

Automated all-sky detection of γ Doradus / δ Scuti hybrids in TESS data from positive unlabelled (PU) learning

Mykyta Kliapets^{1,2}, Pablo Huijse¹, Andrew Tkachenko¹, Alex Kemp¹, Dario J. Fritzewski¹, Daniel Hey³, and Conny Aerts^{1,4,5}

¹ Institute of Astronomy, KU Leuven, Celestijnenlaan 200D, bus 2401, 3001 Leuven, Belgium
e-mail: mykyta.kliapets@kuleuven.be

² Kavli Scholar funded by The Kavli Foundation, 5715 Mesmer Avenue, Los Angeles, CA 90230, USA

³ Institute for Astronomy, University of Hawai'i, Honolulu, HI 96822, USA

⁴ Department of Astrophysics, IMAPP, Radboud University Nijmegen, PO Box 9010, 6500 GL, Nijmegen, The Netherlands

⁵ Max Planck Institute for Astronomy, Koenigstuhl 17, 69117, Heidelberg, Germany

Received TBA; accepted TBA

ABSTRACT

Context. The Transiting Exoplanet Survey Satellite (TESS) mission has observed hundreds of millions of stars, substantially contributing to the available pool of high-precision photometric space data. Among them are the relatively rare γ Doradus / δ Scuti (γ Dor / δ Sct) hybrid pulsators, which have been previously studied using *Kepler* data. These stars are perfect laboratories to probe both inner and outer interior stellar layers thanks to them exhibiting both pressure and gravity modes.

Aims. We seek to classify an all-sky sample of AF stars observed by TESS to find previously undiscovered hybrid pulsators and supply them in a catalogue of candidates. We also aim to compare the light curves produced with the TESS-*Gaia* Light Curve (TGLC) pipeline, currently underused in variability studies, with other publicly available light curves.

Methods. We compared dominant and secondary frequencies of confirmed hybrid pulsators in *Kepler*, extended mission Quick Look Pipeline (QLP) data, and nominal and extended mission TGLC data. We then used a feature-based positive unlabelled (PU) learning classifier to search for new hybrid pulsators amongst TESS AF stars and investigated the properties of the detected populations.

Results. We find that the variability of confirmed hybrids in TGLC agrees well with the one occurring in QLP light curves and has a high recovery rate of *Kepler*-extracted frequencies. Our ‘smart binning’ method allows for robust extraction of hybrids from large unlabelled datasets, with an average out-of-bag prediction for test set hybrids at 93.04%. The analysis of dominant frequencies in high-probability candidates shows that we find more pressure-mode dominant hybrids. Our catalogue includes 62,026 new candidate light curves from the nominal and extended TESS missions, with individual probabilities of being a hybrid in each available sector.

Conclusions. Our catalogue results in a major increase of TESS γ Dor / δ Sct hybrid pulsators, suitable for further asteroseismic studies.

Key words. Asteroseismology – Catalogs – Methods: data analysis – Methods: statistical – Techniques: photometric – Stars: oscillations (including pulsations)

1. Introduction

Stellar variability holds a special place in the theory of stellar structure and evolution (Eyer et al. 2023). Changes in brightness over time reveal a wealth of information about the inner workings of stars themselves, as well as their interactions with companions such as other stars and exoplanets. Over the last several decades, the field of asteroseismology (Aerts et al. 2010) has significantly benefited from nearly continuous high-precision photometry available from space missions. These space telescopes take advantage of the absence of atmospheric effects, absorption of certain parts of the electromagnetic spectrum, day-night cycles, and negative effects of light pollution. This results in a significant increase in the quality and quantity of continuous data available to the variable star community.

Variable star studies have been revolutionized by the high-precision long-baseline data of more than 160,000 stars provided by the *Kepler* (Borucki et al. 2010) satellite, including the K2 mission (Howell et al. 2014), as well as ongoing observations of hundreds of millions of stars by TESS (Ricker et al. 2015). Furthermore, *Gaia* (Brown et al. 2016) has also significantly

contributed to asteroseismic research (De Ridder et al. 2023), despite not being explicitly designed for variable star studies. Many additional discoveries are expected with the upcoming ESA PLAnetary Transits and Oscillations (PLATO, Rauer et al. 2025) mission, particularly in the context of its complementary science programme, which is set to explore all science cases not connected to the asteroseismology of exoplanet-hosting main-sequence stars.

Within the myriad of targets observed from space, OBAF pulsators — both radial and non-radial — are of particular interest to asteroseismologists, enabling precise testing of stellar evolution theories across a wide mass regime (De Ridder et al. 2023). Photometric data with μ mag precision has enabled the study of these objects in various domains, including the investigation of transport processes (e.g. Mathis 2009, Rogers 2015, Eggenberger et al. 2019), calculations of convective core masses (e.g. Mombarg et al. 2024), near-core rotation frequency estimations (e.g. Aerts et al. 2025), magnetic field studies (e.g. Mathis et al. 2021, Li et al. 2022b), multiplicity (e.g. Ijspeert et al. 2021), and many others. For the most complete summary of re-

cent findings, we refer the reader to the reviews by Kurtz (2022) and Aerts & Tkachenko (2024).

The stellar interior is probed by two types of stellar oscillations. In main-sequence stars, pressure (p) modes probe outer stellar regions and have frequencies higher than approximately 4 d^{-1} , while gravity (g) modes are sensitive to near-core regions and have frequencies below approximately 4 d^{-1} (Chapellier et al. 2012). However, the instability regions of the Hertzsprung–Russell diagram can overlap where stars exhibit both p and g modes (see Fig. 1, Aerts 2021). Here, we study one such case of overlapping regimes, particularly γ Dor / δ Sct hybrids. These hybrid pulsating stars (Dupret et al. 2005, Grigahcène et al. 2010) are especially interesting for asteroseismology, as they allow us to more precisely constrain stellar structure and evolution models of stars slightly more massive than the Sun — which have a convective core — thanks to the ability to probe outer and inner stellar interior layers simultaneously (Bradley et al. 2015). It is thus not surprising that this (sub-)class of pulsating stars received attention in both observational asteroseismology (e.g. Bradley et al. 2015) and asteroseismic modelling (e.g. Zhang et al. 2020).

Studies of γ Dor / δ Sct hybrids are limited by the scarcity of objects readily available for analysis. Hundreds of stars previously believed to be pure γ Dor and δ Sct pulsators observed by *Kepler* (Bowman et al. 2016, Li et al. 2020) were later revealed to be hybrids (e.g. Audenaert & Tkachenko 2022). Despite that, many have yet to be discovered in the TESS data (as demonstrated by Skarka et al. 2022 and Skarka & Henzl 2024), which are orders of magnitude more voluminous in the number of targets observed compared to *Kepler*. Manual inspection of millions of light curves of (AF) stars is impractical and thus requires automated processing algorithms to assist with the selection of candidates for further asteroseismic studies.

Asteroseismology has increasingly relied on machine learning when working with large observational datasets. A wide variety of machine learning frameworks have been used in variability studies, with some aimed at γ Dor / δ Sct hybrids specifically (Audenaert & Tkachenko 2022). Machine learning allows one to take population studies and ensemble asteroseismology to novel heights, enabling efficient data handling and proposing candidates for more detailed follow-ups. For a summary of the current state of machine learning in asteroseismology, we refer the reader to the review by Audenaert (2025).

Supervised machine learning, used when a dataset of labelled by experts is available, has seen extensive use in variable star studies. For example, Audenaert et al. (2021) used an ensemble of supervised learners, including a meta-classifier combining the output of other classifiers, on a labelled dataset with eight classes. Hey & Aerts (2024) used a random forest classifier on a dataset of six classes of variable stars, including hybrid pulsators. Some other applications include looking for a specific type of variable stars (van Roestel et al. 2018) and multi-survey classifications (Aguirre et al. 2019).

Unsupervised machine learning, used to discover hidden patterns in the data instead of assigning objects to particular classes, has also seen use in asteroseismology. Audenaert & Tkachenko (2022) used an unsupervised clustering algorithm to discern populations of pulsators in a dataset of γ Dor and δ Sct stars. Rizhko & Bloom (2024) used a self-supervised multi-modal contrastive learning to study a dataset of variable stars. Interestingly, *Gaia* data has been subject to a number of unsupervised machine learning applications, both within the domain of variability studies and beyond. For example, Huijse et al. (2025) performed an all-sky variability analysis on the entire *Gaia* DR3 dataset.

