



HAL
open science

Searching for changing-state AGNs in massive datasets – I: applying deep learning and anomaly detection techniques to find AGNs with anomalous variability behaviours

P. Sánchez-Sáez, Hernan Lira, Luis Martí, Nayat Sanchez-Pi, J. Arredondo, F. E. Bauer, A. Bayo, G. Cabrera-Vives, C. Donoso-Oliva, P. A. Estévez, et al.

► To cite this version:

P. Sánchez-Sáez, Hernan Lira, Luis Martí, Nayat Sanchez-Pi, J. Arredondo, et al.. Searching for changing-state AGNs in massive datasets – I: applying deep learning and anomaly detection techniques to find AGNs with anomalous variability behaviours. *The Astronomical Journal*, 2021, 162 (5), 10.3847/1538-3881/ac1426 . hal-03294849

HAL Id: hal-03294849

<https://inria.hal.science/hal-03294849>

Submitted on 20 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



Searching for Changing-state AGNs in Massive Data Sets. I. Applying Deep Learning and Anomaly-detection Techniques to Find AGNs with Anomalous Variability Behaviors

P. Sánchez-Sáez^{1,2,3} , H. Lira¹ , L. Martí¹ , N. Sánchez-Pi¹ , J. Arredondo² , F. E. Bauer^{3,4,2,5} , A. Bayo^{6,7} ,
G. Cabrera-Vives^{8,2} , C. Donoso-Oliva^{8,2} , P. A. Estévez^{9,2} , S. Eyheramendy^{10,2} , F. Förster^{11,2,12,13} ,
L. Hernández-García^{2,6} , A. M. Muñoz Arancibia^{2,12} , M. Pérez-Carrasco^{8,2} , M. Sepúlveda¹³ , and J. R. Vergara^{14,2} 

¹ Inria Chile Research Center, Av. Apoquindo 2827, Las Condes, Chile; pasanchezsaez@gmail.com

² Millennium Institute of Astrophysics (MAS), Nuncio Monseñor Sótero Sanz 100, Providencia, Santiago, Chile

³ Instituto de Astrofísica, Facultad de Física, Pontificia Universidad Católica de Chile, Casilla 306, Santiago 22, Chile

⁴ Centro de Astroingeniería, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, 7820436 Macul, Santiago, Chile

⁵ Space Science Institute, 4750 Walnut Street, Suite 205, Boulder, CO 80301, USA

⁶ Instituto de Física y Astronomía, Facultad de Ciencias, Universidad de Valparaíso, Av. Gran Bretaña 1111, Valparaíso, Chile

⁷ Núcleo Milenio de Formación Planetaria (NPF), Valparaíso, Chile

⁸ Department of Computer Science, Universidad de Concepción, Edmundo Larenas 219, Concepción, Chile

⁹ Department of Electrical Engineering, Universidad de Chile, Av. Tupper 2007, Santiago 8320000, Chile

¹⁰ Faculty of Engineering and Sciences, Universidad Adolfo Ibañez, Diagonal Las Torres 2700, Peñalolén, Santiago, Chile

¹¹ Data and Artificial Intelligence Initiative (ID&IA), University of Chile, Santiago, Chile

¹² Center for Mathematical Modeling, University of Chile, AFB170001, Chile

¹³ Departamento de Astronomía, Universidad de Chile, Casilla 36D, Santiago, Chile

¹⁴ Department of Computing, Universidad Tecnológica Metropolitana, Santiago, Chile

Received 2021 June 11; revised 2021 July 7; accepted 2021 July 12; published 2021 October 20

Abstract

The classic classification scheme for active galactic nuclei (AGNs) was recently challenged by the discovery of the so-called changing-state (changing-look) AGNs. The physical mechanism behind this phenomenon is still a matter of open debate and the samples are too small and of serendipitous nature to provide robust answers. In order to tackle this problem, we need to design methods that are able to detect AGNs right in the act of changing state. Here we present an anomaly-detection technique designed to identify AGN light curves with anomalous behaviors in massive data sets. The main aim of this technique is to identify CSAGN at different stages of the transition, but it can also be used for more general purposes, such as cleaning massive data sets for AGN variability analyses. We used light curves from the Zwicky Transient Facility data release 5 (ZTF DR5), containing a sample of 230,451 AGNs of different classes. The ZTF DR5 light curves were modeled with a Variational Recurrent Autoencoder (VRAE) architecture, that allowed us to obtain a set of attributes from the VRAE latent space that describes the general behavior of our sample. These attributes were then used as features for an Isolation Forest (IF) algorithm that is an anomaly detector for a “one class” kind of problem. We used the VRAE reconstruction errors and the IF anomaly score to select a sample of 8809 anomalies. These anomalies are dominated by bogus candidates, but we were able to identify 75 promising CSAGN candidates.

Unified Astronomy Thesaurus concepts: [Active galactic nuclei \(16\)](#); [Astronomy data analysis \(1858\)](#); [Surveys \(1671\)](#); [Interdisciplinary astronomy \(804\)](#)

1. Introduction

Active galactic nuclei (AGNs) are powered by the release of gravitational energy related to the accretion of material onto a supermassive black hole (SMBH; Lynden-Bell 1969). They are characterized by their time-variable emission, and these variations can be observed from gamma rays to radio wave bands. The light curves of AGNs appear to be stochastic and have characteristic timescales that range from hours to years (e.g., Ulrich et al. 1997; Padovani et al. 2017).

Changing-state AGNs (CSAGNs; sometimes called optical changing-look AGNs), correspond to AGNs that change their classification as type-1 [objects that present broad permitted emission lines ($\text{FWHM} \gtrsim 2000 \text{ km s}^{-1}$) in their spectra] or type-2 AGNs (objects that lack broad emission lines in their spectra), as well as to sources that present large changes in the flux of their broad emission lines. These changes are normally observed on timescales between months and years (e.g., LaMassa et al. 2015; MacLeod et al. 2016; Ross et al. 2018; Stern et al. 2018; MacLeod et al. 2019; Graham et al. 2020). This transition phase can be accompanied by a drastic change

in the AGN continuum flux by orders of magnitude which is unexpected considering the typical variability amplitudes observed in large AGN samples (~ 0.1 – 0.2 magnitudes per year; e.g., Vanden Berk et al. 2004; Sánchez-Sáez et al. 2018).

The origin of the anomalous optical variability of CSAGNs is still not well understood. Possible physical explanations include changes in obscuration (which produce variable absorption), or changes in the black hole accretion rate or accretion disk structure. Recent observations indicate that this phenomenon (in the optical range) is not due to obscuration since this would imply large cloud size and rapid velocities required to occult a sizable portion of the disk but is rather due to changes in the innermost regions of the accretion disk (e.g., LaMassa et al. 2015; Ross et al. 2018; Stern et al. 2018; MacLeod et al. 2019). Understanding the changing-state phenomenon is therefore crucial to improve our knowledge of the physics behind AGN variability. In particular, we need to understand to what extent the variability properties of CSAGNs are generally outliers from the overall distribution of AGNs (i.e., even before they exhibit their strong changes).

Most of the few known CSAGNs have been discovered “after the fact,” using archival data, which does not allow the possibility of witnessing the change/following it up, in a multiwavelength manner, and thus, innovative strategies are needed to catch them during the transition phase. Ricci et al. (2020) presented observations of the disappearance and re-appearance of the X-ray corona in 1ES 1927+654, a source that was previously classified as a CSAGN by Trakhtenbrot et al. (2019a). This event is particularly interesting because it is a CSAGN both in the optical and X-ray regimes. The authors claim that a tidal disruption event (TDE; tidal disruption of a star by the accreting SMBH) provoked an increase in the accretion rate at the innermost regions of the accretion disk, which emptied the inner disk and led to the destruction of the X-ray corona. This corresponds to the first CSAGN extensively monitored at X-rays and optical through a transition state, placing strong constraints on the physics of the system. That work demonstrated the importance of the rapid detection and rapid follow up of changing-state events, which can give us the opportunity to probe accretion physics.

Previous searches of CSAGN in massive archival data sets have employed simple statistical analysis to find them, using for instance measurements of the variability amplitude (e.g.; MacLeod et al. 2016, 2019) or more complex analysis like Bayesian blocks representation (BB; Scargle et al. 2013; Graham et al. 2020). However, these methods are not suitable for the real-time detection of CSAGN events since they favor the detection of events well after they have occurred and not the detection of the transition phase. More promising techniques involve time series forecasting or the use of anomaly-detection algorithms (AD; e.g., Mehrotra et al. 2017).

AD techniques have been used in astronomy to search for unusual objects within astronomical data sets, however, this search has largely focused to date on transient and periodic classes (e.g., Ishida et al. 2021; Lochner & Bassett 2021; Malanchev et al. 2021; Villar et al. 2021). M. Pérez-Carrasco et al. (2021, in preparation) present the outlier detector for the Automatic Learning for the Rapid Classification of Events (ALeRCE) broker (Carrasco-Davis et al. 2020; Förster et al. 2021; Sánchez-Sáez et al. 2021). Its main goal is to find multiple types of variable and transient sources that are not included in the taxonomy tree of the ALeRCE broker light-curve classifier (Sánchez-Sáez et al. 2021). However, AD algorithms can also be used to detect light curves of objects with previous known classifications, which suddenly start presenting unusual variability behaviors (also known as contextual AD). Recently, Suberlak et al. (2021) showed that CSAGN candidates appear as outliers when they are compared with the damped random walk (DRW) parameters (timescale and variance; Kelly et al. 2009) obtained from light curves of the same objects that cover different timespans (and using different data sets). Thus, we can expect that the early detection of CSAGNs would benefit from the use of AD techniques on current large monitoring surveys.

In this work we present an unsupervised AD algorithm to search for AGN light curves with anomalous variability behaviors in massive data sets, focusing in particular on data from the Zwicky Transient Facility (ZTF; Bellm et al. 2019). The main goal of the proposed algorithm is to find CSAGN candidates at different stages of the changing-state event (i.e., either early or late stages), although the methodology should be generally useful to find other classes of anomalous light curves, like atypical flaring

activity (e.g., Graham et al. 2017; Trakhtenbrot et al. 2019b; Frederick et al. 2020), extremely variable AGNs (e.g., Rumbaugh et al. 2018; Guo et al. 2020; Luo et al. 2020), or sources with incorrect labels or photometric issues, as we show in the following sections. Our method is inspired by the work presented in Tachibana et al. (2020), who used a Recurrent Autoencoder (RAE) architecture to model AGN light curves observed by the Catalina Real-time Transient Survey (CRTS; Drake et al. 2009). We instead use a Variational Recurrent Autoencoder (VRAE; Fabius & van Amersfoort 2014) architecture together with an Isolation Forest (IF; Liu et al. 2008) algorithm, following a two-stage approach, to search for anomalous light curves. As far as we know, this is the first attempt to search for anomalous AGN variability behaviors using AD techniques.

