



HAL
open science

Ask The Machine: Systematic detection of wind-type outflows in low-mass X-ray binaries

D Mata Sánchez, T Muñoz-Darias, J Casares, M Huertas-Company, G Panizo-Espinar

► **To cite this version:**

D Mata Sánchez, T Muñoz-Darias, J Casares, M Huertas-Company, G Panizo-Espinar. Ask The Machine: Systematic detection of wind-type outflows in low-mass X-ray binaries. *Monthly Notices of the Royal Astronomical Society*, 2023, 524 (1), pp.338-350. 10.1093/mnras/stad1895 . hal-04160208

HAL Id: hal-04160208



<https://hal.science/hal-04160208>

Submitted on 24 May 2024

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.

Ask the machine: systematic detection of wind-type outflows in low-mass X-ray binaries

D. Mata Sánchez ^{1,2★}, T. Muñoz-Darias ^{1,2}, J. Casares,^{1,2} M. Huertas-Company^{1,2,3}
and G. Panizo-Espinar^{1,2}

¹*Instituto de Astrofísica de Canarias, E-38205 La Laguna, Tenerife, Spain*

²*Departamento de astrofísica, Universidad de La Laguna, E-38206 La Laguna, Tenerife, Spain*

³*LERMA, Observatoire de Paris, PSL Research University, CNRS, Sorbonne Universit'es, UPMC Univ. Paris 06, F-75014 Paris, France*

Accepted 2023 June 20. Received 2023 June 19; in original form 2023 May 18

ABSTRACT

The systematic discovery of outflows in the optical spectra of low-mass X-ray binaries opened a new avenue for the study of the outburst evolution in these extreme systems. However, the efficient detection of such features in a continuously growing data base requires the development of new analysis techniques with a particular focus on scalability, adaptability, and automatization. In this pilot study, we explore the use of machine learning algorithms to perform the identification of outflows in spectral line profiles observed in the optical range. We train and test the classifier on a simulated data base constructed through a combination of disc emission line profiles and outflow signatures, emulating typical observations of low-mass X-ray binaries. The final, trained classifier is applied to two sets of spectra taken during two bright outbursts that were particularly well covered, those of V404 Cyg (2015) and MAXI J1820+070 (2018). The resulting classification gained by this novel approach is overall consistent with that obtained through traditional techniques, while simultaneously providing a number of key advantages over the latter, including the access to low-velocity outflows. This study sets the foundations for future studies on large samples of spectra from low-mass X-ray binaries and other compact binaries.

Key words: software: data analysis – stars: black holes – X-rays: binaries.

1 INTRODUCTION

The field of X-ray binaries has grown steadily in the last 60 years since the discovery of Sco X-1, the first and brightest stellar X-ray source in the night sky (Giacconi et al. 1962). Over 200 low-mass X-ray binaries (LMXBs, a binary where a compact star accretes mass from a Roche lobe filling companion star via an accretion disc) have so far been identified (see, e.g. Liu, van Paradijs & van den Heuvel 2007) and about a third are proposed to harbour black holes (see Corral-Santana et al. 2016; Tetarenko et al. 2016). In spite of over a thousand papers dedicated to the topic over the past six decades, the evolution of an LMXB outburst (epochs where the brightness of the system increases dramatically over the course of a few weeks to years) remains a key outstanding question in the field. Previous works have established that the outburst properties are governed by two intrinsically connected processes, namely the accretion and ejection of mass (see Fender & Muñoz-Darias 2016 for a review). The two outflow signatures traditionally studied are radio jets, detected during the hard state (e.g. Fender, Belloni & Gallo 2004), and X-ray winds, mostly detected during the soft state (e.g. Neilsen & Lee 2009; Ponti et al. 2012; Díaz Trigo & Boirin 2016; Ponti et al. 2016). More recently, recurrent outflows in the optical spectra of V404 Cygni during its 2015 outburst (Muñoz-Darias et al. 2016; Mata Sánchez et al. 2018; hereafter MD16 and MS18, respectively;

see also Casares et al. 2019; Muñoz-Darias & Ponti 2022) were discovered. Their simultaneous detection with jet emission proved that these two mass ejection channels coexist in the hard state. Since then, optical outflows have been detected in up to nine LMXBs (see the most recent compilation in Panizo-Espinar et al. 2022), effectively proving it is a common feature of the population. These low-ionization (a.k.a. cold) winds are primarily detected during the hard and soft-intermediate states of the outburst (e.g. Muñoz-Darias et al. 2019, hereafter MD19; also Mata Sánchez et al. 2022), though near-infrared observations suggest they might also be present during the soft state (Sánchez-Sierras & Muñoz-Darias 2020). The main tools employed for the identification of outflow features in optical and near-infrared spectra are limited to visual inspection or the so-called excesses diagnostic diagram (MS18). While the former typically relies on simultaneous detection of a number of features (including blue-shifted absorptions, extended emission wings, skewed profiles, and flat-top lines) over different emission lines, the latter assumes a Gaussian-like underlying component from the accretion disc contribution at the wings of the profile (see also Panizo-Espinar et al. 2021). However, as the population of systems with detected outflows continues to grow, it is required to develop automatic methods for the identification and classification of these features, in order to provide scalable and reproducible solutions.

Machine learning (ML) methods have revolutionized most scientific fields over the past decade. In astrophysics, these techniques have also been adopted, which proved to be particularly well suited

* E-mail: matasanchez.astronomy@gmail.com

to analyse large data bases accumulated over years (e.g. Ball & Brunner 2010; Huertas-Company & Lanusse 2023). The arguably most widespread use case is the automatic classification of vast observing samples, to the point that ML has become a standard part of large survey discovery pipelines (e.g. Brink et al. 2013; Wright et al. 2015; Killestein et al. 2021). For the particular field of LMXBs, the application of these techniques has just started to permeate, being so far focused on their X-ray emission, attending to their spectral and timing properties (e.g. Pattnaik et al. 2021; de Beurs et al. 2022; Orwat-Kapola et al. 2022; Ricketts et al. 2023).