Ranaivomanana et al. (2025) used an unsupervised clustering algorithm to detect sub-populations of hot sub-luminous stars. One undisputed advantage of unsupervised algorithms is that they are free of human bias, meaning they are not constrained by the pre-defined labels as in supervised learning, which makes them well-suited to detect novel (sub-)classes of variable stars and check scientific intuition.

While the output of supervised machine learning is easier to evaluate than unsupervised learning, its main advantage — reliance on a large labelled set — is also its disadvantage when it comes to variability studies of rare classes such as γ Dor / δ Sct hybrids. Finding more examples from a rare class of astrophysical objects requires having a large set of labelled examples. However, those labels would have to come from manually inspecting the unlabelled data, a process we have already shown to be impractical. This ‘vicious circle’ slows down the progress of variability classification of rare pulsators, and subsequently their further asteroseismic follow-up studies. There are at least three ways to tackle this problem.

One potential solution is to use algorithms specifically designed to handle class imbalance, known as imbalance learning (Chen et al. 2024). Hosenie et al. (2020), Kang et al. (2023), and Zhang et al. (2023) have used various imbalance learning architectures to classify periodic variables, allowing for a better classification performance in situations where objects of one class could outnumber objects of other classes. Another option to consider is a family of anomaly detection algorithms. This paradigm deals with the detection of data points that are significantly different from the majority of data (‘normal’ observations). These algorithms can be used to both find novel astrophysical phenomena (e.g. Rebbapragada et al. 2009) and look for specific rare objects (e.g. Sánchez-Sáez et al. 2021 and Chan et al. 2022). Interestingly, this paradigm comprises both supervised and unsupervised anomaly detection algorithms, in contrast to imbalance learning algorithms that require labelled instances. For example, Huijse et al. (2025) have used unsupervised outlier detection methods for error analysis of their unsupervised clustering framework. Unsupervised anomaly detection, however, does not have any prior knowledge on anomalies (Pang et al. 2021), so the objects singled out by the algorithm might not belong to the class of interest in the first place, while supervised anomaly detection also relies on time-intensive labelling of both ‘normal’ and anomalous observations (Pang et al. 2021).

Finally, it is possible to use semi-supervised learning by taking advantage of the unlabelled examples to learn the structure and to regularize the classification boundaries learned from the scarce labelled examples (Van Engelen & Hoos 2020). These algorithms remedy situations where labelled data is difficult to obtain, such as in medical imaging and risk assessment applications. This circumvents the need to label most of the dataset, and instead label only some of the positive (class of interest) and/or negative (other) data. Although this approach might appear ideal, semi-supervised machine learning shares evaluation and interpretation challenges of unsupervised learning, which might explain why it has not seen a lot of use in variability studies so far. In this work, we aim to rectify this research gap and demonstrate the potential of semi-supervised learning for rare object detection in asteroseismology.

The most common object of study in variability classification is a light curve. Most data-driven applications for asteroseismology relying on TESS data have so far not used the light curves produced with TGLC (Han & Brandt 2023) pipeline. Until now, the focus has been on other science products available from the

Mikulski Archive for Space Telescopes (MAST),¹ namely QLP (Huang et al. 2020a, 2020b, Kunimoto et al. 2021, 2022) and TASOC (Handberg et al. 2021, Lund et al. 2021) light curves. Our work therefore, has two connected goals:

- To compare the dominant and secondary variability of confirmed hybrid pulsators in the TGLC data from the nominal (sectors 1-26, 30-minute cadence data) and extended (sectors 27-, 10-minute cadence data) TESS mission with other available light curves (QLP and *Kepler* data); and
- To use a semi-supervised machine learning method on a dataset of known hybrids and unlabelled set of TESS observations to create a large all-sky catalogue of candidate hybrid pulsators with individual probabilities of being a γ Dor / δ Sct hybrid in a single sector and multi-mode variability extracted from TGLC light curves. Our new public catalogue can then be used by the asteroseismology community for follow-up studies, including unsupervised machine learning clustering, population studies, mode identification, asteroseismic modelling, and probing for other astrophysical properties.

2. TGLC for hybrid pulsator studies

Several high-level science products (HLSP) based on raw TESS data are available. In this Section, we describe the TGLC (Han & Brandt 2023) pipeline and compare the available data for known hybrid pulsators with other publicly available data, particularly *Kepler* and QLP light curves.

2.1. The TGLC pipeline

A complete description of the TGLC methodology has been given by Han & Brandt (2023), so here we only briefly summarize it for convenience. The TGLC pipeline allows to readily access light curves up to 16th magnitude by taking advantage of *Gaia*'s astrometry and photometry, particularly of the high angular resolution offered by *Gaia* (Brown et al. 2016). This is in contrast to other available pipelines, where only sources up to 13.5 magnitude are readily available for download on MAST.

TGLC provides both point spread function (PSF) and aperture light curves for each source. Both rely on the position and brightness of stars in the field from *Gaia* to construct the final science product. In general, an aperture TGLC light curve delivers higher-quality data than a PSF one in crowded fields. This leads to more reliable amplitude estimations for variable stars (Han & Brandt 2023). In Fig. 1 we show an example of both a PSF and an aperture light curve of a confirmed hybrid pulsator with their respective Lomb-Scargle periodograms. For that target, we find a similar peak structure for both periodograms, although the one for aperture photometry exhibits higher amplitudes. The standard deviation of the flux for the PSF light curve is twice that of the aperture one. This is caused by the noisier floor throughout the entire frequency space.

We additionally manually inspected light curves of hybrids and a random subsample of other sources in TGLC (see Section 4.1). We find no pathological long-term instrumental trends in light curves in both science products. However, we note an occasional systematic at 1 d^{-1} , previously reported by Hey & Aerts (2024), which appears in PSF light curves more often. We further found that calibrated PSF flux is missing for 5.4% of our sample. We therefore opted to use aperture light curves, both to have




Fig. 1: Aperture (top) and PSF (bottom) light curve (grey) and Lomb-Scargle periodogram (overplotted in burgundy, clipped at 30 d^{-1} for visibility) of a confirmed hybrid pulsator (DR3 2147267632621883776) detected in TESS sector 40. Note how this target is p- and g-mode dominated on the aperture and PSF light curves, respectively.

the most reliable amplitude estimations for variable stars and for completeness of the dataset. The remainder of this Section focuses on comparing the two dominant independent frequencies (f_1 and f_2) extracted from the aperture TGLC light curves with other light curves available on MAST for the confirmed hybrids.

2.2. Dominant and secondary variability of labelled hybrids

We preprocessed each aperture TGLC light curve using both the normal TESS² and TGLC quality flags (Han & Brandt 2023), clipping flux outliers above 3 standard deviations, and applying a high-pass Gaussian filter to remove long-term trends caused by the thermomechanical behaviour of the satellite, which does not contain any astrophysical information. More concretely, we subtracted the Gaussian smoothed time series with $\sigma_G = 100$ datapoints to remove long-period instrumental trends. For hybrids from the extended TESS mission, we found that the ratio of the standard deviation of the subtracted signal to the original ($\sigma_{\text{subtr}}/\sigma_{\text{original}}$) is 0.107 ± 0.066 , suggesting that the preprocessing pipeline removes signal comprising one tenth of the entire light curve variability, which could be instrumental. For a subset of non-hybrid sources, the ratio of $\sigma_{\text{subtr}}/\sigma_{\text{original}}$ was expectedly higher at 0.199 ± 0.114 , as signals in a largely non-variable set are dominated by instrumental trends, noise, and systematics removed by preprocessing, rather than genuine (intrinsic) variability that passes the filter.