The paper is organized as follows. In Section 2 we describe the data sets used for this work, including the AGN catalogs considered, and the procedure to define the final sample of ZTF light curves. In Section 3 we describe the VRAE architecture and the AD methodology used. In Section 4 we present a comparison of the results obtained using RAE and VRAE architectures. We show the results obtained using our two-stage AD method, together with a description of the anomalies found. In Section 5 we discuss our main findings and compare our results with previous works. Finally, in Section 6 we summarize and conclude the main results obtained in this work.

2. Data

2.1. The Sample of Known AGNs

In order to detect CSAGNs in real time, we need a sample of AGNs with known classifications. In particular, we need sources optically classified as type 1 or type 2. However, in order to test the quality of our AD algorithm, we decided to include other classes of AGNs such as blazars, X-ray-detected AGNs, and radio-detected AGNs. We used the following AGN catalogs to create our sample of known AGNs:

1. The Roma-BZCAT Multi-Frequency Catalog of Blazars (ROMABZCAT; Massaro et al. 2015; priority 1).
2. The New Catalog of Type-1 AGNs (Oh2015; Oh et al. 2015; priority 2).
3. The Million Quasars (MILLIQUAS; priority 3) Catalog, Version 7.0 (Flesch 2015, 2019).

These catalogs were used to obtain celestial coordinates and classifications of known AGNs. We considered the following 12 classes (labels used in this work are provided in parentheses):

1. Seyfert 1 or host-dominated type-1 AGN (A).
2. Seyfert 1 with radio emission (AR).
3. QSO 1 or core-dominated type-1 AGN (Q).
4. QSO 1 with radio emission (QR).
5. Host- and core-dominated type-2 AGN (type2).
6. Type-2 AGN with radio emission (type2 R).
7. BL Lac AGN (BLLac).
8. Flat Spectrum Radio Quasar (FSRQ).
9. Blazars of uncertain type (BZU).
10. AGN identified by its X-ray emission (X).
11. AGN identified by its radio emission (R).

If a source appeared in more than one catalog, we followed the priorities defined above to select the final label. For instance, if a source appears in ROMABZCAT and in MILLIQUAS we

kept the classification provided by ROMABZCAT. After this, we ended up with a sample of 875,322 known AGNs (sources classified as candidates in MILLIQUAS were not included).

In addition, we obtained physical properties (redshifts, SMBH masses, Eddington ratios, and luminosities) from the Sloan Digital Sky Survey (SDSS; York et al. 2000) for 524,454 sources in the sample of known AGNs using the Oh2015 catalog and the catalog of Spectral Properties of Quasars from SDSS Data Release 14 (Rakshit et al. 2020). This information was used to define training and validation sets balanced by means of their physical properties (see Section 2.3).

2.2. Optical Light Curves

For this work we used the point-spread function (PSF) fit-based optical light curves available in the ZTF public data release 5¹⁵ (ZTF DR5; Masci et al. 2019). In particular, we used the ZTF Application Programming Interface (ZTF API) to obtain light curves for all the sources in our sample of known AGNs described in Section 2.1. We retrieved the light curves of each source, centered in the coordinates of the catalogs presented in Section 2.1, and using a radius of 1.5 arcseconds, which is the standard cross-matching radius used by ZTF. When doing this for a particular object, the ZTF API can return more than one light curve in a given filter. This happens because ZTF treats independently the light curves observed in a particular field, filter, and CCD-quadrant; thus, if a source appears in more than one combination of these three, ZTF will provide more than one light curve. The ZTF quadrants are calibrated independently, thus combining light curves from different fields and CCD-quadrants (but same filter), can produce spurious variability (van Roestel et al. 2021). Therefore, to avoid having multiple light curves for any particular source, we only used the longest light curve (in terms of its timespan) of each filter of a particular object.

ZTF DR5 provides light curves in the g , r , and i ZTF bands, however, for this work we only used light curves observed in the g band. We expect that this band presents less contamination from the host galaxy of each target. Future work will further exploit the potential of using multiband light curves for AD.

We cleaned the ZTF DR5 light curves using the `catflags` quality score, keeping only the epochs with `catflags = 0`, as advised by the ZTF documentation. We also removed all the epochs with a photometric error higher than 1 magnitude.

The ZTF PSF-fit-based light curves are provided for point-like and extended sources, however, PSF-fit photometry measurements are proper for point sources only. Therefore, we used the catalog presented in Tachibana & Miller (2018), which provides a morphological classification score (`ps_score`) using Pan-STARRS1 photometry, to filter out extended sources from our analysis, keeping only sources with `ps_score > 0.5` (i.e., point sources according to Tachibana & Miller 2018).

In addition, to ensure the detection of AGN variability (with timescales of hours to years; Ulrich et al. 1997), we filtered the sample by the timespan (t_{length}) and number of epochs (n_{det}) of the light curves, keeping only sources whose light curves have $t_{\text{length}} \geq 730$ days (2 yr) and $n_{\text{det}} \geq 50$. We also removed from the sample sources with an average g -band magnitude fainter




Figure 1. Number of sources per class in the full sample.

than 20.6 (close to the limiting magnitude of the ZTF images) and brighter than 13.5 (to avoid saturated observations).

Finally, we computed two variability features; P_{var} , the probability that the source is intrinsically variable, and σ_{rms}^2 , the normalized excess variance (for a detailed description, see Appendix A of Sánchez-Sáez et al. 2021), and used them to remove from the sample light curves that do not show any variations. Following a similar approach to Sánchez et al. (2017), we removed from the analysis sources with $P_{\text{var}} < 0.95$ and $\sigma_{\text{rms}}^2 < 0$. After this, we ended up with a sample of 230,451 known AGNs with variable (and long enough) ZTF light curves (hereafter, the “full” sample), with 166,243 of these having spectroscopic properties measured from SDSS (hereafter, the “full spectroscopic” sample).

Figure 1 shows the number of sources per class in the full sample. It is dominated by sources of class “Q,” which correspond to 90.7% of the total.

2.3. Training, Validation, and Testing Sets

It is well known that AGNs with different physical properties show different variability behaviors (e.g., MacLeod et al. 2010; Hernández-García et al. 2016; Caplar et al. 2017; Sánchez-Sáez et al. 2018). For instance, SMBH mass (BH_{mass}) and accretion rate set the luminosity variability properties to first order. Distance sets the flux limit, with rare sources only found in large volumes, and faint sources only found in small volumes. Finally, past follow up and multiwavelength coverage may include its own set of biases. Thus, in order to develop an AD algorithm that is not biased toward the more common physical properties in the sample, we need to train our algorithm using training and validation sets that are balanced by means of these physical properties. Otherwise, sources with physical properties that escape the general distribution will appear as anomalous, even though they do not show anomalous variability behaviors.

The left panel of Figure 2 shows BH_{mass} and the bolometric luminosity (L_{bol}) of the full spectroscopic sample (166,243 sources). From the figure we can see that most of the sources cluster around $\text{BH}_{\text{mass}} \sim 10^9$ and $L_{\text{bol}} \sim 10^{46.5}$. In order to have homogeneous training and validation sets, we defined two subsamples (of the full sample with spectral properties) that are balanced by means of their BH_{mass} and L_{bol} .

In addition, the number of epochs per light curve covers a wide range in our sample, with $50 \leq n_{\text{det}} \leq 935$. Figure 3 shows in black the histogram of n_{det} per source for the full sample. It can be seen that most of the light curves have less than 200 epochs and only a few have more than 800 epochs.

¹⁵ Documentation: <https://www.ztf.caltech.edu/page/dr5>.




Figure 2. Bivariate histograms of the black hole mass (BH_{mass}) and the bolometric luminosity (L_{bol}) of the full spectroscopic sample (left) and the balanced sample (right).




Figure 3. Histogram of the number of epochs (n_{det}) per light curve for the full sample (black) and the balanced sample (purple).

This can produce some issues for our AD algorithm (see discussion in Sections 3.1 and 4.1) and therefore, we decided to balance one of the balanced sub-samples by the number of epochs as well.

Accordingly, we constructed two training and validation sets, one that is balanced only by means of the physical properties (hereafter, the “balanced phys” sample), and another one balanced by means of the physical properties and n_{det} (hereafter, the “balanced phys-epochs” sample).

To construct the balanced phys sample, we defined a grid in the space of BH_{mass} and L_{bol} , where each bin had a width of 0.1 dex, and selected a random sample of 40 targets from each bin (if the bin has less than 40 targets, we kept all the sources). From this, we ended up with a sample of 26,614 objects, with 88% of them belonging to class Q.

By contrast, to construct a balanced phys-epochs sample, we defined a grid in the space of n_{det} , BH_{mass} , and L_{bol} , where each bin had a width of 0.1 dex for BH_{mass} and L_{bol} , and a width of 20 for n_{det} , and selected a random sample of 2 targets from each bin (if the bin has less than 2 targets, we kept all the sources).

By doing this, we ended up with a balanced subset of 24,755 sources. According to the labels of the original catalogs, 95% of this balanced sample is classified as Q. The right panel of Figure 2 shows the distribution of BH_{mass} , and L_{bol} for the balanced phys-epochs sample and Figure 3 shows in purple its histogram of n_{det} . From both figures we can see that this subset is more balanced in its physical and observational properties compared to the original full set. For the case of the balanced phys sample, the distributions of the physical properties are similar to those shown in Figure 2 for the balanced phys-epochs sample.

The balanced phys set is used only for testing purposes, while the balanced phys-epochs set is used to train the final AD model, as explained in the following sections. We used 20% of each balanced sample as a validation set and 80% as the training set.

3. Methods

In this work we followed an unsupervised approach (i.e., without considering the labels of our data set) to search for anomalous AGN variability behaviors. We followed a two step approach, where we first train a VRAE architecture to reconstruct the light curves in our sample, and used the reconstruction error as an AD score. Then, we used the VRAE mean latent space as a vector of features for an IF algorithm and selected anomalies using the anomaly score provided by the IF method. In the following section we provide a detailed description of our AD methodology.