The present work aims at constructing a neural network (NN) with the purpose of automatically classifying spectroscopic features traditionally associated with the presence of wind-type outflows. As a pilot study, this work provides proof of concept on the methodology, technical feasibility, and limitations of the technique. We train the NN on a simulated data base of spectral emission lines generated from theoretical disc models of LMXBs, to which we inject a variety of outflow features (Section 2). This allows us to test the performance of supervised ML methods on the simulated data (Section 3), after which we apply the trained algorithm to observational data from outbursts of two different LMXBs (Section 4). Comparison of the derived results with those obtained through traditional techniques enables us to determine if the ML approach is a convenient tool for the study of outflow spectral features.

2 DATA BASES

This pilot study employs two types of data bases, each serving to a different aim. The purpose of simulated data bases is to train the supervised ML algorithms at the heart of the classification. On the other hand, observational data bases enable us to test the trained classifiers on a real-case scenario, as well as to compare the outcome with traditional techniques dedicated to outflow detection. It follows a detailed description of each of these samples.

2.1 Simulated data bases

We construct a simulated data base, where we emulate emission line profiles arising in LMXB accretion discs, including typical outflow features identified in their observational spectra. Different components must be implemented in order to generate a realistic sample: an underlying disc profile reproducing the classic double-peaked emission line; the outflow features that we aim to detect and classify (such as P-Cygni profiles); and finally, line contaminants that might bias the ML performance (such as known interstellar/telluric lines, or emission lines from distinct but nearby species). We will create separate data bases for the most frequently observed emission lines in the optical range of LMXBs: $H\alpha$, He I 5876, $H\beta$, and He II 4686. The simulated data bases contain a million spectra each, equally split over five different classes: *disc*, *blue-absorption*, *broad*, *P-Cygni*, and *absorbed*. We provide below a detailed description of these classes and their associated parameters. Examples of generated line profiles for each class are displayed in Fig. 1 for the reader convenience.

2.1.1 Disc emission line profile

The optical spectra of LMXBs are dominated by the contribution from their accretion discs, which generates a broad emission line profile either single-peaked or double-peaked, depending on the binary parameters (e.g. orbital inclination) as well as the spectral

resolution of the observations (defined by the particular instrumental set-up).

We employ the model described in equation (4) of Orosz et al. (1994), adapted from Horne & Marsh (1986), to generate the pure *disc* emission line profile. We note that this model was originally developed to reproduce quiescent accretion discs, while the aim of the present project requires comparison with line profiles produced during the outburst (as optical outflows have been detected only during these brighter epochs). During the outburst, the accretion disc expands and heats up, becoming brighter while pushing the region responsible for the optical line formation to an outer (cooler) radius. Under the assumption of a Keplerian distribution of velocities for the gas in the disc, the velocity range of the line-forming region during the outburst should shrink, ultimately leading to narrower line profiles. This is considered the canonical profile evolution, and it has been observed in many LMXBs (e.g. *Swift* J1357.2–0933, Corral-Santana et al. 2013; Mata Sánchez et al. 2015; Torres et al. 2015; MAXI J1820+070, MD19, Torres et al. 2020). Nevertheless, deviations from this scenario have also been reported, with the presence of outflows as a proposed explanation for the widening of the line profile during an outburst (e.g. V404 Cygni, MS18). For these reasons, we decided to employ the aforementioned models to generate our simulated data base, but covering an adequate range of parameters to generate narrow enough profiles comparable to those observed during outburst events. The limitations and biases introduced by this decision are further discussed in Section 5.

The disc model presented here depends on a series of physical parameters for the binary system. To generate a prolific enough data base in order to emulate all kinds of LMXB profiles, we define the following ranges for the different parameters at play:

- (i) α : following Orosz et al. (1994), we set the range for the exponent of the disc emissivity power law to 1.5–1.6. We note that the effect of this parameter on the line profile is small for the inspected range.
- (ii) r_1 : the innermost radius of the accretion disc, in units of the outer disc radius, is suggested to vary between 0.05–0.15 (Orosz et al. 1994). This parameter limits the maximum velocities reached at the wings of the profile and constrains the simulated profiles within a reasonable range.
- (iii) i : the orbital inclination of the system is one of the key parameters controlling the width of the line profile. We allowed this parameter to explore the full range of possible values (0–90°) so as to generate both single- and double-peaked profiles.
- (iv) v_d : the outermost disc velocity effectively sets the double peak separation. We allow it to vary between 300–1200 km s⁻¹, a range that encompasses typical values for outburst spectra (Orosz et al. 1994; Orosz & Bailyn 1995).
- (v) β and ϕ_{orb} : these two parameters allow one to emulate the effect of a hotspot in the data. They essentially modify the relative intensity between the red and blue peaks as a function of the orbital phase. We cover ranges of $\beta = 0.0$ –0.2 (up to a 20 per cent intensity ratio between the two peaks) and $\phi_{\text{orb}} = 0.0$ –1.0 (i.e. all possible orbital phases).

The above parameters describe *disc*, normalized spectra. However, to reproduce experimental spectra, a handful of additional parameters are required:

- (i) μ : the centroid of the line profile, which allows us to simulate the systemic velocity. We cover the range ± 250 km s⁻¹, consistent with the observed kick velocities in LMXBs (Atri et al. 2019).