² https://archive.stsci.edu/files/live/sites/mast/files/home/missions-and-data/active-missions/teess/_documents/EXP-TESS-ARC-ICD-TM-0014-Rev-F.pdf

¹ <https://archive.stsci.edu>

2.2.1. Effect of *Gaia*-informed light curve extraction: TGLC vs. QLP light curves

Differences in light curve extraction pipelines can have a significant impact on the light curve, as well as the frequency content detected in the periodogram, and the amplitudes of pulsations (Han & Brandt 2023, Hubrig et al. 2024, Kunimoto et al. 2024). The QLP light curves remain one of the most used in variability studies and offer greater sector availability on MAST compared to TGLC. We therefore compared the dominant and secondary variability for the hybrid pulsators available in both pipelines by producing frequency-frequency plot from preprocessed data, shown on Fig. A.1. We applied the same preprocessing to both types of light curves for a fair comparison (see Section 2.2). Both light curves agree well for both extracted frequencies, having an agreement of 77.13% and 56.44% within 0.1 d^{-1} of the unity line, respectively. There is noticeable scatter for f_2 (bottom) both above twice unity and below half-unity, where one pipeline detects a high frequency and another a low frequency, or vice versa. This suggests that modes with low amplitudes are susceptible to differences in processing pipelines and noise. The agreement is generally better at higher frequencies, partially caused by the occasional TGLC systematic at 1 d^{-1} , which is extracted as a dominant peak, while a genuine mode might be recovered by the QLP.

2.2.2. Longer baselines for g-mode detection: TESS TGLC and QLP vs. *Kepler* light curves

Single-sector TESS observations are typically too short to correctly detect all g modes. The four-year *Kepler* data, on the other hand, has proven to be a treasure trove for the studies of g-mode oscillators (Uytterhoeven et al. 2011, Tkachenko et al. 2014, Li et al. 2020, Pedersen et al. 2021, Aerts & Tkachenko 2024). We therefore investigated the efficiency of recovering *Kepler*-extracted frequencies. Fig. 2 shows that the frequency recovery rate for f_1 is approximately the same in both TGLC and QLP, although the systematic at 1 d^{-1} affects the recovery rate of the lower frequencies. We found 50.88% and 43.42% recovery rates for QLP and TGLC, respectively. At higher frequencies, the two pipelines agree better, within $<1\%$ differences in measured frequencies, allowing us to conclude that both pipelines are of similar quality for hybrid pulsator studies, noting the potential to find faint hybrids with TGLC light curves. While longer-baseline observations would lead to higher performance of any automated classification pipeline, concatenating several TESS sectors together causes issues of including data with inconsistent quality. We therefore opted to work with the shorter-baseline individual TESS sector data.

2.2.3. Finer sampling for p-mode detection: primary- vs extended-mission TGLC light curves

The detection of p modes is known to benefit from data with shorter cadence (Murphy 2012). We compared the dominant and secondary variability of confirmed hybrids in the extended and nominal TESS missions. We also considered a binned version of the extended mission where every three observations are averaged to simulate the 30-minute cadence. As seen in the middle and left panels of Fig. A.1, the agreement between the extended and nominal mission data for f_1 and f_2 is 74.33% and 44.68%, respectively. Between the 10- and 30-minute cadence extended mission data, the agreement is 90.21% and 67.46% within 0.1 d^{-1} of the bisector for dominant and secondary variability, re-




Fig. 2: Comparison of *Kepler* f_1 (purple histogram) recovery efficiency in QLP (gold), TGLC (genuine and binned to 30-minute cadence in teal and grey, respectively).

spectively. In both cases, more p modes are detected in extended mission TGLC indicated by the scatter below the half-unity line. The quality of the nominal mission data is lower than the down-sampled (binned) extended mission. The better result for 10- versus 30-minute cadence extended mission data comparison can be explained by them coming the same sector, whereas comparisons with nominal mission data are based on different sectors, capturing systematics in some sector data but not the other.

The binned TGLC data also sees a significant drop in *Kepler* recovery rate, where the detection of p modes is affected more than g modes (Fig. 2). The secondary frequency f_2 is more sensitive to the lower frequency resolution because of its lower amplitude (Fig. A.1). While the performance of 30-minute cadence data is undoubtedly lower than for 10-minute cadence data, we argue that it is justified to use both in separate training sets, provided their sufficient purity. Manual or semi-automatic vetting of light curves, ensuring that the hybrid pulsation signal is the one captured in the light curve rather than a systematic in the training set, should allow new candidates to be detected. We also note that sampling differences between the extended and nominal mission warrant separate classifications when frequency- and/or period-domain features are used, as the higher Nyquist limit would allow one to capture more p-mode variability in extended mission data. In practice, the median Nyquist frequency of the labelled data we used from nominal and extended missions amount to 24.00020 ± 0.00011 and $72.0009 \pm 0.0018 \text{ d}^{-1}$, respectively.

3. PU learning

The simplest classification settings in machine learning is binary classification. In a binary classifier, a model is tasked with assigning instances of a tuple (x, y) , where x are the features and $y \in \{0, 1\}$ is a label, to one of two classes: positive ($y = 1$) or negative ($y = 0$) (Müller & Guido 2016). Being a supervised learning method, the classifier must be provided with labelled examples of both positive and negative instances. In the case of rare objects such as γ Dor / δ Sct hybrids, providing a large enough labelled set would require a substantial manual effort, particularly to label a representative set of negative instances.

In Positive Unlabelled (PU) Learning, a family of semi-supervised algorithms, only part of the data is labelled. Instances in PU Learning are instead represented by a tuple (x, y, s) , where $s \in \{0, 1\}$ denotes whether an instance was selected to be labelled

(Jaskie & Spanias 2019, Bekker & Davis 2020). An important differentiating feature of PU Learning is that only positive instances are labelled, but not all of them. This means that an unobserved instance ($s = 0$) can be either positive ($y = 1$) or negative ($y = 0$), while observed instances ($s = 1$) are always positive: $p(y = 1 | s = 1) = 1$. An alternative way to reason about this is that the probability of a negative instance being labelled should be zero: $p(y = 0 | x, s = 1) = 0$ (Elkan & Noto 2008), unless it has been mistakenly labelled as positive by an expert. A selection of labelled instances is guided by one of the three labelling mechanisms, so a labelled distribution is a biased distribution of all positive instances: Selected Completely at Random (SCAR), Selected at Random (SAR), or Probabilistic Gap (a special case of SAR where positive instances are less likely to be labelled if they are closer to the negative ones in their attributes). Additionally, both positive (labelled) and unlabelled datasets are assumed to be independently and identically distributed (i.i.d.) samples of their respective real distributions. The data are assumed to be separable, meaning there is a classifier that can separate positive instances (Bekker & Davis 2020). A schematic representation of such a learning setting is shown in Fig. B.1. Several algorithms exist in the PU Learning family, from simpler two-step methods training binary classifiers to more complex probabilistic approaches and deep PU Learning (Jaskie & Spanias 2019, Li et al. 2022a, Bekker & Davis 2020). For the most complete summary, we refer the reader to the review by Bekker & Davis (2020).

Despite its advantage when searching for objects with few labels in large datasets (Yang et al. 2021), PU Learning has not yet seen much use in astrophysics. One notable exception is the search for Carbon stars in the Sloan Digital Sky Survey (Du et al. 2016). A similar learning approach, One-Class Classification, has been applied to look for quasars in Wide-Field Infrared Survey Explorer data (Solarz et al. 2017). In general, however, PU Learning and semi-supervised learning in general remain rather underutilized in variability studies, likely caused by the evaluation and validation complexity.

4. Application of PU learning to TESS data

In this Section, we outline the four key stages of our machine learning pipeline: data acquisition and preprocessing, feature engineering and extraction, model selection and training, and performance evaluation. As PU Learning requires having both a positive (labelled) and a (typically larger) unlabelled dataset, we begin by describing the two below.

4.1. The dataset

4.1.1. Positive (labelled) dataset of confirmed hybrids

Our initial labelled dataset included published hybrids detected in the northern continuous viewing zone (CVZ) of TESS (Skarka et al. 2022) and in the *Kepler* field of view (Bowman et al. 2016, Li et al. 2020). We used two types of TESS light curves for these objects. First, we used the light curves available for download on MAST, which include 3,292 and 785 light curves from the nominal and extended missions, respectively. Second, because there are less TESS light curves available for the extended mission classification, we used the TGLC code made available by Han & Brandt (2023) to extract TESS light curves from Full Frame Images for the extended mission.

The confirmed hybrids from Skarka et al. (2022) indicate the dominant variability of the target (γ Dor / δ Sct or δ Sct / γ Dor hybrids). For the purposes of this study we do not differentiate

between the two, simply labelling all datasets as hybrids. The dominant variability of hybrids (Fig. A.1) across the entire labelled set reveals that there are more p-mode dominated hybrids in the training set. This is expected to propagate into the classification (see Section 5.1) by favouring δ Sct / γ Dor hybrids from the unlabelled set. We suggest that this is partly due to the short baseline of TESS, which makes the detection of g modes difficult.