3.1. VRAE Architecture

Autoencoders (AEs) are artificial neural networks that are able to learn compressed representations of an input data set, called “latent representations,” in an unsupervised way. The latent representations are then used to reconstruct or predict the original data. Basically, AEs are composed by an encoder set of layers that compress the data, a latent space that contains the compressed information, and a decoder set of layers that




Figure 4. Proposed VRAE architecture. The model receives as input Δt , normalized magnitude, and the normalized magnitude error, and outputs a prediction of the normalized magnitude for the original Δt vector.

produces a reconstruction or prediction. RAEs, on the other hand, are recurrent AEs that use recurrent neural networks (RNN) to encode and decode sequential data and are normally used for sequence-to-sequence learning.

Tachibana et al. (2020) presented a novel algorithm to empirically model AGN light curves using an RAE (e.g., Naul et al. 2018) architecture. One of the main advantages of using this kind of technique is that it does not assume any underlying statistical process when fitting the light curves (as, for instance, DRW modeling does; e.g., Kelly et al. 2009), but instead, it follows a nonparametric approach that learns by itself the most important attributes that describe the data. This allows the algorithm to learn properties that might be hard for any expert to recognize, characterize, and/or parameterize within a model (statistically or physically motivated).

Instead of using an RAE architecture, we adopt a VRAE architecture to extract a set of latent attributes or features from the AGN light curves since it has been demonstrated that variational autoencoders (VAEs; Kingma & Welling 2013) are more suitable for AD tasks (e.g., Portillo et al. 2020; Villar et al. 2021).

VAEs¹⁶ correspond to a modification of the more classical AE architectures. In the case of VAE architectures, the latent representations are described by multivariate normal distributions, where each attribute or feature in the latent space is described by a latent mean (μ) and a latent variance (σ^2), which can be used to randomly sample a set of attributes. Then, we can use this sampled vector of latent attributes to generate a prediction or reconstruction of the input data. Since we can generate multiple latent vectors from a given distribution of μ and σ^2 , we can use VAEs as generative models and generate (or sample), for instance, new light curves that can be used for data augmentation (e.g., Moreno-Barea et al. 2020). By constructing the latent space in this way, we are also forcing the latent space to be smooth and to cluster together subsets with similar reconstructions (e.g., light curves with similar shapes or properties are closer in the latent space), thus facilitating the detection of anomalous data.

VRAEs (Fabius & van Amersfoort 2014) are RNNs based on VAEs that encode sequential data into a variational distribution over latent variables. One of the main advantages of using

RNNs is that they can deal with time series of different lengths and thus are ideal for light-curve analyses. Tachibana et al. (2020) used long short-term memory (LSTM; Hochreiter & Schmidhuber 1997) layers, while we use gated recurrent-unit (GRU; Cho et al. 2014) neurons instead, as they provide similar results and are faster to train (Yang et al. 2020).

Figure 4 shows the VRAE architecture used in this work. It is similar to the one used by Tachibana et al. (2020), but, as already mentioned, instead of an AE we used a VAE, and instead of LSTM units we used GRUs. In order to deal with the unevenly sampled ZTF light curves, we used the following input values:

1. Δt : differences between sampling times, normalized by 365 days (to use units of years).
2. Normalized magnitude (m_{norm}): magnitude normalized by the average magnitude and the standard deviation of the light curve [$m_{\text{norm}} = (m - \bar{m}) / \text{std}(m)$].
3. Normalized magnitude errors (err_{norm}): photometric errors (err) normalized by the standard deviation of the light curve [$\text{err}_{\text{norm}} = \text{err} / \text{std}(m)$].

We used two bidirectional GRU layers (BGRU¹⁷) for the encoder and two GRU layers for the decoder, each one with 32 units, with dropout of 0.25, and with linear activation functions. The latent space mean and variance correspond to dense layers (regular, fully-connected neural networks) of size 16, from which a sampled latent vector of size 16 is randomly generated. In the initial layers of the decoder, the sampled latent vector is repeated (in a copy-paste manner) N_T times, with N_T being the number of epochs in the light curves, which in this case, we define as the maximum number of Δt values in the given batch being evaluated (934 when using the full sample at once), and the values of each Δt are concatenated to this repeated vector (for shorter light curves the missing Δt values are replaced by zero). This is done to allow interpolation and/or extrapolation of the light curves if needed.¹⁸ This

¹⁷ Where a sequence is passed through a forward GRU and through a backward GRU and the outputs of the two are used to compute the output of the layer (e.g., Chaini & Kumar 2020).

¹⁸ This particular property of our architecture is not used in this work, but it will be used in future analyses and will allow us to do, for instance, light-curve forecasting.

¹⁶ Portillo et al. (2020) provide an astronomer-friendly description of the mathematics behind AEs and VAEs.

concatenated matrix is then the input of the decoder GRU layers. Finally, we end up with a vector of predicted (or reconstructed) normalized magnitudes for each input light curve (m_{rec}).

VAEs in general are trained using a loss function that includes two terms, one that takes into account the reconstruction error, called reconstruction loss, and another one, called latent loss, that forces the latent space to look as if it was sampled from a prior distribution, which in this case, we assume to be a normal distribution. Thus, for the reconstruction loss we used the weighted mean-squared error where the weights correspond to the inverse of the normalized photometric errors ($w = 1/\text{err}_{\text{norm}}$) and for the case of the latent loss, we used the standard Kullback–Leibler (KL) divergence:

$$\mathcal{L} = \frac{1}{N_{lc}} \sum_k \sum_i^{N_{lc}} [(m_{\text{norm},i} - m_{\text{rec},i})w_i]^2 - \frac{1}{200} \frac{1}{N_{lc}} \sum_k \sum_j^{N_{lc}} 0.5[1 + \log(\sigma_j^2) - \sigma_j^2 - \mu_j^2] \quad (1)$$

with N_{lc} being the number of light curves in the batch and N_l the dimensionality of the latent space. Note that we divide the latent loss by 200. This is done to ensure that the latent and reconstruction losses have similar scales since, when we compute the reconstruction loss, we are computing the mean over the total number of epochs of each light curve. We selected 200 because it is a representative number of the total number of epochs of the light curves. This number can be tuned as a hyper-parameter (e.g., β -VAE; Higgins et al. 2017), but we leave that for future work, as good results are achieved by assuming this particular value.

The model was trained in Google Colaboratory (Bisong 2019), taking advantage of the graphics processing units (GPUs) offered by its serverless Jupyter notebook environment. We used Keras (Chollet et al. 2015) with a TensorFlow backend (Abadi et al. 2016). We minimize the loss function using the Adam optimizer (Kingma & Welling 2013), with standard parameter values ($\beta_1 = 0.9$ and $\beta_2 = 0.999$), a learning rate of 10^{-3} , and a batch size of 512. We tested modifications of the original proposed architecture (e.g., different number of units or layers), but we did not observe any relevant improvements in the results, thus we decided to keep the original number of units and layers. Finally, to train our architecture, we used the balanced sets described in Section 2.3.

3.2. AD Methodology

A common way to do time series AD with RAEs or VRAEs is to train a model with a training set that is assumed to be normal and then use this model to reconstruct the data of a test set and measure the reconstruction error. If the reconstruction error is above a given threshold, then the time series is labeled as anomalous (e.g., Malhotra et al. 2016; Pereira & Silveira 2018).

Another option is to use the RAE or VRAE latent space as a feature for an AD algorithm, such as IF. IF is a tree-based ensemble method that isolates anomalies by making a random selection of features and dividing the data set according to randomly selected thresholds. In general, anomalies tend to be isolated in fewer steps compared to normal objects, thus an anomaly score can be defined as proportional to the number of splits required to isolate a given object. IF has been used

successfully for AD of astronomical time series (e.g., Ishida et al. 2021; Malanchev et al. 2021; Villar et al. 2021).

In this work we followed both approaches to select anomalous AGN light curves: we used the VRAE model to select anomalous light curves by means of their reconstruction error and we used the latent space as features of an IF algorithm.

We first trained our VRAE model using the balanced phys-epochs set (the balanced phys set is used for testing purposes only). We assume that in general, the balanced phys-epochs set is composed of normal light curves, as it mostly includes type-1 QSOs. Some of the light curves in the balanced phys-epochs set could be anomalous, if they were, for instance, associated with CSAGN events or there were problems in the ZTF PSF-fit-based photometry. However, we do not expect that these anomalous light curves correspond to a high fraction of the balanced phys-epochs set since we are filtering the light curves using the `catflags` quality score and since CSAGN events are very rare.

After training, we apply the VRAE model to the full sample (230,451 sources). This sample contains not only regular type-1 QSOs, but also blazars and Seyferts that can be used to validate the efficiency of our algorithm selecting anomalous light curves. We expect that blazars have light curves that are intrinsically different from the light curves of regular type-1 QSOs since the optical emission is dominated by emission from the relativistic jet, and rapid time-variable beaming effects can be present (Netzer 2013). Seyferts, however, have optical emission that corresponds to a combination of the flux coming from the active nucleus and flux coming from the host, thus, the optical variable signal of Seyferts may look quite different from the variable signal of regular QSOs if the contribution of the host galaxy is considerably high. They may also have systematically lower $\text{BH}_{\text{mass}} (< 10^7)$ and lower accretion rates, which could lead to more rapid fluctuations (which are then damped by the host galaxy).

To measure the reconstruction error, we sampled 10 different latent vectors and generated 10 VRAE reconstructions for each light curve, then computed the average reconstructed normalized light curve [$\overline{m_{\text{rec}}}$] and the standard deviation of the reconstruction [$\text{std}(m_{\text{rec}})$]. Since the inputs of the VRAE model are the normalized light curves, we have to denormalize the VRAE reconstructions by applying the same average magnitude and standard deviation used when we computed m_{norm} and err_{norm} . Thus, we compute the final light-curve prediction (m_{pred}) and its error (err_{pred}) as:

$$m_{\text{pred}} = \overline{m_{\text{rec}}} * \text{std}(m) + \overline{m}, \quad (2)$$

$$\text{err}_{\text{pred}} = \text{std}(m_{\text{rec}}) * \text{std}(m). \quad (3)$$

We used as reconstruction error (R) the mean-squared weighted deviation (MSWD; also known as the reduced chi-square statistic)

$$R = \frac{1}{N_T} \sum_i^{N_T} \frac{(m_i - m_{\text{pred},i})^2}{\text{err}_i^2 + \text{err}_{\text{pred},i}^2}. \quad (4)$$

Figure 5 shows the reconstruction error for the sources in the balanced phys-epochs set. Fourteen percent of the sources in this sample have $R < 1$, with most of them (79%) having average ZTF photometric errors larger than 0.1, thus their light curves are dominated by noise, which explains the low R values. The peak of R is around ~ 1.25 , and only a small




Figure 5. Histogram of the reconstruction error (R) per light curve for the balanced phys-epochs set. The threshold of $R = 3$ is shown with a black dashed line.

fraction (2%) have $R \geq 3$. This is not surprising, as we can expect that a small fraction of the balanced phys-epochs set present anomalous behaviors. Thus, we used $R = 3$ as our threshold to label a light curve as anomalous. Note that this value can be modified as desired to reduce or increase the number of anomalies to be presented to the final user of the method.