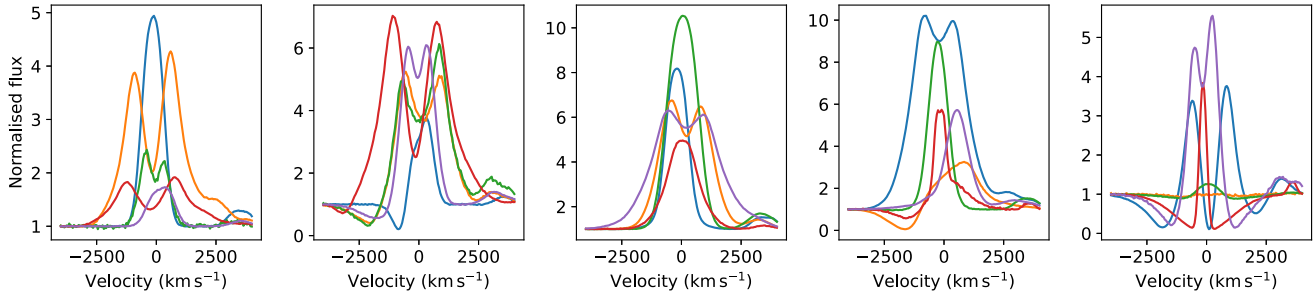


Figure 1. Examples of the final, simulated spectra for the $H\beta$ line, using the instrumental setup corresponding to GTC R2500R. Each panel corresponds to the five different classes defined in this pilot study. From left to right: *disc*, *blue-absorption*, *broad*, *P-Cygni*, and *absorbed*. Note that the emission line of He I 4922 (located at $\sim 3700 \text{ km s}^{-1}$) has been also included in the data base, to better simulate the real observations.

(ii) h : the height of the normalized line profile ultimately depends on the relative contribution between the emission line and the underlying continuum. As this ratio is known to change during the outburst (e.g. due to differing ionization levels), we decided to model a range of heights consistent with those typically observed in LMXBs. In particular, depending on the modelled emission line, the maximum allowed height was different, being the largest value for $H\alpha$ (<20) and the lowest for He I 5876 (<2.0). The minimum value was defined in terms of the signal-to-noise ratio (SNR; see Section 2.1.5), imposing a minimum $>3\sigma_N$ detection.

(iii) $\text{FWHM}_{\text{broad}}$: a number of effects, whether intrinsic or external to the LMXB, can lead to an additional broadening of the profile (characterized by its full width at half-maximum, $\text{FWHM}_{\text{broad}}$). The most straightforward is the instrumental set-up, which inevitably introduces a broadening component in the observed spectrum due to the limited instrumental resolution ($\text{FWHM}_{\text{inst}}$; we used either 200 km s^{-1} or 300 km s^{-1} for this pilot study, see Section 2.2). Additionally, an intrinsic broadening component when comparing the disc models and observations has been recently reported (Casares et al. 2022). Even though the origin of this local broadening remains unclear, we simulate it by allowing for $\text{FWHM}_{\text{local}}$ values in the range $50\text{--}650 \text{ km s}^{-1}$ (following the results presented in Casares et al. 2022). We combine both effects by convolution of the disc spectrum with a Gaussian kernel of $\text{FWHM}_{\text{broad}} = \sqrt{\text{FWHM}_{\text{inst}}^2 + \text{FWHM}_{\text{local}}^2}$.

Each combination of the above parameters generates a single emission line profile which is evaluated in the radial-velocity space, defined with respect to the corresponding emission line rest wavelength. We decided to cover a range of $\pm 4000 \text{ km s}^{-1}$ (which include the fastest outflow profiles ever detected, see MD16 and MS18), and impose a pixel size determined by the dispersion of the observational set-up (either 50 or 100 km s^{-1} , see Section 2.2).

We then use a Monte Carlo approach to generate, for a given observing set-up, a data base containing a million of simulated spectra. The parameter values employed to construct each spectrum were randomly drawn from uniform distributions within the aforementioned limits, except for the orbital inclination, for which a uniform distribution in $\cos(i)$ was employed instead (corresponding to an isotropic distribution of i). Out of the complete data base generated and with the aim to define five balanced classes, 20 per cent will be kept untouched as pure *disc* profiles, which will be classified as *disc* type spectra. The remaining 80 per cent spectra will be modified to describe the four remaining classes detailed below.

2.1.2 Outflow features

The observational signatures of wind-type outflows can result in a wide variety of profiles in the optical spectra, depending on a number of conditions such as the outflow geometry, its density and temperature. Complex outflow models are out of the scope of this pilot study, which aims at detecting typical features generated by either optically thin or thick ejecta. As such, we focus on reproducing the most commonly observed profiles:

(i) *Blue-absorption*: an optically thick outflow launched towards the observer produces an absorption at blue-shifted wavelengths (with respect to that of the transition). We emulate this feature with a Gaussian of negative height (between -50 per cent of the *disc* profile normalized flux and the $3\sigma_N$ detection level), with blue-shifted centroid (covering the range of -500 to -2500 km s^{-1} , relative to the centroid of the *disc* profile μ), and σ between that constrained by the spectral resolution, and the v_d value (but always below 1000 km s^{-1} , to avoid unrealistically wide profiles).

(ii) *Broad*: an optically thin, isotropic and homogeneous nebula would produce instead an emission component across a broad range of projected velocities. This has been observed in a number of systems experiencing outflows, such as nova events (e.g. Iijima & Esenoglu 2003), and more recently in LMXBs outbursts (MD16). We implement this feature as a Gaussian in emission with height within the range between $3\sigma_N$ detection and 4 (up to 20 for $H\alpha$, as it has shown the most prominent examples); centroid fixed at the *disc* profile μ value, and σ limited to the range between $v_d/2$ and $v_d/2 + 1000 \text{ km s}^{-1}$ (up to $v_d/2 + 1300 \text{ km s}^{-1}$ for $H\alpha$).

(iii) *P-Cygni*: the most robust outflow signatures are arguably P-Cygni profiles, present in a variety of systems (e.g. Smith & Hartigan 2006; Thöne et al. 2017). It is naturally produced by an homogeneous and isotropic expanding shell of matter, and produces a blue-shifted emission (due to the material ejected towards the observer) superimposed to a broad emission component (produced through line recombination over the entire shell). We will implement this profile as a combination of the two features described above (*blue-absorption* and *broad*), using the exact same range of values.

These features are injected into the *disc* profiles by randomly selecting parameter combinations from uniform distributions constrained within the previously defined limits. For each of the outflow classes defined above, we implement the corresponding outflow feature in 20 per cent of the simulated spectra (i.e. 60 per cent of the simulated data base contains some kind of outflow signature).