Additionally, not every source labelled as hybrid in one survey will manifest as a hybrid in another. Since part of the labelled sample comes from *Kepler*, some of the g modes will not be resolved in TESS light curves due to different observational baselines. Removing these sources is a necessary step to avoid confusing the classifier and assigning high probabilities to other (non-hybrid) light curves, which defeats the purpose of an automated detection pipeline. We therefore manually vetted each TGLC light curve and their respective periodograms to ensure that only high-quality observations with low noise levels and systematics are used for training. After manual inspection of both publicly available and manually generated light curves, our labelled dataset comprises 480 unique sources, totalling 1,404 light curves in the nominal and 2,376 light curves in the extended mission.

4.1.2. Unlabelled dataset of AF stars in TESS

The unlabelled dataset comprises the aperture TGLC data for stars that are available to download on MAST (up to TESS sector 44 as of 31 March 2025). We used the same filtering approach to work with the TESS Input Catalog as described in Ijspeert et al. (2021), except that our cut is based on the TIC effective temperature (T_{eff}) instead of colours. We limited our selection to stars with $10,000 \geq T_{\text{eff}} \geq 4,000$ K, where the lower limit is generously placed below the temperature limit of solar-like oscillators to account for potential errors in measured temperatures. The upper limit rejects most Slowly Pulsating B-type Stars / β Cep (SPB / β Cep) hybrids and p- and g-mode pulsating subdwarfs (sdBVs), which are outside the scope of this work.

Applying the T_{eff} cut yielded 29,309,213 light curves for dwarf stars, 16,168,167 in the nominal mission and 13,141,046 in the extended mission. The number of light curves differs significantly from sector to sector, with the smallest set coming from sector 4 (49,759) and the largest from sector 12 (1,716,982). Instances in the positive (labelled) set were removed from the unlabelled set before training regardless of the sector to avoid rediscovering hybrids that have already been found.

Application of a PU Learning algorithm to this dataset is based on an assumption that there are instances in this unlabelled dataset that belong to the positive class of hybrid pulsators, which have not been labelled (discovered) yet (Fig. B.1). In our case, it is based on the lack of all-sky hybrid searches in the TESS data, as previous studies have been limited to the northern (Skarka et al. 2022) or southern (Skarka & Henzl 2024) CVZ of TESS, or a limited number of TESS sectors (Audenaert & Tkachenko 2022).

4.2. Classification features

The search for hybrid pulsators, whether manual or automatic, relies on variability analysis, often resulting in the use of frequency-domain features, such as frequencies or periods, amplitudes, and phases (Debosscher et al. 2007, VanderPlas 2018).

As hybrid pulsators exhibit variability in two different frequency regimes, previous studies used binning to extract this information, which involves dividing the frequency space into smaller parts, from which features are extracted. There are several approaches to periodogram binning: (i) bins can either be equal or vary in size; and (ii) bins can be overlapping or non-overlapping. Skarka et al. (2022) and Skarka & Henzl (2024) previously used two bins to look for γ Dor / δ Sct hybrids within a population of AF stars: $< 5 \text{ d}^{-1}$ and $> 5 \text{ d}^{-1}$. Hey & Aerts (2024) used three partially overlapping bins in their analysis of a population of OBAF pulsators: $< 1 \text{ d}^{-1}$, $< 5 \text{ d}^{-1}$, and $> 5 \text{ d}^{-1}$. A graphic representation of these two approaches is shown on Fig. 3 (top and middle).

In this work, we used a dynamic binning approach to automatically find the most optimal bins for hybrid detection. For each labelled hybrid, we preprocessed the light curve, computed a frequency spectrum using a Lomb-Scargle periodogram method (VanderPlas 2018) with $f_{\min} = 10^{-4} \text{ d}^{-1}$ and $f_{\max} = \frac{1}{2\Delta t}$. This low lower limit is inspired by the fact that pulsators with low-frequency g modes have beating periods that may span many years, notably longer than the 4-year nominal *Kepler* light curves (Van Beeck et al. 2021). We then identified peaks that are the local optima with `scipy.signal.findpeaks` and created a power histogram with m bins. From the histogram, we calculated the cumulative power distribution across the m frequency bins, normalized it, and applied k power-law scaling. We then applied linear interpolation of n linearly spaced bins to map them on the cumulative power distribution. This resulted in n non-overlapping non-equal bins, which are finer in more ‘important’ parts of the frequency space (where more peaks are present) and sparser bins elsewhere. We repeated this process for all labelled hybrids, averaging the bin edges for the entire labelled dataset, separate for nominal and extended missions hybrids. A graphic representation of our approach with $m = 50$, $n = 10$, and $k = 1$ (no power law scaling) is shown on Fig. 3 (bottom). The averaged bins were then applied to the entire unlabelled set to extract features for classification, ensuring that the bins are best-fit to look for a specific type of pulsators. From each of the bins we extracted the amplitude of the highest peak and the four statistical moments: mean, variance, skewness, and kurtosis, with higher feature values in multiple relevant bins meant to separate hybrids from pure p- or g-mode pulsators. The same was repeated for the period space, resulting in two sets of features for each light curve.

There are two reasons why we only extracted the amplitudes from the bins and not the highest frequencies themselves. First, the bin widths themselves serve as proxies for frequencies and periods, especially in the finer-spaced bins. Amplitudes therefore allow for the extraction of additional information about variability, although a certain degree of mutual information is expected with means extracted from the bins. Second, an ideal-case variability study would focus on the entirety of the frequency (or period) space as a feature. However, it is limited by practical issues, such as the curse of dimensionality. Using the compressed periodogram as a basis for features allows to scale the pipeline to work with millions of light curves. We note that using the full periodogram as a feature is an interesting area of research to explore for smaller-scale classification problems.

Both individual and averaged bin edges can reflect the underlying physics of hybrid pulsators. We observed finer g-mode bins due to rotational modulation, as rotational peaks can contribute significantly to the cumulative power distribution which does not discriminate between the sources of power. In stars with no rotational modulation, g-mode regime was typically binned sparser.




Fig. 3: Graphic representation of binning approaches applied to a confirmed hybrid pulsator (DR3 1653802003712212992). Top plot: bins from Skarka et al. (2022) and Skarka & Henzl (2024), centre plot: bins from Hey & Aerts (2024), bottom plot: this work.

Including a significant number of stars with rotational peaks in the training set results in finer low-frequency binning, especially with power law scaling.

4.3. PU bagging classifier

Given that positively labelled instances (known hybrids) have been selected based on their attributes, our learning problem builds on the SAR — Selected at Random — assumption. This limits the available pool of PU Learning algorithms that can be applied. SAR-based classifiers do not incorporate priors, meaning that we do not inject the expected occurrence rate of hybrids into the model. In our framework, we used a PU bagging classifier (Mordet & Vert 2014). Bagging is an example of ensemble learning, which involves using multiple weak learners — classifiers making individual classifications — and combining their results. Bagging specifically involves learning T_{weak} weak learners from i_{boot} bootstrapped samples of the entire dataset.

Due to the large size of the unlabelled set, we first randomly sampled 15% of observations from each sector, to not introduce potential sector-specific biases. We then trained an ensemble of $T_{\text{weak}} = 4000$ decision trees (weak learners), on a bootstrapped sample i_{boot} of unlabelled instances equal to the size of the labelled sample, separately for nominal and extended TESS missions, and separately only using frequency- and period-domain

features. Each of the trees makes an out-of-bag prediction — prediction on objects not in i_{boot} — resulting in a distribution of the probability of being a hybrid for every unlabelled example. We then used the trained model to predict the probability of a hybrid nature for all observations from each sector.

The final probability for each unlabelled target being a hybrid pulsator was determined as an average of n_{oob} out-of-bag predictions associated with it. Notably, due to the all-sky nature of this study, multiple predictions were made for the same star in all TESS sectors it was available. We follow the ‘single sector - single instance’ approach used by Hey & Aerts (2024), meaning that a star had multiple opportunities to be labelled a hybrid. We did not combine the predictions from different sectors to make a final prediction for a particular object, to avoid decreasing probabilities for genuine hybrids due to instrumental noise characteristics occurring in a single sector. Our final probability for a single light curve was an average of two predictions made separately using frequency- and period-domain features.