In a second stage, we used the latent mean (μ) as a vector of features and selected anomalous sources with IF, using the implementation available in `scikit-learn` (Pedregosa et al. 2011) with 500 base estimators and a contamination of 2%. The IF model was trained with the balanced phys-epochs set, then we applied the model to the full sample. The `scikit-learn` implementation of IF provides an anomaly score (which is proportional to the mean depth of the leaf containing the observation) where the lower the value, the more abnormal the source is. We computed the IF scores (IF_{score}) and used the threshold value $\text{IF}_{\text{th}} = -0.57633$ (defined for a contamination of 2% in the training set) returned by the model to label the light curves as anomalous. As for the case of R , this threshold value can be modified to change the number of anomalies presented to the final user, and we used here a fixed value to ease the analysis of the results obtained using our method.

Hence, following this two-stage approach, we label a light curve as anomalous if its reconstruction error is $R \geq 3$ or if it has an IF score smaller than the threshold ($\text{IF}_{\text{score}} \leq \text{IF}_{\text{th}}$). This is done because the IF scores obtained for light curves with large reconstruction errors are not trustworthy since the latent space vectors of these light curves do not compress properly the information associated with the real variability signal.

4. Results

4.1. Comparison of the Performance of the VRAE and RAE Architectures


At the beginning of this project, we began by evaluating the anomaly-detection performance of an RAE architecture similar to the one presented in Tachibana et al. (2020), using the same number of layers and units but excluding the forecasting step. In our initial tests of that RAE architecture, we used the balanced phys set to train the model. When we applied this model to the full sample, we noticed a high correlation between

the values of the attributes of the latent space (RAE e_i , with i the number of the attribute) and the number of epochs per light curve. Panel (a) of Figure 6 (purple) shows the value of each attribute versus the number of epochs per light curve (n_{det}). The Spearman rank coefficient (ρ_s) is provided on top of each subplot. It is clear from the subplots and from the correlation coefficient ρ_s that there is considerable correlation between some attributes and n_{det} , particularly for RAE e_1 , RAE e_4 , RAE e_5 , RAE e_{10} , and RAE e_{14} . This effect could produce problems when we use the vector of attributes as features for an AD algorithm. When we tried to use this vector of features for the IF algorithm, most of the light curves labeled as anomalous had a large number of epochs. Panel (a) in Figure 7 (purple) shows the IF score obtained for this RAE architecture. The IF score is more negative (i.e., the sources are more anomalous) for light curves with more epochs. This happens because, as shown in Figure 3, most of the light curves in the full sample have less than 200 epochs, thus light curves with larger numbers of epochs are considered anomalous by the IF model.

VAEs have shown to be able to deal with AD problems properly (e.g., Portillo et al. 2020; Villar et al. 2021), hence we repeated the experiment using the VRAE architecture described in Section 3.1 to test whether it could better deal with the high imbalance in n_{det} of the full sample, using the balanced phys set to train the model. The results of this test are presented in panel (b) of Figure 6 and in panel (c) of Figure 7 (orange plots in both figures). In this case, we used the mean latent space as a vector of attributes (VRAE e_i , with i the number of the attribute). The values of the vector of attributes show much fewer correlations in general across all the layers with n_{det} , with the exception of the attributes VRAE e_{11} and VRAE e_{12} , where a weak correlation can be observed. We also used these attributes as features for our IF model (see panel c of Figure 7), and noticed that the correlation of the IF score with n_{det} is highly reduced. These results tell us that the proposed VRAE architecture is better than the original RAE architecture not only because of its smooth latent space representation but also because it offers a latent space robust against diversity in the length of the time series; basically more motivated by physical differences than by those in the observing strategy.


Finally, we tested whether the correlation with the n_{det} is eliminated when we train both RAE and VRAE architectures with training and validation sets that are balanced by means of their physical properties and n_{det} . We used the balanced phys-epochs set described in Section 2.3. The results for the full sample when using the RAE architecture are shown in panel (c) of Figure 6 and in panel (b) of Figure 7 (black plots in both figures). We can see in both figures that the correlation with n_{det} is highly reduced. A similar result is observed for the VRAE architecture, whose results are shown in panel (d) of Figure 6 and in panel (d) of Figure 7 (blue plots in both figures). In general, the correlation with n_{det} is nonexistent or very weak (e.g., VRAE e_2 , and VRAE e_{10}), although the improvement with respect to the previous VRAE model is small, and the main differences are observed in the IF_{score} obtained for light curves with $n_{\text{det}} > 600$ (see panels c and d of Figure 7).

From these results, we can conclude that a VRAE architecture provides better performance when dealing with data sets that are highly imbalanced in their observational properties, like n_{det} , and that for both RAE and VRAE architectures better results can be obtained when balancing the




(a) RAE balanced phys

(b) VRAE balanced phys



(c) RAE balanced phys-epochs



(d) VRAE balanced phys-epochs

Figure 6. Latent space attributes as a function of the number of epochs for an RAE (panel a; in purple) and a VRAE (panel b; in orange) architecture trained with the balanced phys set and for an RAE (panel c; in black) and a VRAE (panel d; in blue) architecture trained with the balanced phys-epochs set. In the four panels, brighter areas correspond to overdensities (as in Figure 2, but we omitted the color bars). The Spearman rank correlation coefficient is shown as reference on top of each subplot. For all the subplots, the p -value of the coefficient is $p_{\text{val}} \leq 10^{-8}$.




Figure 7. Number of epochs per light curve as a function of the IF score for the full sample, obtained using features extracted by an RAE architecture trained with a data set balanced by means of their physical properties (panel a; purple; RAE balanced phys), by an RAE architecture trained with a data set balanced by means of their physical properties and number of epochs per light curve (panel b; black; RAE balanced phys-epochs), by a VRAE architecture trained with a data set balanced by means of their physical properties (panel c; orange; VRAE balanced phys), and by a VRAE architecture trained with a data set balanced by means of their physical properties and number of epochs per light curve (panel d; blue; VRAE balanced phys-epochs). In the four panels, brighter areas correspond to overdensities (as in Figure 2, but we omitted the color bars).

training and validation sets by means of n_{det} . We decided to use the VRAE architecture trained with balanced phys-epochs set as our final model since it provides results that are not biased by the number of epochs and because the smooth latent space facilitates the detection of anomalies.

4.2. Quality of the Reconstructed Light Curves

We used the VRAE architecture trained with the balanced phys-epochs set to reconstruct all the light curves in the full sample. Figure 8 shows a random selection of original light curves and reconstructions from the full sample. Figure 9 shows the normalized histogram of the reconstruction errors obtained for the light curves in the full sample and light curves with classes Q, A, and BLLac. Most of the sources in the full sample have class Q (see Figure 1), thus it is not surprising that the distribution of R for the full sample is quite similar to the distribution of class Q.

Thirteen percent of the total sample has $R < 0.1$. Similar to the balanced sample, a high fraction of these sources have ZTF light curves with large photometric errors (82% have average photometric errors larger than 0.1). After a visual inspection, we noticed that several of them seem to have overestimated photometric errors, which explains the low R values measured for them.

The high R values obtained for BLLacs are not surprising, as these sources are expected to be intrinsically different from regular quasars. The high R values obtained for some objects with class A could be produced by three factors: (1) contamination from the host is damping the light curves, (2) there are still some issues with the photometry, even though we removed the more extended sources (e.g., problems in the photometry not identified by the ZTF pipeline) from the sample, and (3) the light curves of class A and Q sources are intrinsically different due to, e.g., different physical properties (BH_{mass} , accretion rate, redshift). The observed results are probably due to a combination of these three factors since the lower luminosity leads to more host-galaxy contribution and also host-dominated sources normally appear as extended sources in the ZTF images.

By visually inspecting the reconstructed light curves, we noticed that long-term variations are properly recovered by the reconstructions; however, in some cases, more rapid variations are missed by the reconstruction. An example of this can be




Figure 8. Original light curves (black) and reconstructed light curves (red) for a random selection of light curves from the full sample.




Figure 9. Normalized histogram of the reconstruction error (R) per light curve for the full sample (blue solid), the Q class (orange dashed), the A class (green dotted), and the BLLac class (pink thick dashed). The threshold of $R = 3$ is shown with a black dashed line.

observed in the top middle panel of Figure 8. Tachibana et al. (2020) observe a similar result for the RAE architecture. This may be related to the cadence and photometric noise of the




Figure 10. Fraction of light curves per class classified as anomalous. The black dashed line highlights the position where the fraction is equal to 2%.

light curves, but it could also reflect the fact that in general, sources belonging to class Q do not show very rapid variations (the timescale of the variations is of the order of months or years since they are more luminous and more massive). The results obtained for class A sources are probably explained by this since they are expected to show more rapid variations as they are less luminous and less massive (MacLeod et al. 2010; Simm et al. 2016; Sánchez-Sáez et al. 2018).

In any case, the results obtained by the VRAE architecture are quite robust, regardless of the diversity of cadences, photometric errors, and variability behaviors and the fact that we are neither interpolating the data, nor correcting by redshift, nor doing any light-curve manipulation, as other works that use a similar architecture do (e.g., Naul et al. 2018; Villar et al. 2020, 2021).

4.3. Selection of Anomalous Light Curves

As described in Section 3.2, we applied our trained VRAE model to the full sample and used the mean of the latent space as attributes for our IF algorithm. Then we classified a given light curve as anomalous if its VRAE reconstruction error was $R \geq 3$ or if its $IF_{\text{score}} \leq IF_{\text{th}}$.