2.1.3 Broad absorption

Broad absorption components have been observed in a number of outbursting LMXBs, such as Nova Muscae 91 (della Valle, Jarvis & West 1991), GRO J0422+32 (Callanan et al. 1995; Casares et al. 1995; Shrader et al. 1997), Nova Velorum 93 (della Valle et al. 1997; Masetti, Bianchini & della Valle 1997), GRS J1655–40 (Bianchini et al. 1997; Soria, Wu & Hunstead 2000), XTE J1118+480 (Dubus et al. 2001), MAXI J1807+132 (Jiménez-Ibarra et al. 2019), MAXI J1803–298 (Mata Sánchez et al. 2022), and more recently MAXI J1348–630 (Panizo-Espinar et al. 2022); to name a few. It can be described as a broad absorption where the emission line is embedded, with varying depth and width as the outburst evolves. Depending on its depth and centroid velocity (with respect to the emission line), it can give rise to a profile which might be confused with the *blue-absorption* outflow case. The origin of this feature is still under debate, with the most accepted explanation invoking self-absorption by the accretion disc atmosphere (Dubus et al. 2001). The automatic identification of this feature will be a key tool to study its origin through future population studies, but it is also required in the present paper to avoid flagging it as an outflow. We simulate it as a low velocity, broad absorption Gaussian with centroid in the range of $\mu \pm 400 \text{ km s}^{-1}$, height varying between -50 per cent of the *disc* profile normalized flux at the absorption centroid and the $3\sigma_N$ detection level, and σ between $2v_d$ (with an absolute minimum of 300 km s^{-1}) and 1100 km s^{-1} (up to 1300 km s^{-1} for $\text{H}\alpha$). We apply such broad absorption feature to the remaining 20 per cent of the generated spectra.

2.1.4 Line contaminants

Real spectral lines of LMXBs might exhibit other features close to the transition of interest that effectively contaminate the observations. In this work, we decided to simulate separate data bases for the following emission lines: $\text{H}\alpha$, He I 5876, $\text{H}\beta$, and He II 4686. The main contaminants for each of these lines, as well as the process to include them in the simulated data bases, are as follows:

(i) $\text{H}\alpha$: a telluric line is present at 6613 \AA , which could influence the profile of particularly broad $\text{H}\alpha$ lines. For this reason, we simulated the contaminant as a Gaussian absorption component at the corresponding wavelength, σ defined by the spectral resolution, and height randomly picked from a uniform distribution between -0.1 and 0.0 .

(ii) He I 5876: the main contaminant to this line is the Na doublet at 5890 \AA , a pair of absorption lines with a mixed origin (telluric and interstellar). We simulate them as two Gaussians centred at the corresponding wavelengths, with σ determined by the spectral resolution, and height randomly picked from a uniform distribution between -0.4 and -0.2 .

(iii) $\text{H}\beta$: the He I line at 4921.929 \AA , even if typically weaker, is close enough to $\text{H}\beta$ to influence its profile on certain spectra. Therefore, we simulate it as an accretion disc line with the exact same model parameters employed for the $\text{H}\beta$ line, but allowing for slightly wider peak separation (randomly picked from a uniform distribution between v_d and $1.5 v_d$), due to the line being produced at a slightly different radius, and smaller height (0.07 – 0.2 times that of $\text{H}\beta$).

(iv) He II 4686: this is the trickiest line of our sample, due to the presence of both the He I 4713 line, and the Bowen blend at $\sim 4640 \text{ \AA}$. The former is typically weak, and we model it as an accretion disc line with the same parameters as He II 4686, but height within the

0.05 – 0.25 range. The latter is a mixture of C III and N III emission lines, which appears blended into a single broad component close to He II 4686 and of comparable intensity. We attempt to simulate this complicated profile as a broad Gaussian emission line, with σ between 2000 – 2700 km s^{-1} , and height of 0.1 – 0.55 times that of He II 4686.

2.1.5 Simulated noise

In order to create a more realistic data base, we decided to add noise to the simulated spectra. We implemented it by using a single parameter: SNR. For a given spectrum, at a particular velocity pixel, we sample a value from a uniform distribution in the range of $\text{SNR} = 30$ – 330 , and we create a Gaussian distribution centred at the normalized spectrum flux value (f_{norm}) with a standard deviation of $\sigma_N = \sqrt{f_{\text{norm}}/\text{SNR}}$. We randomly draw samples from this distribution, which are added on top of the original f_{norm} of the *disc* spectrum. Note that this approximation assumes a constant underlying continuum under the emission line, a reasonable assumption for such a small wavelength range. It also produces noise with $\sigma_N \sim 1/\text{SNR}$ at the normalized continuum level while improving on the relative uncertainties at the peak of the emission line.

2.1.6 A zoo of profiles

During the generation of the simulated data base, a number of the randomly generated line profiles attracted our interest. The most conspicuous among them are shown in Fig. 2, and include flat-top, triangular, and skewed emission lines. These have been previously observed in LMXBs and associated to the presence of outflows (e.g. Sánchez-Sierras & Muñoz-Darias 2020, Panizo-Espinar et al. 2022). We are aware that similar profiles have also been observed in other systems with outflows but without accretion discs (e.g. supernovae, see Smith, Mauerhan & Prieto 2014), which implies that the double-peaked component is not mandatory to produce the aforementioned outflow signatures. Indeed, these can arise as a result of different combinations of outflow geometries and optical thickness. Nevertheless, this also implies that our simulated data base, in spite of not employing a physical model to simulate the outflow components, is already able to produce such profiles. While describing the physical formation mechanism of the outflow profiles is beyond the scope of this paper, we wanted to highlight that a simple combination of no more than two Gaussians (respectively in absorption and emission) on top of the double-peaked *disc* profile is already able to naturally produce these characteristic shapes, and therefore, they will be considered during the algorithm training.

