, with recall values below 25%. For the remaining ATLAS-only classes, recall reaches 50% or higher. In contrast, the ZTF-only baseline maintains recall above 25% for all classes, with the lowest recall observed for SNIIn (36%) and SLSN (26%). All other ZTF-only classes achieve recall of 50% or higher.

Averaging the predicted probabilities from the single-survey baselines (Figure 4c) results in recall increases in 17 out of 21 classes compared to the ZTF-only baseline. Recall decreases were observed in four classes where the ATLAS-only baseline recall was already below 25%: SESN, SLSN, TDE, and Microlensing. For these, the averaged-probabilities fusion yielded lower recall than the ZTF-only baseline. The EB/EW class maintained the same recall as in the ZTF-only baseline. Overall, the averaged-probabilities fusion appears to leverage complementary information between surveys, producing more balanced class-wise performance compared to using only ZTF data.

The MS-SWINV2-Linear model (Figure 4d) further increases class-wise recall in 19 out of 21 classes compared to the ZTF-only baseline, suggesting that the model can effectively leverage complementary information between surveys. The four most challenging classes, SESN, SLSN, TDE, and Microlensing, exhibited low recall in the ATLAS-only baseline (19%, 25%, 8%, and 7%, respectively), which limits the contribution of complementary information in these cases and makes improvements over the ZTF-only baseline more difficult. While the averaged-probabilities fusion did not improve recall in any of these classes, MS-SWINV2-Linear achieved a +3 percentage point improvement over the ZTF-only baseline in two of them: SESN and SLSN, while no improvement was observed for TDE and Microlensing, which are also the minority classes in the dataset. This suggests that joint multi-survey learning can extract useful features even when one survey contributes limited information.