Figure 10 shows the fraction of light curves labeled as anomalous for each of the classes described in Section 2.1. The low fraction of anomalous objects belonging to class Q is not surprising since we trained our model with a data set that is dominated by this class. The fraction of anomalous Q sources is close to the 2% that we defined as threshold in our AD strategy. However, the classes with the largest fraction of anomalies are the three blazar classes (BLLac, BZU, and FSRQ), the host-dominated classes (A and AR), and the type-2 AGNs with radio emission (type2 R).

As mentioned in Section 2.1, blazars were included in the analysis to test the efficiency of our AD methodology and we expected these classes to be classified as anomalous. The results obtained for blazars therefore confirm the capacity of our model to detect anomalous AGN light curves. Interestingly, the fraction of anomalous FSRQs is $\sim 35\%$, and thus much smaller than the fraction observed for BLLacs ($\sim 65\%$). From this we might conclude that FSRQs have optical emission that is intrinsically more similar to the optical emission of regular quasars, compared to BLLacs, but further analysis is required to confirm this result.

Something interesting that can be noticed from Figure 10 is that in general the classes with detected radio emission (AR,

type2 R, and QR) tend to have a larger fraction of light curves classified as anomalous compared to their counterpart classes without detected radio emission (A, type 2, and Q). From this we might conclude that AGNs with radio emission have intrinsically different optical variability behaviors compared to AGNs without radio emission. Further analysis is required to confirm this result as we are using the classifications provided by MILLIQUAS and thus we cannot ensure that all the sources without radio emission do not emit in the radio wave bands.

4.4. Types of Anomalies Found

In this work we used the AD methodology as a recommendation system to search more easily for CSAGN candidates. Thus, in order to select CSAGN candidates, we visually inspected the sources classified as anomalies by our method, excluding the three blazar classes and class R since these classes were included only for testing purposes. We define a source as a promising CSAGN candidate if its light curve shows during visual inspection of flaring activity or abrupt increase/decrease of the optical flux. By making this selection, we cannot ensure that the candidates are in fact CSAGNs as they can also correspond to examples of atypical flaring activity (e.g., Graham et al. 2017; Trakhtenbrot et al. 2019b; Frederick et al. 2020) or extremely variable AGNs (e.g., Rumbaugh et al. 2018; Guo et al. 2020; Luo et al. 2020) and that is why we called them promising CSAGN candidates.

Here we provide a general description of the types of anomalies found for classes Q, QR, A, AR, type2, type2 R, and X. A sample of promising CSAGN candidates is presented in Section 4.5.

1. *Class Q (5774 anomalies)*: around 40% of the anomalies have very noisy light curves where the reported photometric errors seem to be underestimated (since they show outlier observations with small photometric errors), or show evidence of problems in the photometry. We inspected the ZTF stamps (using the ZTF Time Series Tool) of some of these sources and found that in some cases there are artifacts in the images that are not removed when using `catflags=0`. Figure 11 shows a set of stamps for two of these candidates. All these stamps have `catflags=0`, however, the artifacts in the images are easily identified. In addition, the anomalies in this class have a peak in redshift at $z \sim 0.5$, while the Q sources labeled as normal have in general $z > 1$, and thus, it is probable that a high fraction of the anomalies have light curves distorted by the emission of their host galaxies or are affected by time dilation and rest-frame differences. Moreover, we found two sources that are in fact variable stars but were classified as quasars by the SDSS pipeline (ZTF DR5 ID 794110400002914 and 801111300003144). We identified 49 promising CSAGN candidates.
2. *Class QR (747 anomalies)*: a high fraction of the anomalies have variability behaviors similar to those observed for BLLacs. In particular, 63% of the QR anomalies have classifications in the ALERCE light-curve classifier (Sánchez-Sáez et al. 2021) and 26% of these are classified as blazars. In addition, 36 of these anomalies are classified as BLLac in the SIMBAD database (Wenger et al. 2000). After the visual inspection we identified 10 promising CSAGN candidates.




Figure 11. Set of ZTF stamps (with a cutout size of 3 arcminutes) in the g band for two anomalies with class Q, with ZTF DR5 ID 498110300023393 (top) and 646103300003619 (bottom). The orange circles show the location of the target in the stamps. The stamps were obtained using the ZTF Time Series Tool.

3. *Class A (873 anomalies)*: the anomalies in this class have lower redshifts than those labeled as normal and therefore are probably highly dominated by the emission from the host galaxies. Their light curves show in general low-amplitude rapid variations, which are not properly recovered by the VRAE reconstructions. We found one source classified as an AGN by the SDSS pipeline (ZTF DR5 ID 408114200003141) but classified as a YSO in SIMBAD. We identified 10 promising CSAGN candidates.
4. *Class AR (190 anomalies)*: as with class A, the anomalies in this class have lower redshifts than those labeled as normal, thus their light curves could be more contaminated by emission from the hosts. Similar to the QR class, some sources show variability behavior close to those observed for BLLacs. Eighty-eight percent of these anomalies have classification in the ALerCE light-curve classifier and 18 are classified as blazars. In addition, 10 are classified as blazars in the SIMBAD database. We identified three promising CSAGN candidates.
5. *Classes type2 and type2 R (20 anomalies)*: several of the type2 anomalies are probably type 1.8 or 1.9 sources since their SDSS spectra show evidence of a broad emission line component, which may help to explain their optical variations. We identify two promising CSAGN candidates.
6. *Class X (246 anomalies)*: when we inspected the light curves of these anomalies, we noticed that several of them seem to be variable stars. Forty-six percent of the anomalies have classification in the ALerCE light-curve classifier and 87 are classified as variable stars (including several cataclysmic variables and young stellar objects). In addition, 65 have classification in the Simbad database and only 4 of them are classified as AGNs or AGN candidates; the rest are labeled as variable stars. We identify one promising CSAGN candidate.

4.5. CSAGN Candidates

To select the sample of CSAGN candidates we rejected very noisy light curves, the sources identified as variable stars in SIMBAD, and the sources that can be classified as variable stars according to their light curves. We selected as promising

CSAGN candidates those anomalies that present evidence of flares, and/or abrupt increment or decrement in the luminosity.

We identified 75 promising CSAGN candidates (65% of them belong to class Q); 73 of these have alerts in the ZTF alert stream. The full list of candidates is presented in the [Appendix](#). Four of these candidates appear in the list of CSAGNs presented in Graham et al. (2020) and another two have been confirmed spectroscopically (M. Graham 2021, private communication). Moreover, 28 of the candidates have been selected as CSAGN candidates using CRTS data and/or ZTF DR3 data (M. Graham et al. 2021, in preparation). All these candidates are highlighted in [Appendix](#).

Sixty candidates have at least one SDSS spectrum available, with 12 observed relatively recently (MJD > 58,000) such that it overlaps with the overall baseline of the ZTF light curves. We leave for future work any detailed spectroscopic analysis of the sample as this type of study is beyond the scope of this work and we provide here a brief description of the properties of the most interesting sources.

The light curves and reconstructions of the 12 candidates with recent SDSS spectra are shown in Figure 12. We can see that the light curves show a variety of shapes. We noticed that 72% of the total sample of candidates present periods of low-amplitude variations (sometimes even constant emission), preceded or followed by abrupt variations and we called this the “plateau” feature (e.g., ZTF DR5 ID 566112100004425, 646113100002570, and 649113400004567). These types of light-curve shapes have been previously detected in X-ray light curves of X-ray binaries and are associated with changes in the accretion state (e.g., Remillard & McClintock 2006; Done et al. 2007; Sobolewska et al. 2011) and also recently in X-ray light curves of CSAGNs during outburst activity (Ruan et al. 2019a).

Five of the sources with recent SDSS spectra show very narrow $H\alpha$ profiles (ZTF DR5 IDs: 674113200005713, 678115400001961, 693106300007466, 712112100003746, and 719114100002677). One of these (ZTF DR5 ID 674113200005713) was previously classified as a narrow-line Seyfert 1 (NLS1) by Rakshit et al. (2017). Additionally, three of the candidates with old SDSS spectra (MJD < 58000) are classified as NLS1 in Rakshit et al. (2017; ZTF DR5 IDs: 479108100002617, 582108200000456, and 676103300001128). Previous works have shown that NLSy1 can present rapid enhanced flaring activity (e.g., Miller et al. 2000; Frederick et al. 2020), thus we need to follow up the candidates to




Figure 12. As in Figure 8 but for CSAGN candidates with recent SDSS spectra. For the reconstruction we show the average and standard deviation (error bars) of the 10 reconstructions. The shaded regions show the minimum and maximum values obtained for the reconstruction at each epoch.

confirm whether they correspond to CSAGN events or are examples of these new optical transients detected in NLSy1.

Finally, there are nine candidates with atypical spectra (ZTF DR5 IDs: 427114100001870, 474105300001176, 526115100009333, 578114100000128, 616105200002737, 626104200000260, 66411300002498, 789103200005093, and 818114300000932), with two of them (ZTF DR5 IDs: 626104200000260, and 66411300002498) showing very asymmetric broad line profiles, similar to those observed in NGC 3516, a known CSAGN (Oknyansky et al. 2021).

Due to the current COVID-19 pandemic and subsequent lengthy closure of the observatories, we have not been able to follow up our candidates, but we provide to the astronomical community the candidate list so that future follow-up campaigns can target them.

5. Discussion

5.1. About the Use of AD Techniques to Search for CSAGNs

In this work we used an AD approach to search for CSAGNs. A high fraction of the anomalies found by our method are bogus candidates (e.g., noisy light curves, or light curves with problems in the photometry), or sources with incorrect labels in the original catalogs (e.g., variable stars classified as quasars by the SDSS pipeline, or X-ray-detected variable stars). This type of issue has been found in previous works that used AD algorithms to search for anomalous light curves (e.g., Malanchev et al. 2021, M. Pérez-Carrasco et al. 2021, in preparation). Thus, our algorithm can be a powerful tool to identify light curves with photometric issues in massive data sets and to identify mislabeled sources.

Moreover, as mentioned in Section 4.2, the light curves of host-dominated AGNs are not properly reconstructed by our VRAE architecture, producing a large fraction of anomalous candidates, which hinders the discovery of true outliers such as CSAGNs from this population. Tachibana et al. (2020) found similar results and decided to remove host-dominated AGNs from their analysis. In our case, we were able to detect promising CSAGN candidates in the A and AR classes, thus, for the sake of inclusivity, we do not rule out host-dominated AGNs from this work. A possible solution to this problem would be to remove the host-galaxy emission from the light curves or to use light curves obtained from difference images. We leave this type of analysis for future work.