2.2 Observational data sets

In order to test the ML algorithm trained on the simulated data base, comparison with real, observed normalized spectra from LMXBs is required. For this pilot study, we decided to focus on two systems for which particularly prolific data bases were compiled during their latest outbursts.

(i) V404 Cygni: This LMXB, which contains the first dynamically confirmed stellar-mass black hole (Casares, Charles & Naylor 1992) exhibited an extremely bright outburst in 2015 after ~ 25 years of quiescence (Barthelmy et al. 2015). Intensive multi-wavelength follow up of the event led to dozens of publications (e.g. Kimura et al. 2016; Motta et al. 2017). Of particular relevance for this work is the first systematic analysis of the appearance and evolution of

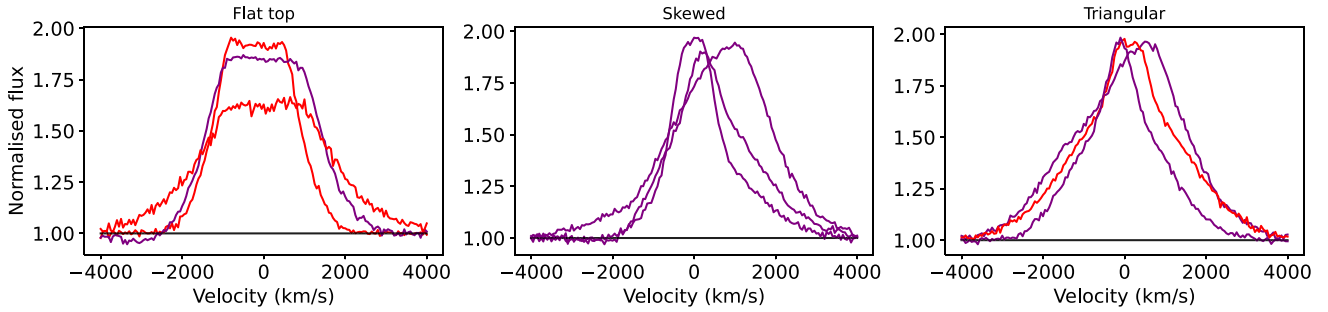


Figure 2. Simulated line profiles randomly generated with peculiar shapes reminiscent of more complex outflow-related profiles. Red colour corresponds to *broad* type and purple to *P-Cygni* type simulated spectra.

optical outflows in a LMXB (MD16), which sparked the interest of the community in such features and their influence in the outburst evolution. In this pilot study, we will limit our analysis to a subsample of the data base presented in MS18. In particular, we only consider the 510 spectra obtained with the Gran Telescopio Canarias (GTC) and grism R1000B, which cover the first 15 days of the outburst.

(ii) MAXI J1820+070 (hereafter MAXI J1820): the discovery outburst of this LMXB (Tucker et al. 2018) was also followed with great interest, due to both its brightness and accessibility from all over the globe. The numerous studies dedicated to this event allowed to unveil both a high orbital inclination and a black hole nature for the compact object (see, e.g. Torres et al. 2019, 2020). The detection of outflow features in both optical and near-infrared spectra has also been reported (MD19, Sánchez-Sierras & Muñoz-Darias 2020), and a significant data base of spectra has been collected. In this pilot study, we focus on the subsample of 48 spectra observed with GTC and grism R2500R presented in MD19, which cover both the hard and soft state of the outburst spanning over eight months.

Both data bases have been previously described thoroughly in dedicated studies to each object, so we refer the reader to the aforementioned papers in order to learn the data reduction details.

2.3 Data bases homogenization

Direct comparison of the simulated and observational data bases is not straightforward. To this end, we extracted cuts of the observed spectra centred at the lines of interest: $H\alpha$, He I 5876, $H\beta$, and He II 4686. We binned all the spectra in velocity space to a uniform pixel size consistent with the dispersion of each instrument (50 and 100 km s⁻¹ for R2500R and R1000B, respectively) and covering the range of ± 4000 km s⁻¹. We then created separate, simulated data bases using the corresponding dispersion values mentioned above and the associated spectral resolution with each grism configuration (FWHM_{inst} = 200 km s⁻¹ and 300 km s⁻¹ for R2500R and R1000B, respectively).

3 METHODS AND RESULTS

The aim of this work is to apply ML techniques to identify outflow features in LMXBs emission lines. Thanks to the explosive growth of the ML field and its application across different disciplines, new tools have been developed, tested, and made available to the community. All the ML classifiers implemented in this work belong to the NN type and have been developed making use of the open source deep

learning library KERAS (Chollet et al. 2015) with the TENSORFLOW (Abadi et al. 2015) back-end.

3.1 NNs classifiers

We decided to focus on Deep Neural Networks (DNNs), which, since their original success for solving computer vision problems (Krizhevsky, Sutskever & Hinton 2012), have been applied to an extensive range of cases (see e.g. Szegedy et al. 2015). DNNs performance seems to particularly excel on speech recognition (e.g. Sainath et al. 2013) and natural language processing tasks (e.g. Mikolov et al. 2013), both of which are characterized by their sequential nature, which is also a common property for time series and spectra.

We explore three types of supervised classifiers of the DNN class, aiming to classify our simulated data base into the five different classes defined in Section 2.1. We provide below a short description of the parameters defining the different layers compounding each algorithm, while referring the reader to Ismail Fawaz et al. (2018) and references therein for an in-depth description of the algorithms themselves:

(i) Multi layer perceptrons (MLP): this network treats independently time series/spectral elements from each other, which means the sequential aspect of the data is ignored. We constructed a network containing four dense layers whose outputs are fully connected. They are three hidden layers of 500 neurons and Rectified Linear Unit (ReLU) activation function, and a final softmax classifier. Each layer is preceded by a dropout (Srivastava et al. 2014) with rates of 0.1, 0.2, 0.2, and 0.3, respectively, a form of regularization preventing overfitting.

(ii) Fully convolutional neural network (FCN): a type of convolutional neural network, proposed in Wang, Yan & Oates (2017) and composed of three convolutional blocks. Each consists of a convolution, followed by a batch normalization (Ioffe & Szegedy 2015; a normalization layer to help the network converge quickly), and a ReLU activation function. The three convolution layers have 128, 256, and 128 filters; with respective lengths of 8, 5, and 3.