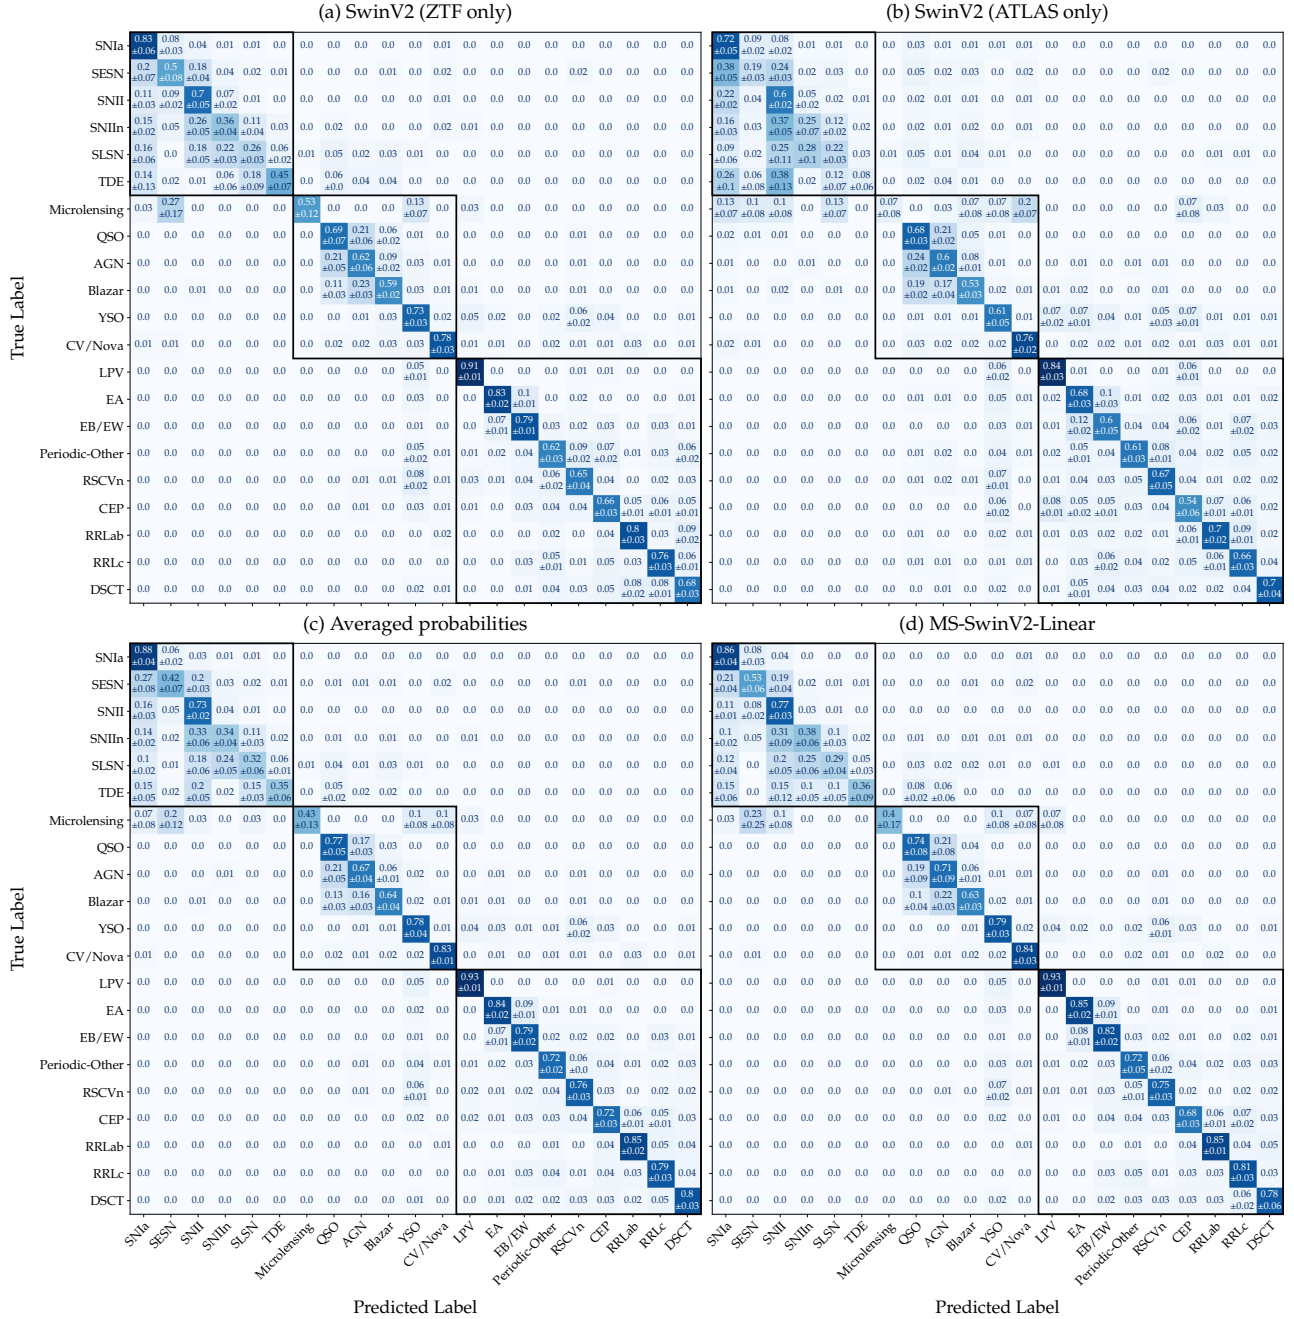


Figure 4. Confusion matrices for the different strategies explored to combine ZTF and ATLAS data for photometric classification. **(a)** Single-survey SWINV2 model trained on ZTF light curves. **(b)** Single-survey SWINV2 model trained on ATLAS light curves. **(c)** Prediction-level fusion: combined predictions obtained by averaging the output probabilities of the single-survey ZTF and ATLAS models. **(d)** Multi-survey SWINV2 architecture: jointly trained on both ZTF and ATLAS light curves, with separate input streams and a final Linear Layer classifier. Black lines delineate the three main groups of objects defined by the ALerCE Team: Transient, Stochastic, and Periodic groups, from left to right. Standard deviations across cross-validation folds are displayed for values greater than or equal to 0.05.