These results show us that we cannot blindly use the output of our AD algorithm to define a sample of CSAGN candidates and demonstrate that interaction between AD algorithms with domain experts is needed in order to find reliable samples of candidates (Ishida et al. 2021; Malanchev et al. 2021; M. Pérez-Carrasco et al. 2021, in preparation).

In addition, the candidates presented in this work require further confirmation, as atypical variations have also been observed in other AGN processes, like flaring activity in NLS1s (Trakhtenbrot et al. 2019b; Frederick et al. 2020), or extremely variable AGNs (Rumbaugh et al. 2018; Guo et al. 2020; Luo et al. 2020).

Finally, it is encouraging that we are finding promising CSAGN candidates in light curves with timespans that range between 2 and 2.8 years. We expect that much better results will be obtained by our AD algorithm in the future with the upcoming ZTF data releases and also with data obtained by the Vera Rubin Observatory Legacy Survey of Space and Time (LSST; Ivezić et al. 2019) or by combining different data sets to extend the baselines by ~ 1 order of magnitude and in different filters.

5.2. The Plateau Feature

A high fraction (72%) of our CSAGN candidates show light-curve shapes with a fast rise or decay and a “plateau” that spans hundreds of days; we called this shape the “plateau feature” (e.g., sources with ZTF DR5 ID 566112100004425, 646113100002570, and 64911340000456 in Figure 12). These light-curve shapes are reminiscent of those previously observed in X-ray binaries at X-ray energies (Remillard & McClintock 2006; Done et al. 2007; Sobolewska et al. 2011). In these types of sources, it is well established that a coupling exists between the accretion disk and the jet, giving rise to different states that can be followed up on human timescales thanks to the small size of the black holes (Fender et al. 2004). It has been suggested that AGNs might be scaled-up versions of X-ray binaries, thus the different states could also be present (McHardy et al. 2006), but with subsequently longer timescales between them.

If we extrapolate the accretion behaviors observed in X-ray binaries to regular AGNs (Fender et al. 2004), the predicted timescales of these accretion state transitions are $\sim 10^4$ – 10^5 years (if the changes are related to viscous timescales), thus




Figure 13. As in Figure 8, but for two flaring events presented in Frederick et al. (2020; ZTF DR5 IDs 629115200005576 and 634110400007229), and for the CSAGN event presented in Ricci et al. (2020; ZTF DR5 ID 827113200013781).

much larger than the ones observed in our CSAGN candidates. However, these state transitions have been observed in CSAGN events in the X-ray domain (Ruan et al. 2019a; Ricci et al. 2020), thus it is possible to use CSAGNs to obtain insights about the similarities between the accretion flows of X-ray binaries and AGNs (e.g., Ruan et al. 2019b). The physical mechanism behind this plateau feature remains an active investigation, and we leave that for future analysis.

5.3. Comparison with Previous Works

Frederick et al. (2019) and Ricci et al. (2020) detected CSAGN events (six events and one event, respectively) that are visible in the ZTF DR5 data, and Frederick et al. (2020) presented a sample of five NLSy1 detected in the ZTF alert stream with atypical flaring activity. None of these sources were included in our analysis since all of them are classified as extended sources by Tachibana & Miller (2018; $ps_score < 0.5$). Despite this, we applied our AD methodology to these sources and all of them were classified as anomalies. The ZTF PSF-fit-based light curves of these targets are quite noisy, thus we consider their detection with our methodology with caution, except for three sources: the CSAGN event presented in Ricci et al. (2020) and two flaring events from Frederick et al. (2020) whose light curves are shown in Figure 13. These sources have very well sampled light curves in the ZTF alert stream and variability amplitudes larger than 1 magnitude and thus were easy to identify using only the alert stream. Our method, by contrast, allows the detection of anomalous variations with any level of variability amplitudes, which favors the early detection of CSAGN events and the detection of CSAGN events with smaller amplitudes that could be missed by the ZTF alert stream.

Graham et al. (2017, 2020) used more than eight years of CRTS data to search for AGNs with extreme optical variability, finding a sample of 51 sources with major flares and 111 CSAGN candidates, respectively. Graham et al. (2017) used a Weibull distribution to characterize the flares and Graham et al. (2020) used BB representation (Scargle et al. 2013) to identify significant changes in variability behaviors. The large baseline of the CRTS data allowed them to use these techniques to find extreme variability behaviors with timescales of several years, which are still impossible to detect with the existent baseline of the ZTF data, thus we do not consider applying such techniques to the current ZTF light curves. With the upcoming ZTF data releases, we might be able to complement our AD methodology with these other techniques.

MacLeod et al. (2016) used SDSS and Pan-STARRS1 (PS1; Chambers et al. 2016) data to select CSAGN candidates. They used a sample of quasars with repeated SDSS spectra

and defined those sources with large photometric variations ($|\Delta g| > 1$) as CSAGN candidates. The use of a sample of well characterized quasars avoided the introduction of contamination from incorrect classifications. However, this way to select CSAGN candidates based solely on their variability amplitudes cannot be extrapolated to larger data sets, since the contamination from noisy light curves, or from sources with incorrect labels (e.g., variable stars, or even blazars) would be considerably larger.

6. Conclusions

In this work we presented an AD algorithm to search for AGNs with anomalous light curves in massive data sets (like ZTF DR5), whose main aim is to find CSAGN candidates. We used an architecture inspired by the work with RAE architecture presented by Tachibana et al. (2020) but with two modifications: we used GRUs instead of LSTMs and we used a VRAE architecture instead of RAE.

Our data set is composed of 230,451 g-band light curves from ZTF DR5. The light curves are associated with different classes of AGNs, although 90.7% are regular quasars (class Q). The rest of the sources are type-2 AGNs, Seyfert galaxies, blazars, and sources selected by their X-ray or radio emission. We used the ZTF PSF-fit-based light curves from DR5 with more than 2 years of data and more than 50 epochs.

We first tested the original RAE architecture and noticed that the attributes of the latent space were highly correlated with the number of epochs per light curve, particularly when the training and validation sets were not balanced by means of the number of epochs (using the balanced phys set). We therefore tested whether a VRAE architecture is able to solve this issue, finding a considerable reduction in the level of correlation with the number of epochs when using the latent mean of the VRAE as a vector of attributes. We decided to train the RAE and VRAE architectures with training and validation sets that were balanced by means of their spectroscopic properties (obtained from SDSS data) and by means of the number of epochs per light curve (using the balanced phys-epochs set) to avoid any correlation with the cadence of the light curves. We found that when using the balanced phys-epochs set, the latent space of the RAE architecture reduced the correlations with the number of epochs considerably, while for the case of the VRAE architecture there is a small improvement. From these results, we decided to use the VRAE architecture trained with the balanced phys-epochs set for our AD algorithm as a VRAE architecture to have the potential of solving any other balancing issues present in the data (e.g., redshift, host-galaxy contamination, presence of a jet, among others). This allowed us to work with the varied cadences present in the ZTF light curves, which have large differences in the number of epochs per light curve and also allowed us to use the latent space as a vector of features for an IF algorithm.

We used a two-stage approach to classify anomalies. First we reconstructed all the light curves in our sample with our VRAE architecture and measured the reconstruction error R for each light curve. We classified as anomalies all the light curves with $R \geq 3$. Then, we used the mean latent space as features for an IF algorithm and classified all the sources with an IF score smaller than the threshold defined for a contamination of 2% for the training set ($IF_{score} \leq -0.57633$) as anomalies. With this selection we ended up with a sample of 8,809 anomalies, with 65.5% them belonging to class Q.

The set of anomalies is dominated by sources with noisy light curves (due in some cases to problems in the images that are not identified by the ZTF pipeline) and by sources with incorrect labels in the original catalogs (e.g., we found several variable stars classified as AGNs in the literature). Therefore, we highlight that our AD algorithm is useful not only to find intrinsic anomalous light curves but also to detect problems in the data sets used for AGN variability analyses.

We visually inspected the anomalies found by our AD algorithm and selected a sample of 75 promising CSAGN candidates. The full list of candidates is presented in the [Appendix](#). Visual inspection is required, as a large fraction of our anomalies correspond to bogus candidates (like light curves with photometric issues or incorrect labels reported in the literature). Further spectroscopic confirmation would be required to confirm the real nature of our sample.

A high fraction of our candidates present light-curve shapes similar to those observed in X-ray light curves of X-ray binaries, which are generally associated with changes in the accretion state. The timescales of these variations in our set of candidates are of the order of hundreds of days, much smaller than the typical viscous timescales for accretion disks around SMBHs (order of thousands of years), thus we need further analyses to discover the origin of these atypical variations.

This work corresponds to our first attempt to search for CSAGN candidates in massive data sets. In the future, we plan to use this architecture, together with other techniques, to forecast the light curves of known type-1 and type-2 AGNs, and from this, detect variability behaviors that deviate from the normal variations expected for each source. For this aim, we will use data from the ZTF data releases, but also data from the ZTF alert stream, which should allow us to detect CSAGN candidates in real time and trigger follow-up to better pin down the nature of this enigmatic phenomenon.

This work was funded by: project CORFO/ANID International Centers of Excellence Program 10CEII-9157 Inria Chile (PSS, HL, LM, NSP); ANID—Millennium Science Initiative Program—ICN12_009 (PSS, JA, GCV, PE, LHG, AMA, FEB, FF); ANID—Millennium Science Initiative Program—NCN19_171 (AB); CATA-Basal—AFB-170002 (FEB); the Competition for Research Regular Projects, year 2019, code LPR19-22, Universidad Tecnológica Metropolitana and the high-performance computing system of PIDi-UTEM (SCC-PIDi-UTEM CONICYT-FONDEQUIP-EQM180180; JV); ANID FONDECYT Postdoctorado No 3200250 (PSS); ANID

FONDECYT Iniciación No 11191130 (GCV); ANID FONDECYT Regular No 1171678 (PE), 1190818 (FEB), 1200495 (FEB), and 1200710 (FF).

Based on observations obtained with the Samuel Oschin 48 inch Telescope at the Palomar Observatory as part of the ZTF project. ZTF is supported by the National Science Foundation under grant No. AST-1440341 and a collaboration including Caltech, IPAC, the Weizmann Institute for Science, the Oskar Klein Center at Stockholm University, the University of Maryland, the University of Washington, Deutsches Elektronen-Synchrotron and Humboldt University, Los Alamos National Laboratories, the TANGO Consortium of Taiwan, the University of Wisconsin at Milwaukee, and Lawrence Berkeley National Laboratories. Operations are conducted by COO, IPAC, and UW.