(iii) Residual Network (ResNet): this architecture, also proposed by Wang et al. (2017), consists of three residual blocks. Each of them is composed by three convolution layers with a ReLU activation function and preceded by a batch normalization, whose output is fed to the input of the residual block. The convolutions within a residual block all have the same number of filters (which is 64, 128, or 128, for each different residual block), and different filter lengths (8, 5, and 3; for each set of convolution layers within the block).

Table 1. Performance metrics for the MLP, FCN, and ResNet classifiers evaluated over the test data set, corresponding to the averaged values over all classes.

MLP					
Line	Grism	Precision	Recall	F ₁	Accuracy
H α	R1000B	0.85	0.82	0.81	0.82
	R2500R	0.81	0.78	0.77	0.78
He I 5876	R1000B	0.91	0.90	0.90	0.90
	R2500R	0.88	0.87	0.87	0.87
H β	R1000B	0.90	0.88	0.88	0.88
He II 4686	R1000B	0.83	0.81	0.82	0.81
FCN					
Line	Grism	Precision	Recall	F ₁	Accuracy
H α	R1000B	0.95	0.95	0.95	0.95
	R2500R	0.96	0.96	0.96	0.96
He I 5876	R1000B	0.94	0.94	0.94	0.94
	R2500R	0.95	0.95	0.95	0.95
H β	R1000B	0.92	0.92	0.92	0.92
He II 4686	R1000B	0.89	0.89	0.89	0.89
ResNet					
Line	Grism	Precision	Recall	F ₁	Accuracy
H α	R1000B	0.95	0.95	0.95	0.95
	R2500R	0.97	0.96	0.96	0.96
He I 5876	R1000B	0.95	0.94	0.94	0.94
	R2500R	0.96	0.96	0.96	0.96
H β	R1000B	0.94	0.93	0.93	0.93
He II 4686	R1000B	0.93	0.92	0.92	0.92

3.2 Training and performance assessment

We independently trained the three algorithms described above on the simulated spectral data bases previously introduced. This implies the independent training of six data bases: four for the grism set-up R1000B (one per line profile: H α , He I 5876, H β , and He II 4686) and two for the R2500R (H α , He I 5876). As part of the training, we evenly split each data base into training and test samples. We flagged 30 percent of the training data set as a validation data set in order to assess the training performance. We compiled the models with an ADAM optimizer and a CATEGORICAL CROSSENTROPY loss function. We set an early stopping condition defined by the lack of improvements over the last 50 epochs on the loss function, measured in the validation set. The performance is then evaluated by applying the trained classifier to the unseen test data base, which produces the metrics reported in Table 1 (see also an example of a confusion matrix in Fig. 3). We note that the actual output of the NN is an array containing the probability for each particular spectrum to be associated with each of the five classes (as determined by the final softmax layer). We assign to each processed spectrum the class with the highest probability for the metric comparison purposes described above. Inspection of the individual Receiver Operating Characteristic (ROC) curves for each class and set-up reveals area under the curve (AUC) values higher than 0.98. We also explored the thresholds on the classification probability based on the ROC curves that optimize the true positive rate (tpr) and the false positive rate (fpr), and derived individual results for each set-up, always keeping tpr > 0.93 and fpr < 0.07. The resulting threshold values are comparable between line types and set-ups, but vary significantly between the different classes, with 0.5 (*absorbed*) and 0.1 (*P-Cygni*) being at the extremes, coinciding with the best and worst performing classes, respectively. To simplify our selection criteria, we will hereafter keep the most

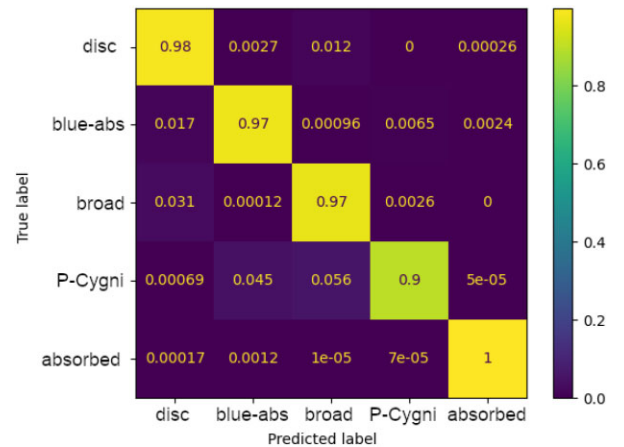


Figure 3. Confusion matrix for the ResNet classifier when applied to the H α R2500R data base. Each cell value has been normalized by the total number of spectra per class.

restrictive threshold value of 0.5. Classifications below this threshold should be taken with caution, and when possible, independently confirmed.

The comparison of the performance of the different methods (attending only to their average metrics) reveals that the MLP script is producing the less accurate classification, down to 0.78 for some data bases. The average accuracy of both the FCN and ResNet algorithms is significantly better (0.89–0.97), with a slightly better performance favouring the ResNet case (0.92–0.97). A closer inspection of each individual confusion matrix shows that the main bulk of miss-classified spectra actually corresponds to blends of outflow classes (i.e. *blue-absorption*, *broad* and *P-Cygni* class). Spectra from the *P-Cygni* type leak into the *blue-absorption* and *broad* classes (with typical recall values for the *P-Cygni* class of ~ 0.85 – 0.90), or alternatively, *blue-absorption* and *broad* class spectra are classified as *P-Cygni* (with recalls of ~ 0.90 – 0.97 and ~ 0.92 – 0.97 for the *blue-absorption* and *broad* classes, respectively). This is not unexpected, as the *P-Cygni* class is generated as a combination of profiles from the two other outflow categories. Therefore, if one of the two components (either the blue-shifted absorption or the broad wing emission) is too shallow (we simulate outflow features with depth/high over the $> 3\sigma$ threshold defined by the SNR), the algorithm might not be able to classify them correctly. On the other hand, the *disc* and *absorbed* classes show a much more consistent performance, with stable recall values of > 0.96 .