An important note to make is that a normal grid search or cross-validation is not available for hyperparameter tuning in the PU Learning setting as there is no precision and therefore no f-score (mean of precision and recall). An alternative f-metric exists for some PU Learning methods, but only for the ones following the SCAR — Selected Completely at Random — assumption. On the other hand, our architecture has an advantage that it is based on decision trees, which are computationally inexpensive to learn. We therefore chose $T_{weak} = 4000$, resulting in reduced uncertainty due to a high number of predictions for each unlabelled instance.

4.4. Classification evaluation & error analysis

Evaluation of this classification framework is challenging as precision is not readily available, and consequently, nor is the f-score metric. While precision can be estimated (Claessen et al. 2015), it is not exact as for binary classification and is therefore less reliable. We instead analysed our classification by inspecting high-probability predictions of some of the targets from our sample, making predictions for labelled non-hybrids from Hey & Aerts (2024), and analysing feature importance.

4.4.1. Recall

Recall — the fraction of true positive examples recovered from the unlabelled set — is the only traditional evaluation metric available for SAR-based PU Learning methods. As PU Learning relies on an idea that some unlabelled positive instances exist in the unlabelled set (Fig. B.1), recall allows us to understand how good the model is at finding new genuine hybrids. Not knowing all positive instances contained in the unlabelled dataset, however, means that this metric is not as exact as for the binary classification but rather a good proxy for it. To estimate it, we set aside a 15% of the labelled hybrids to be used as a test set.

Our classifier reported a high recall, demonstrated in Fig. 4, both in the nominal and extended TESS missions. The average out-of-bag prediction for all known hybrids was 93.04%. 92.08% receive average probabilities above 80%, 88.56% above 85%, 79.75% above 90%, and 61.62% above 95%. Some hidden hybrids from the nominal mission received lower scores than extended mission hybrids. We assume that this is due to the larger volume of the unlabelled set, which generally requires more weak learners for a more reliable probability estimation. From the entire unlabelled subset used in training, only 1.00% received




Fig. 4: Combined out-of-bag predictions for the unlabelled instances from nominal and extended missions with 15% of the total data (golden) and hidden hybrids from nominal (teal) and extended (purple) missions. Teal (94.25%) and purple (90.98%) dotted lines represent mean predictions for hidden hybrids from nominal and extended missions, respectively. Golden dotted line represents the grand mean for all unlabelled targets (7.57%) used for training.

probabilities above 80%, 0.61% above 85%, 0.31% above 90%, and 0.09% above 95%. This suggests that hybrids construe only a small fraction of dwarf stars; we revisit this in Section 5.1.

4.4.2. Estimating contamination from rotational modulation and pure g- and p-mode pulsators

Distinguishing γ Dor / δ Sct hybrids in a vast unlabelled sample is especially challenging due to the presence of other types of pulsators and, most critically, their parent classes: pure γ Dor and δ Sct variables. Another potential source of confusion is the class of rotational variables, which is known to be cross-contaminated with g-mode pulsators from supervised variable star classification results by Audenaert et al. (2021) and Hey & Aerts (2024). Although the unresolved comb of frequencies of g-mode pulsators is relatively easy to differentiate manually from one or three rotational peaks, it is not a trivial task for an automated method, especially without harmonic analysis and prewhitening (Debusscher et al. 2007, Richards et al. 2011) which we explicitly did not use in this study due to the volume of the unlabelled dataset and undesired additional data manipulation. In our framework, this issue leads to the expectation that some high probabilities may be assigned to δ Sct stars with rotational modulation instead of being a hybrid pulsator. To test that, we predicted the hybrid class probability (Fig. 5) on a subset of the three potentially confusing classes: γ Dor, δ Sct, and rotational variables, from Hey & Aerts (2024). We also predicted probabilities for a class of eclipsing binaries, as there could potentially be a (hybrid) pulsation component in a binary system, whose orbital period is found in the realm of (sub-)multiples of g modes.

As expected, the class containing δ Sct stars received the highest out-of-bag scores, with a mean of 0.45. Pure g-mode pulsators have a mean score of 0.37. We note that this class includes both γ Dor and SPB stars that can not be differentiated from each other without T_{eff} (Audenaert et al. 2021, Aerts et al. 2023), so predicting on it with our classifier is essentially the same as predicting on pure γ Dor stars. Eclipsing binaries and rotational variables received mean scores of 0.39 and 0.32, re-




Fig. 5: Out-of-bag predictions for pure δ Sct stars (purple), eclipsing binaries (burgundy), pure γ Dor / SPB pulsators (teal), and rotational variables (golden) from TESS sectors 1 and 14 from Hey & Aerts (2024).

spectively. This allows us to conclude that our classifier should in principle be able to differentiate genuine hybrids from variable stars with variability in just one of the frequency regimes.

Given the distribution of dominant variability of the labelled set (Fig. A.1), we inspected the highest probability predictions for this non-hybrid dataset. The highest probabilities have all been assigned to the δ Sct class, but only three sources receive scores above 80%. This suggests that placing a probabilistic threshold at 80% would allow one to avoid most of the contamination coming from variable but non-hybrid sources. For the remainder of this paper we thus used the 80% cut-off for our candidate targets.

4.4.3. Feature importance

Feature importance allow us to understand how much each feature contributes to the model's predictions. While it does not directly evaluate classification performance, it provides a valuable insight into whether the model is working as intended. The importances of features of our models shared certain similarities. Inspecting individual feature importances revealed that the most important features for frequency-based models were associated with higher-frequency bins responsible for capturing p-mode variability. The importance of features of period-based models were the highest for short-period bins, also capturing p-mode variability. This can be seen as a result of the predominance of hybrids where f_1 is a p mode in the training set (Fig. A.1).

Summing the importances of the five features associated with each of the fifteen frequency bins (Fig. 6) confirmed that both models favoured high-frequency regimes. For frequency-based classifiers, the low-frequency regime was captured by the first three bins, which had moderately high importance (> 0.05). High importance was attributed to the first two high-frequency bins. There is an increasing importance trend in the frequency space above 10 d^{-1} , where one might expect genuine p modes or their aliases mirrored around the Nyquist limit (Hey et al. 2021), which results from the cumulative power distribution summing the power regardless of its origin. The importance of bins of the period-based classifiers also had a high preference for p-mode variability similar to their frequency counterparts.

From the feature importance and the summed importance of the bins, we could expect both a predominance of p-mode dom-




Fig. 6: Normalized importance of the frequency (top) and period (bottom) space, plotted at the centre of each frequency bin of nominal (burgundy) and extended (grey) mission models. Note that the final frequency bin goes to the Nyquist limit and the final period bin goes to approximately the inverse of the closing edge of the first frequency bin — they are plotted in their respective places for visibility.

inated hybrids in the high-probability candidates from the unlabelled set and p-mode contaminants. However, combining the prediction of two classifiers for a sample is expected to alleviate some of this bias, producing a more balanced classification result.

5. All-sky candidate hybrids catalogue

5.1. Frequency analysis

After classification, we filtered the data to remove potential contamination by false positives and contamination by instrumental effects. From a Lomb-Scargle periodogram for high-probability candidate hybrids ($p(\text{hybrid}) > 0.8$) we filtered out those that did not have at least one significant peak in frequency regimes $> 4 \text{ d}^{-1}$ (f_g) and $< 4 \text{ d}^{-1}$ (f_p) with a $< 1\%$ false-alarm probability (FAP). For this filter, f_p was the highest peak that was not an integer or half-integer multiple in the range of $[1.5, 15.5]$ of f_g to also avoid potential contamination by pure g-mode pulsators with harmonics in higher frequencies and eclipsing binaries with no pulsation component. We also removed candidates where a signal with $> 1\%$ FAP check was a near-zero frequency (mainly for the nominal mission), 0.9, or 1 d^{-1} (mainly for the extended mission), within twice the TESS sampling frequency f_{sample} . These three filters left us with 154,778 out of 284,325 light curves. The final sample was further reduced to 62,026 by applying a $T_{\text{eff}} \geq 6,000$ cut from Gaia DR3 for homogeneity with the training set (see Section 5.2). This resulted in a final fraction of candidate hybrids of 1.9% and 2.2% of the total analysed targets in nominal and extended mission datasets, respectively.