Software: Google Colaboratory (Bisong 2019), Keras (Chollet et al. 2015), Scikit-learn (Pedregosa et al. 2011), and Tensorflow (Abadi et al. 2016). For graphical representations and data manipulation: Jupyter (Kluyver et al. 2016), Matplotlib (Hunter 2007), Numpy and Scipy (van der Walt et al. 2011), Pandas (McKinney et al. 2010), Python (Van Rossum & Drake 1995), and Seaborn (Waskom et al. 2017).

Appendix List of Promising CSAGN Candidates

Table 1 shows the coordinates, redshifts (reported by MILLQUAS), previous classifications, and ZTF DR5 and alerts object IDs of the 75 CSAGN candidates described in Section 4.5. The column “plateau” indicates whether the light curve of the source shows the plateau feature in its light curve. The column “SDSS spectrum” indicates whether the source has SDSS spectra available. Sources with recent spectra are marked with \blacklozenge in this column. There are four candidates in this sample that were reported in the list of CSAGNs of Graham et al. (2020; objects with ZTF alert ID ZTF19aalmwdr, ZTF18aaxmbqw, ZTF18aahhdas, and ZTF18aaqmll). In addition, two candidates have already been confirmed spectroscopically as CSAGNs (objects with ZTF alert ID ZTF18acbzncb and ZTF18aavsdjf; M. Graham 2021, private communication). These six confirmed candidates are marked with \dagger in column “ID ZTF alerts.” Finally, 28 candidates are identified as candidates from their CRTS light curves and/or their ZTF DR3 light curves, using the techniques presented in Graham et al. (2020; M. Graham et al. 2021, in preparation); these sources are marked with \star in column “ID ZTF alerts.”

Table 1
Selection of Promising CSAGN Candidates

R.A. (deg)	Decl. (deg)	Redshift	Class	ID ZTF DR5	ID ZTF alerts	Plateau	SDSS spectrum
34.47318	−19.972092	0.46	Q	299104400000216	ZTF19abkfpqk	Yes	No
66.427416	−10.535333		X	354108100003201	ZTF19acnkupu	No	No
152.664006	−12.737412	0.42	type2	366103400000077	ZTF20aadswdl	No	No
35.812817	−5.317624	0.96	type2	401103200007922	ZTF19aadbbtz	No	No
151.046448	−3.613936	0.45	Q	417106400000034	ZTF19aagoauz \star	No	No
221.653773	0.782385	0.72	QR	427114100001870	ZTF18acuszrd \star	No	Yes
229.643977	−1.090381	0.65	Q	428109200001687	ZTF19aascyqe \star	No	Yes
310.764753	−1.023893	1.19	Q	440111200002868	ZTF18abvvtxk	Yes	Yes
28.853979	2.471266	0.08	AR	450101200003599	ZTF18abqvdu \star	No	No
34.740162	2.461193	1.73	Q	451102100000333	ZTF19abfqigm	Yes	Yes
31.842494	2.526327	0.66	Q	451103200005537	...	Yes	Yes \blacklozenge
38.714227	2.774584	1.17	Q	452107300008623	ZTF19abjgcof	Yes	Yes

Table 1
(Continued)

R.A. (deg)	Decl. (deg)	Redshift	Class	ID ZTF DR5	ID ZTF alerts	Plateau	SDSS spectrum
129.477711	4.002007	0.56	Q	465107200003387	ZTF19aufsep	No	Yes
143.825809	2.07099	0.64	QR	467103100002301	ZTF18acezjdj*	Yes	Yes
174.278534	1.663312	0.19	Q	471101300000081	ZTF19aalmwdr†	Yes	Yes
183.858915	3.842626	0.93	Q	473107200004410	ZTF18acvgvou*	Yes	Yes
194.984234	3.409168	0.74	Q	474105300001176	ZTF20aavavif*	Yes	Yes
194.992985	7.223641	1.5	Q	474113300001013	ZTF18acvwlvv*	Yes	Yes
228.771345	4.170062	1.26	Q	479105200007871	ZTF19aailvld*	Yes	Yes
224.519576	4.054482	0.69	Q	479108100002617	...	No	Yes
236.681566	2.138313	0.6	QR	481104200003568	ZTF19abbtqpb	Yes	Yes
336.760588	6.224071	1.23	Q	495111100005833	ZTF20abcxfmu*	No	Yes◆
333.860942	6.253339	0.22	Q	495112200000743	ZTF18aceiofb*	Yes	No
193.365443	14.915562	0.25	Q	526115100009333	ZTF18aahhdas†	Yes	Yes
206.858147	11.600358	0.84	Q	528107200000356	ZTF18ablqds*	Yes	Yes
310.077349	10.995846	0.08	Q	543108100012056	ZTF18aazmwqp	No	No
113.92998	20.215779	2.1	Q	566112100004425	ZTF18abwclfc	Yes	Yes◆
144.242347	19.499183	0.38	Q	570111400005490	ZTF18accnmoe*	No	Yes
202.796024	22.656575	0.96	Q	578114100000128	ZTF18aavsdj†	Yes	Yes
227.421616	18.766963	0.32	A	582108200000456	ZTF20aaznovd	Yes	Yes
228.182002	19.979198	0.8	Q	582112100014255	ZTF20aauwvfb	Yes	Yes
244.439728	16.548838	2.11	Q	584103100003077	ZTF19aailumj	Yes	Yes
333.436661	21.787462	0.65	Q	596114400000305	ZTF19acgfpvb	Yes	Yes
26.371489	28.705946	0.64	QR	603113300005168	ZTF19abdkgly	Yes	No
123.570806	25.493102	0.6	Q	616105200002737	ZTF18aagevkc*	Yes	Yes
192.926178	24.076456	0.18	AR	626104200000260	ZTF18acahbqb*	No	Yes
204.008817	23.63965	0.49	Q	627102200007292	ZTF19aailxfk	Yes	Yes
349.263518	29.416571	1.26	Q	646113100002570	ZTF20abwddhk	Yes	Yes◆
4.257901	30.070426	1.3	Q	648101400002016	ZTF18abvztgk	Yes	Yes◆
358.545135	35.605134	0.22	Q	648115300004761	ZTF18abmawjk	Yes	No
12.46894	33.53162	0.37	A	649109400004074	ZTF19abujzxs	Yes	No
12.624976	35.436805	1.56	Q	649113400004567	ZTF20abmdzgu	Yes	Yes◆
47.684888	32.658078	0.12	A	654106200004523	ZTF19aacqjiv*	Yes	No
120.834624	29.763564	0.52	Q	663101300003892	ZTF18aagccsf	Yes	Yes
124.148074	33.796856	0.51	AR	664111300002498	ZTF20acfxnec*	Yes	Yes
184.860032	32.023781	0.63	Q	672108300002204	ZTF19actnycp	No	No
185.403099	34.59464	0.29	A	672112100002964	ZTF18aaxmbqw†	Yes	Yes
206.738679	36.547599	0.24	A	674113200005713	ZTF18aacdexp	Yes	Yes◆
218.187173	30.243154	0.35	QR	676103300001128	ZTF18aautsvy*	No	Yes
215.7201	30.274046	2.2	Q	676104300000686	ZTF20aawwxkg	Yes	Yes
235.017706	35.847258	0.16	A	678115400001961	ZTF18aakhtca*	Yes	Yes◆
260.736568	32.240587	1.17	Q	681106400012690	ZTF18aayaxrf	No	Yes
348.418519	30.217598	0.86	Q	692101400001447	ZTF18abtclue*	Yes	Yes
353.578057	32.08145	0.32	Q	693106300007466	ZTF18acbxoul	Yes	Yes◆
134.863998	43.929443	0.35	A	709113200000994	ZTF20aayvktl	Yes	Yes
155.107655	41.561772	0.36	Q	712112100003746	ZTF18acbzncb†	Yes	Yes◆
180.15749	44.177165	1.37	QR	715116100000207	ZTF18aaqleut	Yes	Yes
192.080372	37.107331	1.52	Q	716103400002868	ZTF19adcfwam	Yes	Yes
188.569167	42.559662	1.54	Q	716116400003725	ZTF20aagxkl	Yes	Yes
219.800624	37.085547	0.79	Q	719102400002793	ZTF19adcfvjv	Yes	Yes
220.51178	43.619104	0.23	Q	719114100002677	ZTF18aaqmll†	No	Yes◆
237.316978	41.452011	0.93	Q	721110400013285	ZTF20aagiimu	Yes	Yes
243.249306	42.32787	0.23	QR	722111100000279	ZTF18aanlzzf*	No	Yes
5.028384	45.133848	0.81	QR	736101200006573	ZTF20abcxodv*	Yes	No
15.913292	47.5503	0.19	A	737105200000494	ZTF18abosuim*	No	Yes
230.847407	49.377265	1.57	Q	760112100000758	ZTF18aczerbt*	No	Yes
233.355747	50.192888	1.1	Q	760115400002305	ZTF20aalrlww	Yes	Yes
298.14922	49.97055	0.46	A	766113400010935	ZTF18aawwilc	Yes	No
177.603194	52.873391	0.7	QR	789103200005093	ZTF18aceyygm*	No	Yes
220.21153	52.079437	0.31	A	793104100003683	ZTF19aaxezow*	Yes	Yes
222.201328	53.340535	1.55	Q	793107300003189	ZTF20aapbkig	Yes	Yes
225.030254	56.60022	0.88	QR	793110200000811	ZTF19abzmptx	Yes	Yes
248.321942	54.399034	1.37	Q	795106100002111	ZTF19aawmmfv*	Yes	Yes◆
167.121585	64.858964	0.71	Q	818114300000932	ZTF19aaeduzf*	Yes	Yes
257.828352	65.503293	1.09	Q	825115100005176	ZTF19aaqagri	Yes	Yes