To further investigate the origin of these miss-classifications, we produced histograms of the parameters employed during the data base simulation, and compared them with those incorrectly classified. This allows us to inspect possible caveats in the parameters range determination when creating the data base, as well as to explore the limitations of the ML classifiers. We conclude that their accuracy is mainly affected by the following cases:

(i) Low velocity and shallow P-Cygni profiles: the low velocity P-Cygni profiles ($\lesssim 1000 \text{ km s}^{-1}$), specially if they are also shallow, are harder to identify as they fall between the peaks of the *disc* emission line profile. Rather than a clear absorption, this effectively creates an asymmetry in the line profile, which might be confused with those produced by a hotspot.

(ii) Low height and narrow emission wings: the detection of a shallow broad wing component is specially challenging if the outflow components are relatively narrow ($\sigma \lesssim 500 \text{ km s}^{-1}$). In such a case,

Table 2. Results from applying the ATM classifier to real data bases, where we report the number of spectra belonging to each class. Note that for each combination of line and grism, ATM was trained with a slightly different simulated data base in order to better reproduce the corresponding setup conditions. We also report within parentheses those spectra whose classification is uncertain, as it falls below the 0.5 threshold described in Section 3.2.

Data base	Line	<i>Disc</i>	<i>Blue-absorption</i>	<i>Broad</i>	<i>P-Cygni</i>	<i>Absorbed</i>
V404 Cygni	H α	6 (4)	0 (0)	235 (40)	269 (4)	0
	He I 5876	203 (4)	193 (6)	48 (1)	65 (2)	1 (0)
	H β	246 (0)	80 (2)	87 (3)	87 (4)	10 (0)
	He II 4686	67 (5)	268 (3)	51 (6)	65 (3)	59 (3)
MAXI J1820	H α	36 (0)	0 (0)	12 (0)	0 (0)	0 (0)
	He I 5876	28 (1)	19 (1)	0 (0)	0 (0)	1 (0)

only the central region of the *disc* profile is altered, an effect that can be mimicked through slight modifications of the *disc* profile parameters.

(iii) Low SNR: the detection of features in noisy spectra is a challenging task, but it does not seem to be a critical parameter for the range considered (30–330), as the number of miss-classified spectra between the extremes does not vary drastically (typically a factor <3). It is only the dominant factor for the *disc* class spectra.

All of the above are expected limitations not exclusive to the ML approach. On this regard, it is worth remarking that the ML technique does not appear to introduce any new biases, other than those inherent to the data quality and particular outflow properties.

4 APPLICATION TO OBSERVATIONAL SPECTRA

After assessing the performance of the different NNs on the test and validation data sets, we select ResNet as the most efficient among them. Further discussion will be limited to results produced by this classifier, which will be hereafter referred to as the Ask The Machine (ATM) algorithm.

We applied the trained ATM classifiers to the real spectroscopic data sets described in Section 2.2. For the V404 Cygni data set, this includes four classifiers dedicated to each of the lines detected with the R1000B grism set-up, while for MAXI J1820 data, it only includes the two lines covered by the R2500R grism. A summary of the results from each classification is collected in Table 2, but also discussed below for each data set and line separately.

We will compare the ATM classification with the two most widely used traditional techniques for outflow detection: the excesses diagnostic diagram and the visual inspection. The excesses diagnostic diagram was first introduced in MS18 and further refined in Panizo-Espinar et al. 2021. It is based on the assumption that departures in the wings of the inspected profile from a Gaussian function are due to the presence of outflows. After a Gaussian fit is subtracted from the line profile and the core of the line is masked (on a case-by-case basis), equivalent widths (EWs) are measured on the residuals at regions corresponding to the red and blue wings of the profile (EW_r and EW_b, respectively). The position of a particular spectrum in the excess diagnostic diagram (EW_b against EW_r) determines the presence and class of the outflow feature. On the other hand, detection of outflow signatures through visual inspection typically requires the presence of simultaneous features associated with outflows in different lines.

4.1 V404 Cygni

The large data base collected during the 2015 outburst of V404 Cygni is extremely diverse, showing a wide range of profiles for each of the inspected lines. We report in Fig. 4 the results from applying ATM to a particular epoch (day 6, compound of 85 individual spectra and showing a large variety of outflow profiles) as a visual example of the classification results. We also exhibit the complete excesses diagnostic diagrams for the entire V404 Cygni data base on H α , He I 5876, and H β (see Fig. 5). In addition, we include their ATM classification in order to compare both techniques. It follows a detailed description for each line.

4.1.1 H α

Out of the complete sample consisting of 510 spectra, all but six are classified by ATM as having outflows in H α (see Table 2). In particular, the algorithm detects outflows of two types: either *broad* (46 per cent) or *P-Cygni* (53 per cent) profiles. The scarce number of spectra without outflow features is shocking, as such features are typically elusive in most LMXBs outburst events. Indeed, these outflows have only been detected systematically during the past few years through dedicated observations with the largest available telescopes. At this point, it is worth considering that a flawed design/training of ATM classification algorithm might be at play, biasing it to flag outflows in most spectra. However, it will become clear in the following sections that the ATM script is indeed able to identify *disc* profiles (without outflows) in other lines and targets. Furthermore, previous works on the 2015 outburst of V404 Cygni (MD16, MS18, Casares et al. 2019) have shown the dramatic evolution of the H α line profile, including extreme EWs, the detection of high-velocity outflows and a FWHM above the quiescence level during most of the event (opposite to the canonical behaviour, see Section 2.1.1). For these reasons, we believe the results from the ATM classification to be correct, and instead reveal V404 Cygni outburst as an extraordinarily rich event where outflows are always present. As a matter of fact, X-ray observations contemporaneous to those used here also show the presence of X-ray winds with similar observational properties to the optical ones (Muñoz-Darias & Ponti 2022; see also King et al. 2015).