In Fig. 7 we show the frequencies associated with the highest frequencies in the $<4 \text{ d}^{-1}$ and $>4 \text{ d}^{-1}$ frequency (top left) and corresponding period (bottom left) regimes with only candidates, with only $p(\text{hybrid}) > 0.95$ plotted for visibility. The frequency-frequency relation (top left) shows a multi-modal distribution with respect to both f_g and f_p , potentially linked to differences in internal structure of these stars. In the period space (bottom left), we found two distinct groups of pulsators with short — a dense population with $P_g \lesssim 1$ with a wide range of P_p — and long-period — a less numerous subset with $P_g \gtrsim 1$ and typically lower P_p values — g modes, potentially indicative of differences in the spherical degree l of their dominant pulsation modes (Chapter 7 of Aerts et al. 2010). As future work, we intend to investigate period spacing patterns of targets in these two clusters and compare them with stars with known identified modes (Li et al. 2020). We also note that a number of hybrids in the training set are fast rotators and some of the manually inspected high-probability candidates show rotational splittings. This makes these stars suitable candidates for further studies to determine internal rotational properties and angular momentum transport mechanisms in AF stars (Aerts et al. 2025).

The harmonic avoidance criterion applied together with the FAP filter is responsible for a large portion of the targets cut. We show the harmonic behaviour of candidates without such a filter on Fig. C.1, which depicts a period-period relation where P_p is simply the inverse of the highest peak in the $>4 \text{ d}^{-1}$ frequency regime. The targets with low log-amplitude ratios are sources where an extracted P_p was a harmonic of P_g instead of a genuine p mode. We manually inspected sources from $P_p = 2 \times P_g$, $P_p = 3 \times P_g$, $P_p = 4 \times P_g$, and $P_p = 5 \times P_g$ (one example of each is shown on Fig. C.2). Some were eclipsing binaries, potentially with a pulsational component. This is because a feature discriminative of this type of stars is in the time domain (Hey & Aerts 2024), which we did not consider in this study. We note that increasing the probabilistic threshold (e.g. to $p(\text{hybrid}) > 0.95$ on Fig. C.1, right) filtered these sources out, meaning they receive lower probabilities than genuine hybrids. This suggests that the initial probabilistic cut at $p(\text{hybrid}) > 0.8$ might be too generous, depending on the nature of the foreseen asteroseismic follow-up study.

As can be seen on Fig. 8, most labelled hybrids are p-mode dominated, and so are the detected candidates. The distribution of power in the candidate catalogue is partly linked to biases propagated through the model. Our model, however, did not worsen this bias, evidenced by the distribution (Fig. 8, top). The shapes of distributions for both f_1 of two populations are not identical but resemble each other, and so are the ones for A_1 . The distribution of A_1 for the unlabelled set is shifted to the right, suggesting that we found higher-amplitude targets than we have in the labelled set.

This is confirmed by Fig. 9 (only sources with $p(\text{hybrid}) > 0.9$ are plotted for visibility), which shows that there are very few candidates with low amplitudes in both p- and g-mode regimes (cluster on the bottom left), while more labelled hybrids were found in the same region. It also shows that more hybrids cluster to the left, meaning that the detected population is dominated by oscillators with a lower-amplitude g-mode component. Additionally, we found more hybrid candidates with a strong component in one of the regimes and a weaker one in the other than stars with equally strong p and g modes, as was also the case for the training set examples.

The density of high-amplitude p modes is higher than the one for g modes — further confirming that most candidates are p-mode dominated. There are several explanations for this.

First, single-sector TESS observations do not have sufficient frequency resolution to fully resolve g modes, as concluded in Section 2.2.2. One potential way to remedy this would be to stitch together observations from multiple sectors and classify longer light curves. Second, pure δ Sct stars are somewhat rare in space photometry (Grigahcène et al. 2010) as they have could have a g mode if the quality of the light curve is high enough, so a p-mode pulsator is more likely to be labelled a hybrid. Finally, the features used in this study relied on the cumulative power distribution, which does not discriminate against the origin of the variability, be that a genuine pulsation, a rotational peak, or a systematic at low frequencies. This gave a p-mode pulsator more chances to be confused with a hybrid than a g-mode pulsator that already has peaks in the same frequency regime.

Following the approach of Hey & Aerts (2024), we performed two-sample Kolmogorov–Smirnov tests with 95% confidence level for the distributions of amplitudes and frequencies of the dominant modes on labelled hybrids, the final list of high-probability candidates following our three filtering stages, and pure p- and g-mode pulsators from Hey & Aerts (2024). The null hypothesis was that the two distributions are sampled from the same distribution. Each of our comparisons reported a p value around 10^{-4} . For the comparison of labelled hybrids and high-probability candidates, this reinforced the underlying SAR assumption of PU Learning, suggesting that the labelled hybrids had not been chosen randomly from the total population of hybrids but rather for their specific attributes. When comparing both labelled hybrids and high-probability candidates with pure p- and g-mode pulsators, this can be understood in terms of several excitation mechanisms contributing to the oscillations of hybrids (Hey & Aerts 2024, Mombarg et al. 2024). Further exploitation of the sample could focus on comparing the distributions of fundamental parameters or pulsation properties beyond f_1 and A_1 .

5.2. Fundamental parameters

For our final candidate selection, we investigated the fundamental parameters, namely effective temperature T_{eff} , stellar luminosity L , and radius R . We first computed L following the method from Fritzewski et al. (2024) by calculating extinction-corrected absolute magnitudes and applying bolometric corrections to get the bolometric magnitude using *Gaia*³ photometry and GSP-Phot parameters. We then computed the radius from the Stefan-Boltzmann Law. We estimated the uncertainties of these parameters using standard uncertainty propagation rules from *Gaia* DR3 properties and assuming that $\sigma_{BC} = 0.05$. As stated in Section 5.1, we initially found that the distribution of T_{eff} for detected candidates extended further to cooler stars than for the labelled set (similar to what was found in De Ridder et al. 2023), which is a result of us including stars with $T_{\text{eff}} \geq 4,000$ K in the unlabelled set and TIC and *Gaia* DR3 reporting different temperatures. In this study, high-probability candidates with $6,500 \geq T_{\text{eff}} \geq 4,000$ were later removed for homogeneity with training hybrids.

We also observed (Fig. 10, left) that a number of candidates showed $T_{\text{eff}} \geq 10,000$ K. This is in line with De Ridder et al. (2023), Aerts et al. (2023), Hey & Aerts (2024), and Mombarg et al. (2024), who have already found that the transition from γ Dor / δ Sct to SPB / β Cep pulsators is not a sharp cut-off but rather a continuum, in contrast to predictions from instability computations. These studies also already reminded us

³ <https://gea.esac.esa.int/archive/>




Fig. 7: Frequency-frequency (top) and period-period (bottom) plot of candidate hybrids with $p(\text{hybrid}) > 0.95$ (left) and labelled hybrids (right) for highest peak detected in frequency regimes $<4 \text{ d}^{-1}$ (f_g) and $>4 \text{ d}^{-1}$ (f_p) with a $<1\%$ FAP. The colour bar represents $\log_{10}$ amplitude ratios of the two frequencies (periods).

that the impact of fast rotation is to make the stars look seemingly cooler in the Hertzsprung-Russell diagram (De Ridder et al. 2023). Given that our classifier had no information on the temperature of targets, it is therefore not surprising they are found in the high-probability dataset. Considering their significant astero-seismic potential (Pedersen et al. 2021), we opted not to remove them from the candidate list. We finally note that the distribution of radii and luminosity for both labelled hybrids and the high-probability candidates follow each other closely.

We prepared a publicly available catalogue of high-probability candidates ($p(\text{hybrid}) > 0.8$) with TIC and *Gaia* DR3 identifiers, coordinates, TESS sector, highest frequencies in the low (f_g) and high (f_p) frequency regimes, their amplitudes (A_g and A_p , respectively), global properties (T_{eff} , L , R), and probabilities of being considered a hybrid (Table 1). Examples of light curves are provided in Fig. D.1. From our catalogue, there is an overlap of 1,619 individual sources labelled as hybrid and 464 as other classes, mainly p-mode pulsators, from the nominal mission catalogue by Hey & Aerts (2024) and an overlap of 21 pulsating eclipsing binaries from IJspeert et al. 2021 and Kemp et al. (subm.), of which 9 are labelled as g-mode, 10 p-mode, and 2 hybrid pulsating binaries. We additionally prepared high-probability candidates that passed the T_{eff} and f_g filtering but not the FAP threshold in one of the frequency regimes as respective catalogues of candidate p- (Table 2) and g-mode (Table 3) pulsators.