ORCID iDs

P. Sánchez-Sáez  <https://orcid.org/0000-0003-0820-4692>
 H. Lira  <https://orcid.org/0000-0002-5155-587X>
 L. Martí  <https://orcid.org/0000-0003-2789-5062>
 N. Sánchez-Pi  <https://orcid.org/0000-0002-5037-9974>
 J. Arredondo  <https://orcid.org/0000-0002-2045-7134>
 F. E. Bauer  <https://orcid.org/0000-0002-8686-8737>
 A. Bayo  <https://orcid.org/0000-0001-7868-7031>
 G. Cabrera-Vives  <https://orcid.org/0000-0002-2720-7218>
 P. A. Estévez  <https://orcid.org/0000-0001-9164-4722>
 S. Eyheramendy  <https://orcid.org/0000-0003-4723-9660>
 F. Förster  <https://orcid.org/0000-0003-3459-2270>
 L. Hernández-García  <https://orcid.org/0000-0002-8606-6961>
 A. M. Muñoz Arancibia  <https://orcid.org/0000-0002-8722-516X>
 M. Pérez-Carrasco  <https://orcid.org/0000-0003-4644-8698>
 J. R. Vergara  <https://orcid.org/0000-0001-6699-4181>

References

- Abadi, M., Barham, P., Chen, J., et al. 2016, in OSDI'16: Proc. 12th USENIX Conf. Operating Systems Design and Implementation, 265, <https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf>
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, *PASP*, **131**, 018002
- Bisong, E. 2019, Google Colaboratory (Berlin: Springer), 59
- Caplari, N., Lilly, S. J., & Trakhtenbrot, B. 2017, *ApJ*, **834**, 111
- Carrasco-Davis, R., Reyes, S. G., Valenzuela, C., et al. 2020, arXiv:2008.03309
- Chaini, S., & Kumar, S. S. 2020, arXiv:2006.12333
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv:1612.05560
- Cho, K., van Merriënboer, B., Gulcehre, C., et al. 2014, in Proc. 2014 Conf. Empirical Methods in Natural Language Processing (EMNLP) (Doha: Association for Computational Linguistics), 1724
- Chollet, F. 2015, Keras, <https://keras.io>
- Done, C., Gierliński, M., & Kubota, A. 2007, *A&ARv*, **15**, 1
- Drake, A. J., Djorgovski, S. G., Mahabal, A., et al. 2009, *ApJ*, **696**, 870
- Fabius, O., & van Amersfoort, J. R. 2014, arXiv:1412.6581
- Fender, R. P., Belloni, T. M., & Gallo, E. 2004, *MNRAS*, **355**, 1105
- Flesch, E. W. 2015, *PASA*, **32**, e010
- Flesch, E. W. 2019, arXiv:1912.05614
- Förster, F., Cabrera-Vives, G., Castillo-Navarrete, E., et al. 2021, *AJ*, **161**, 242
- Frederick, S., Gezari, S., Graham, M. J., et al. 2019, *ApJ*, **883**, 31
- Frederick, S., Gezari, S., Graham, M. J., et al. 2020, arXiv:2010.08554
- Graham, M. J., Djorgovski, S. G., Drake, A. J., et al. 2017, *MNRAS*, **470**, 4112
- Graham, M. J., Ross, N. P., Stern, D., et al. 2020, *MNRAS*, **491**, 4925
- Guo, H., Peng, J., Zhang, K., et al. 2020, *ApJ*, **905**, 52
- Hernández-García, L., Masegosa, J., González-Martín, O., Márquez, I., & Perea, J. 2016, *ApJ*, **824**, 7
- Higgins, I., Matthey, L., Pal, A., et al. 2017, in 5th Int. Conf. Learning Representations (Toulon: ICLR)
- Hochreiter, S., & Schmidhuber, J. 1997, *Neural Comput.*, **9**, 1735
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Ishida, E. E. O., Kornilov, M. V., Malanchev, K. L., et al. 2021, *A&A*, **650**, A195
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111
- Kelly, B. C., Bechtold, J., & Siemiginowska, A. 2009, *ApJ*, **698**, 895
- Kingma, D. P., & Welling, M. 2013, arXiv:1312.6114
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in Positioning and Power in Academic Publishing: Players, Agents and Agendas, ed. F. Loizides & B. Schmidt (Amsterdam: IOS Press), 87
- LaMassa, S. M., Cales, S., Moran, E. C., et al. 2015, *ApJ*, **800**, 144
- Liu, F. T., Ting, K. M., & Zhou, Z.-H. 2008, in 2008 Eighth IEEE Int. Conf. on Data Mining (Piscataway, NJ: IEEE), 413
- Lochner, M., & Bassett, B. A. 2021, *A&C*, **36**, 100481
- Luo, Y., Shen, Y., & Yang, Q. 2020, *MNRAS*, **494**, 3686
- Lynden-Bell, D. 1969, *Natur*, **223**, 690
- MacLeod, C. L., Green, P. J., Anderson, S. F., et al. 2019, *ApJ*, **874**, 8
- MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, *ApJ*, **721**, 1014
- MacLeod, C. L., Ross, N. P., Lawrence, A., et al. 2016, *MNRAS*, **457**, 389
- Malanchev, K. L., Pruzhinskaya, M. V., Korolev, V. S., et al. 2021, *MNRAS*, **502**, 5147
- Malhotra, P., Ramakrishnan, A., Anand, G., et al. 2016, arXiv:1607.00148
- Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, *PASP*, **131**, 018003
- Massaro, E., Maselli, A., Leto, C., et al. 2015, *Ap&SS*, **357**, 75
- McHardy, I. M., Koerding, E., Knigge, C., Uttley, P., & Fender, R. P. 2006, *Natur*, **444**, 730
- McKinney, W. 2010, in Proc. of the 9th Python in Science Conf. 445, 51
- Mehrotra, K., Mohan, C., & Huang, H. 2017, Anomaly Detection Principles and Algorithms (Berlin: Springer)
- Miller, H. R., Ferrara, E. C., McFarland, J. P., et al. 2000, *NewAR*, **44**, 539
- Moreno-Barea, F. J., Jerez, J. M., & Franco, L. 2020, *Expert Syst. Appl.*, **161**, 113696
- Naul, B., Bloom, J. S., Pérez, F., & van der Walt, S. 2018, *NatAs*, **2**, 151
- Netzer, H. 2013, The Physics and Evolution of Active Galactic Nuclei (Cambridge: Cambridge Univ. Press)
- Oh, K., Yi, S. K., Schawinski, K., et al. 2015, *ApJS*, **219**, 1
- Oknyansky, V. L., Brotherton, M. S., Tsygankov, S. S., et al. 2021, *MNRAS*, **505**, 1029
- Padovani, P., Alexander, D. M., Assef, R. J., et al. 2017, *A&ARv*, **25**, 2
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *J. Mach. Learn. Res.*, **12**, 2825
- Pereira, J., & Silveira, M. 2018, in 2018 17th IEEE Int. Conf. on Machine Learning and Applications (ICMLA) (Piscataway, NJ: IEEE), 1275
- Portillo, S. K. N., Parejko, J. K., Vergara, J. R., & Connolly, A. J. 2020, *AJ*, **160**, 45
- Rakshit, S., Stalin, C. S., Chand, H., & Zhang, X.-G. 2017, *ApJS*, **229**, 39
- Rakshit, S., Stalin, C. S., & Kotilainen, J. 2020, *ApJS*, **249**, 17
- Remillard, R. A., & McClintock, J. E. 2006, *ARA&A*, **44**, 49
- Ricci, C., Kara, E., Loewenstein, M., et al. 2020, *ApJL*, **898**, L1
- Ross, N. P., Ford, K. E. S., Graham, M., et al. 2018, *MNRAS*, **480**, 4468
- Ruan, J. J., Anderson, S. F., Eracleous, M., et al. 2019a, arXiv:1909.04676
- Ruan, J. J., Anderson, S. F., Eracleous, M., et al. 2019b, *ApJ*, **883**, 76
- Rumbaugh, N., Shen, Y., Morganson, E., et al. 2018, *ApJ*, **854**, 160
- Sánchez, P., Lira, P., Cartier, R., et al. 2017, *ApJ*, **849**, 110
- Sánchez-Sáez, P., Lira, P., Mejía-Restrepo, J., et al. 2018, *ApJ*, **864**, 87
- Sánchez-Sáez, P., Reyes, I., Valenzuela, C., et al. 2021, *AJ*, **161**, 141
- Seargle, J. D., Norris, J. P., Jackson, B., & Chiang, J. 2013, *ApJ*, **764**, 167
- Simm, T., Salvato, M., Saglia, R., et al. 2016, *A&A*, **585**, A129
- Sobolewska, M. A., Papadakis, I. E., Done, C., & Malzac, J. 2011, *MNRAS*, **417**, 280
- Stern, D., McKernan, B., Graham, M. J., et al. 2018, *ApJ*, **864**, 27
- Suberlak, K. L., Ivezić, Ž., & MacLeod, C. 2021, *ApJ*, **907**, 96
- Tachibana, Y., Graham, M. J., Kawai, N., et al. 2020, *ApJ*, **903**, 54
- Tachibana, Y., & Miller, A. A. 2018, *PASP*, **130**, 128001
- Trakhtenbrot, B., Arcavi, I., MacLeod, C. L., et al. 2019a, *ApJ*, **883**, 94
- Trakhtenbrot, B., Arcavi, I., Ricci, C., et al. 2019b, *NatAs*, **3**, 242
- Ulrich, M.-H., Maraschi, L., & Urry, C. M. 1997, *ARA&A*, **35**, 445
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- van Roestel, J., Duev, D. A., Mahabal, A. A., et al. 2021, *AJ*, **161**, 267
- Van Rossum, G., & Drake, F. L., Jr 1995, Python Reference Manual (Centrum voor Wiskunde en Informatica Amsterdam), <https://ir.cwi.nl/pub/5007>
- Vanden Berk, D. E., Wilhite, B. C., Kron, R. G., et al. 2004, *ApJ*, **601**, 692
- Villar, V. A., Cranmer, M., Berger, E., et al. 2021, *ApJS*, **255**, 24
- Villar, V. A., Hosseinzadeh, G., Berger, E., et al. 2020, *ApJ*, **905**, 94
- Waskom, M., Botvinnik, O., O’Kane, D., et al. 2017, mwaskom/seaborn: v0.8.1 (September 2017), v0.8.1, Zenodo doi:10.5281/zenodo.883859
- Wenger, M., Ochsenbein, F., Egret, D., et al. 2000, *A&AS*, **143**, 9
- Yang, S., Yu, X., & Zhou, Y. 2020, in 2020 Int. Workshop on Electronic Communication and Artificial Intelligence (IWECAL) (Piscataway, NJ: IEEE), 98
- York, D. G., Adelman, J., Anderson, J. E., Jr., et al. 2000, *AJ*, **120**, 1579