6. Conclusion

In this paper, we analysed TESS light curves of known γ Dor / δ Sct hybrid pulsators from datasets of confirmed hybrids and a large dataset of stars using TGLC light curves, resulting in a substantial increase in the available target pool of candidate hybrids. A comparison of TGLC light curves and their dominant and secondary periodic variability extracted from them generally agrees very well with other available light curves. TGLC light curves are at least as suitable as QLP light curves for hybrid pulsator studies. However, we note that the former have a magnitude limit that potentially allows to increase the current pool of available sources: 19.89% of our candidates are fainter than the magnitude limit of QLP. The two pipelines are both equally capable of recovering dominant *Kepler*-extracted frequencies for the labelled targets. Both are, however, better at detecting high-frequency signal than low frequencies, as single-sector TESS observations do not have enough frequency resolution to detect g-mode variability in an optimal way.

We performed a semi-supervised classification of hybrid pulsators on 29 million TGLC light curves with a PU Learning classifier.⁴ This classification is a challenging learning problem as it essentially combines two different goals — differentiating pulsating stars from a largely non-variable set and a specific class from other variable stars, including its parent classes. This is nor-

⁴ The code used in the framework of this study is available at https://github.com/nikkliapets/smartbinning_pulsators.




Fig. 8: Densities of f_1 (top) and A_1 (bottom) for hybrids in the labelled dataset (teal) and high-probability candidates (purple). Histograms are plotted from 1,000 samples for each candidate within the uncertainty range and a Kernel Density Estimator is plotted directly from point estimates of f_1 and A_1 . A burgundy dotted line represents a 4 d^{-1} frequency dividing the bimodal distribution into the g- and p-mode regimes.




Fig. 9: Amplitude-amplitude plot of the highest peak detected in frequency regimes $< 4 \text{ d}^{-1}$ (A_g) and $> 4 \text{ d}^{-1}$ (A_p) with a $< 1\%$ FAP of candidate hybrids with $p(\text{hybrid}) > 0.9$ (purple) and labelled hybrids (black).

maximally achieved by relying on different features. Our method additionally overcomes the need to have a representative labelled set in traditional supervised machine learning architectures. Our study demonstrated that by leveraging a small dataset of known rare objects, we are able to find many new candidates by using features specifically tailored to a class of interest and a machine learning method suitable for learning problems under class

imbalance. Such projects of looking for objects of a particular rare class can be seen as the first necessary step for population studies. The volume of data analysed in this project also demonstrated the practical aspects of employing automated pipelines in the age of space missions.

Our classification resulted in a catalogue of over 62,000 new candidate main-sequence hybrid pulsator light curves with significant peaks in both high- and low-frequency regimes (with additional catalogues of pure p- and g-mode pulsator candidates). These are ideal starting points for many asteroseismic follow-up studies. While our hybrid catalogue is dominated by hybrid candidates with higher p modes, it is an expected result of the TESS sampling rates, the composition of our labelled training set, and the occasional low-frequency instrumental systematics.

The population of detected hybrids comprises various sub-populations with distinct pulsation properties. Future exploitation of these stars can focus on the following:

- From a machine learning perspective: unsupervised clustering of the sample to reveal different sub-classes of hybrids, i.e. γ Dor / δ Sct hybrid, δ Sct / γ Dor hybrids, hybrids with rotational and magnetic splittings, hybrids with rotational modulation, etc.; and
- From an astrophysical perspective: exploitation of internal physical properties utilizing information provided by both p and g modes, particularly to measure differential rotation, mode coupling, and more precise measurements of asteroseismic ages and other astrophysical parameters from asteroseismology (Aerts 2021, Aerts & Tkachenko 2024).

Other types of hybrid pulsators (SPB / β Cep hybrids and p- and g-mode pulsating sdBV's) were not included in this study. However, our methodology is also applicable to these stars too. Moreover, we found B-type hybrids, as demonstrated by the distribution of T_{eff} of detected candidates. In this context, it is interesting to look into developing machine learning architectures that can differentiate different classes of hybrids (e.g. γ Dor / δ Sct hybrids from SPB / β Cep hybrids) by utilizing observables from other surveys, such as *Gaia* data (Brown et al. 2016) or including other modalities, i.e. spectroscopy, such as *Gaia* (Huijse et al. 2025) or SDSS-V (Kollmeier et al. 2019) spectra. Furthermore, additional efforts can be put into more precisely determining the probabilistic threshold of what is considered a hybrid, which would allow us to put a prior on the occurrence of these stars, opening avenues for using semi-supervised learning methods build on the SCAR assumption, which are normally easier to evaluate (Bekker & Davis 2020).

Finally, with the upcoming PLATO (Rauer et al. 2025) mission that will observe brighter stars with at least two year-long baseline, methods specifically tailored to the study of hybrids with dominant g modes are of particular interest to the variable star community. We stress that 4,283 of our high-probability candidates fall in PLATO's first field of view (Nascimbeni et al. 2025, Janssen et al. 2025), offering an opportunity for their asteroseismic exploitation to prepare for future PLATO data releases.

7. Data availability

Tables 1, 2, and 3 are only available electronically at the CDS: <https://cdsarc.cds.unistra.fr/viz-bin/cat/J/A+A/703/A240>.

Acknowledgements. MK and CA acknowledge The Kavli Foundation for their financial support in the framework of the Kavli Scholarship given to MK. CA and PH acknowledge financial support from the European Research Council (ERC) under the Horizon Europe programme (Synergy Grant agreement




Fig. 10: Distribution of fundamental parameters — T_{eff} (left), R (middle), and L (right) — of the final filtered sample of high-probability candidates (purple) and labelled hybrids (teal). Histograms are plotted from 1,000 samples for each candidate within the uncertainty range and a Kernel Density Estimator is plotted directly from point estimates of T_{eff} , R , and L .

N°101071505: 4D-STAR). While partially funded by the European Union, views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. AK, DJF, and CA acknowledge financial support from the Flemish Government under the long-term structural Methusalem funding programme by means of the project SOUL: Stellar evolution in full glory, grant METH/24/012 at KU Leuven. This paper includes data collected by the TESS mission. Funding for the TESS mission is provided by the NASA's Science Mission Directorate. This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

References

- Aerts, C. 2021, *RMP*, 93, 015001
- Aerts, C., Christensen-Dalsgaard, J., & Kurtz, D. W. 2010, *Asteroseismology* (Springer Science & Business Media)
- Aerts, C., Molenberghs, G., & De Ridder, J. 2023, *A&A*, 672, A183
- Aerts, C. & Tkachenko, A. 2024, *A&A*, 692, R1
- Aerts, C., Van Reeth, T., Mombarg, J. S., & Hey, D. 2025, *A&A*, 695, A214
- Aguirre, C., Pichara, K., & Becker, I. 2019, *MNRAS*, 482, 5078
- Audenaert, J. 2025, *Ap&SS*, 370, 1
- Audenaert, J., Kuszlewicz, J. S., Handberg, R., et al. 2021, *ApJ*, 162, 209
- Audenaert, J. & Tkachenko, A. 2022, *A&A*, 666, A76
- Bekker, J. & Davis, J. 2020, *Machine Learning*, 109, 719
- Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Science*, 327, 977
- Bowman, D. M., Kurtz, D. W., Breger, M., Murphy, S. J., & Holdsworth, D. L. 2016, *MNRAS*, 460, 1970
- Bradley, P. A., Guzik, J. A., Miles, L. F., et al. 2015, *ApJ*, 149, 68
- Brown, A. G., Vallenari, A., Prusti, T., et al. 2016, *A&A*, 595, A2
- Chan, H.-S., Villar, V. A., Cheung, S.-H., et al. 2022, *ApJ*, 932, 118
- Chapellier, E., Mathias, P., Weiss, W. W., Le Contel, D., & Debosscher, J. 2012, *A&A*, 540, A117
- Chen, W., Yang, K., Yu, Z., Shi, Y., & Chen, C. P. 2024, *Artificial Intelligence Review*, 57, 137
- Claesen, M., Davis, J., De Smet, F., & De Moor, B. 2015, *arXiv preprint arXiv:1504.06837*
- De Ridder, J., Ripepi, V., Aerts, C., et al. 2023, *A&A*, 674, A36
- Debosscher, J., Sarro, L., Aerts, C., et al. 2007, *A&A*, 475, 1159
- Du, C., Luo, A., Yang, H., Hou, W., & Guo, Y. 2016, *PASP*, 128, 034502
- Dupret, M.-A., Grigahcène, A., Garrido, R., Gabriel, M., & Scuflaire, R. 2005, *A&A*, 435, 927
- Eggenberger, P., Deheuvels, S., Miglio, A., et al. 2019, *A&A*, 621, A66
- Elkan, C. & Noto, K. 2008, in *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, 213–220
- Eyer, L., Audard, M., Holl, B., et al. 2023, *A&A*, 674, A13
- Fritzewski, D., Aerts, C., Mombarg, J., Gossage, S., & Van Reeth, T. 2024, *A&A*, 684, A112
- Grigahcène, A., Antoci, V., Balona, L., et al. 2010, *ApJ*, 713, L192
- Han, T. & Brandt, T. D. 2023, *ApJ*, 165, 71
- Handberg, R., Lund, M. N., White, T. R., et al. 2021, *ApJ*, 162, 170
- Hey, D. & Aerts, C. 2024, *A&A*, 688, A93
- Hey, D., Montet, B. T., Pope, B. J., Murphy, S. J., & Bedding, T. R. 2021, *ApJ*, 162, 204
- Hosenie, Z., Lyon, R., Stappers, B., Mootoovallu, A., & McBride, V. 2020, *MNRAS*, 493, 6050
- Howell, S. B., Sobeck, C., Haas, M., et al. 2014, *PASP*, 126, 398
- Huang, C. X., Vanderburg, A., Pál, A., et al. 2020a, *arXiv preprint arXiv:2011.06459*
- Huang, C. X., Vanderburg, A., Pál, A., et al. 2020b, *RNAAS*, 4, 204
- Hubrig, S., Jayaraman, R., & Järvinen, S. 2024, *RNAAS*, 8, 261
- Huijse, P., De Ridder, J., Eyer, L., et al. 2025, *A&A*, 701, A150
- Ijspeert, L. W., Tkachenko, A., Johnston, C., et al. 2021, *A&A*, 652, A120
- Jannsen, N., Tkachenko, A., Royer, P., et al. 2025, *A&A*, 694, A185
- Jaskie, K. & Spanias, A. 2019, in *2019 10th international conference on information, intelligence, systems and applications (iisa)*, IEEE, 1–8
- Kang, Z., Zhang, Y., Zhang, J., et al. 2023, *PASP*, 135, 094501
- Kollmeier, J., Anderson, S., Blanc, G., et al. 2019, *BAAS*
- Kunimoto, M., DeRocco, W., Smyth, N., & Bryson, S. 2024, *arXiv preprint arXiv:2404.11666*
- Kunimoto, M., Huang, C., Tey, E., et al. 2021, *RNAAS*, 5, 234
- Kunimoto, M., Tey, E., Fong, W., et al. 2022, *RNAAS*, 6, 236
- Kurtz, D. W. 2022, *RA&A*, 60, 31
- Li, F., Dong, S., Leier, A., et al. 2022a, *Briefings in bioinformatics*, 23, bbab461
- Li, G., Deheuvels, S., Ballot, J., & Lignières, F. 2022b, *Nat.*, 610, 43
- Li, G., Van Reeth, T., Bedding, T. R., et al. 2020, *MNRAS*, 491, 3586
- Lund, M. N., Handberg, R., Buzasi, D. L., et al. 2021, *ApJS*, 257, 53
- Mathis, S. 2009, *A&A*, 506, 811
- Mathis, S., Bugnet, L., Prat, V., et al. 2021, *A&A*, 647, A122
- Mombarg, J. S., Aerts, C., Van Reeth, T., & Hey, D. 2024, *A&A*, 691, A131
- Mordelet, F. & Vert, J.-P. 2014, *Pattern Recognition Letters*, 37, 201
- Müller, A. C. & Guido, S. 2016, *Introduction to machine learning with Python: a guide for data scientists* ("O'Reilly Media, Inc.")
- Murphy, S. J. 2012, *MNRAS*, 422, 665
- Nascimbeni, V., Piotto, G., Cabrera, J., et al. 2025, *A&A*, 694, A313
- Pang, G., Shen, C., Cao, L., & Hengel, A. V. D. 2021, *ACM computing surveys (CSUR)*, 54, 1
- Pedersen, M. G., Aerts, C., Pápics, P. I., et al. 2021, *Nat. Astron.*, 5, 715
- Ranaivomanana, P., Uzundag, M., Johnston, C., et al. 2025, *A&A*, 693, A268
- Rauer, H., Aerts, C., Cabrera, J., et al. 2025, *Exp. Astron.*, 59, 26
- Rebbapragada, U., Protopapas, P., Brodley, C. E., & Alcock, C. 2009, *Machine learning*, 74, 281
- Richards, J. W., Starr, D. L., Butler, N. R., et al. 2011, *ApJ*, 733, 10
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *JATIS*, 1, 014003
- Rizhko, M. & Bloom, J. S. 2024, in *Neurips 2024 Workshop Foundation Models for Science: Progress, Opportunities, and Challenges*
- Rogers, T. 2015, *ApJ*, 815, L30
- Sánchez-Sáez, P., Lira, H., Martí, L., et al. 2021, *ApJ*, 162, 206
- Skarka, M. & Henzl, Z. 2024, *A&A*, 688, A25
- Skarka, M., Žák, J., Fedurco, M., et al. 2022, *A&A*, 666, A142
- Solarz, A., Bilicki, M., Gromadzki, M., et al. 2017, *A&A*, 606, A39
- Tkachenko, A., Degroote, P., Aerts, C., et al. 2014, *MNRAS*, 438, 3093
- Uytterhoeven, K., Moya, A., Grigahcène, A., et al. 2011, *A&A*, 534, A125
- Van Beeck, J., Bowman, D. M., Pedersen, M. G., et al. 2021, *A&A*, 655, A59
- Van Engelen, J. E. & Hoos, H. H. 2020, *Machine learning*, 109, 373
- van Roestel, J., Kupfer, T., Ruiz-Carmona, R., et al. 2018, *MNRAS*, 475, 2560
- VanderPlas, J. T. 2018, *ApJS*, 236, 16
- Yang, P., Yang, G., Zhang, F., Jiang, B., & Wang, M. 2021, *Archives of Computational Methods in Engineering*, 28, 917
- Zhang, J., Zhang, Y., Kang, Z., et al. 2023, *PASA*, 40, e037
- Zhang, X., Chen, X., Zhang, H., Fu, J., & Li, Y. 2020, *ApJ*, 895, 124

Appendix A: Frequency comparisons




Fig. A.1: Comparison of f_1 (top) and f_2 (bottom) of confirmed hybrids in extended mission TGLC with extended mission QLP (left), nominal mission TGLC (middle), and extended mission TGLC downsampled to nominal mission quality (right). The black line indicates unity and two grey lines are double and half unity. The colour bar represents $\log_{10}$ amplitude ratios of the two compared observations. Only sources with frequencies of up to 25 d^{-1} are shown for visibility.

Appendix B: PU learning schematic




Fig. B.1: Graphic representation of a 2D projection of a PU Learning classifier (burgundy) looking for unobserved positive instances within a larger dataset of unlabelled instances (grey) with the help of a smaller labelled set (teal). Datasets not to scale.

Appendix C: Harmonic behaviour




Fig. C.1: Period-period plot of candidate hybrids with $p(\text{hybrid}) > 0.8$ (left) and $p(\text{hybrid}) > 0.95$ (right) for the inverse of the highest peak detected in frequency regimes $< 4 \text{ d}^{-1}$ (f_g) and $> 4 \text{ d}^{-1}$ (f_p) with a $< 1\%$ FAP without harmonic avoidance criterion. Note the diagonal lines where P_p is an integer multiple of P_g . The colour bar represents $\log_{10}$ amplitude ratios of the two periods.




Fig. C.2: Light curve (left), periodogram (middle), and zoomed-in periodogram (right) of a single source where $P_p = 2 \times P_g$, $P_p = 3 \times P_g$, $P_p = 4 \times P_g$, and $P_p = 5 \times P_g$, respectively.

Appendix D: Examples of candidates




Fig. D.1: Light curve (left), periodogram (middle), and zoomed-in periodogram (right) of ten high-probability candidates of being a hybrid pulsator. Note the different Nyquist limit for nominal (rows 1, 5-10) and extended (rest) mission candidates.