Direct comparison of the results of the ATM with those drawn in MS18 from the excesses diagnostic diagram shows that this technique also found numerous P-Cygni-like and nebular phase-like profiles (*P-Cygni* and *broad* classes in this work, respectively). We show an updated version (following Panizo-Espinar et al. 2021 prescription) of the H α excesses diagnostic diagram in Fig. 5 (top left-hand panel), which remains consistent with the original presented in MS18. This

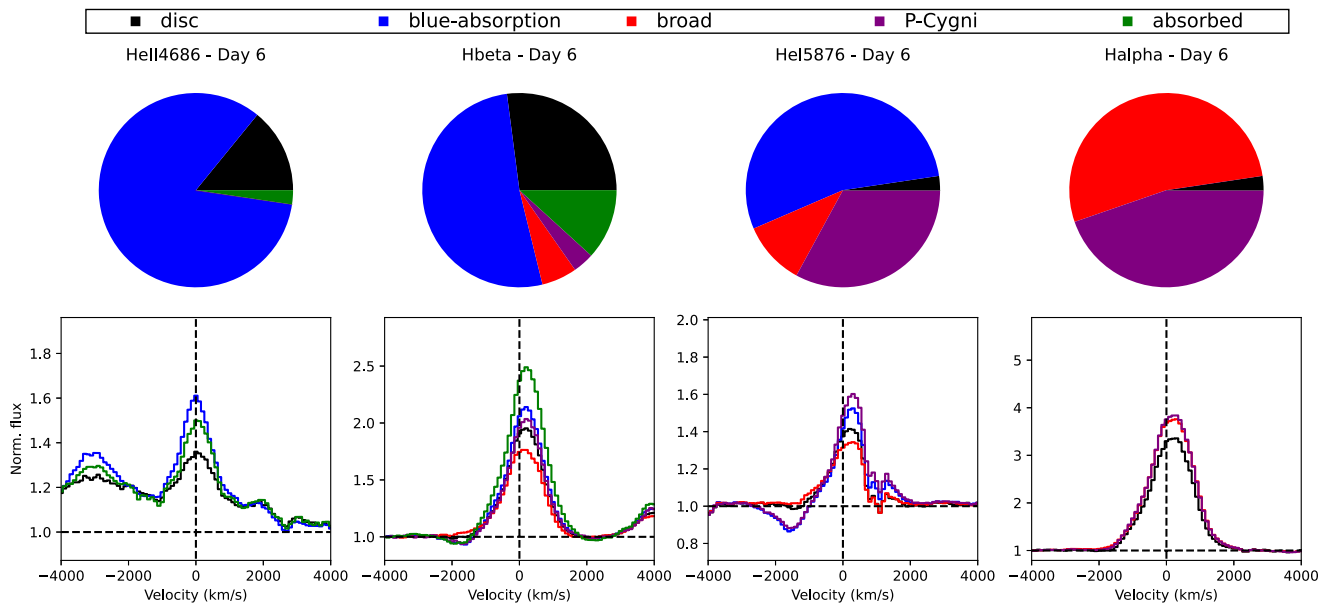


Figure 4. Top row: pie chart representation of the classification with ATM for V404 Cygni spectra observed during day 6, where each colour corresponds to a different class as described in the legend. Bottom row: the resulting averaged spectrum from the combination of all the observed spectra during this run, separated by the identified classes and following the same colour code.

reveals an overall agreement with the ATM results: *broad* spectra are mainly clustered in the top-right region of the diagram, while *P-Cygni* type spectra dominate the bottom-right region (i.e. the *P-Cygni* region). It is worth remarking that the excesses diagnostic diagram only evaluates the wings of the line by comparing them with a Gaussian model, while it disregards asymmetries of the core of the line (which is effectively masked). Therefore, it is not surprising to find spectra in the nebular phase region being classified by ATM as *P-Cygni* type: this suggests that, while broad wings are probably dominant at the edge of the profile, blue-shifted absorptions are also required to reproduce the observations. Another limitation of the traditional excesses diagnostic diagram, already highlighted in MS18, was the detection of apparent inverse *P-Cygni* profiles on day 6 (top-left region of the diagram). Visual inspection revealed they are actually caused by the presence of the intense *P-Cygni* profiles of the nearby He I 6678 line, as well as a clearly asymmetric $H\alpha$ line. On the other hand, ATM directly classifies day 6 spectra as either *broad* or *P-Cygni* types (see Fig. 4), with the sole exception of two spectra that are classified as *disc* at a probability below the threshold level of 0.5 (see Section 3.2). This shows the ability of ATM to overcome previous classification issues.

4.1.2 He I 5876

ATM classification of this line profile reveals the following distribution among classes: 40 per cent *disc* spectra, 60 per cent with outflow features (distributed as 63 per cent *blue-absorption*, 16 per cent *broad* and 21 per cent *P-Cygni* profiles), and a single spectra of the *absorbed* type. The detection of a significant number of *disc* profiles is reassuring, supporting the aforementioned argument that the continuous presence of outflows in $H\alpha$ is real, rather than derived from fundamental issues in ATM design and training. Most of the profiles classified as outflows contain a blue-shifted absorption, sometimes accompanied by a broad emission component. This is consistent with the results of the excesses diagnostic diagram (see

top right-hand panel in Fig. 5), where deep, *P-Cygni* features are detected in a significant fraction of the spectra, while no nebular phase spectra are reported outside the 5σ confidence region. The handful of spectra classified by ATM as *broad* type do lie in the nebular region of the excesses diagnostic diagram, though their inclusion within the 3σ confidence region would not trigger a detection from the traditional method alone. This might be a consequence of ATMs inspecting the full line profile rather than only the wings, eventually leading to more sensitive detection. The fact that only one spectrum was classified as *absorbed* suggests that the inclusion of a simulated Na I doublet feature in the training data set was able to mitigate potential misclassifications to a remarkable degree.

4.1.3 $H\beta$

Classification of these profiles reveals 48 per cent *disc* spectra, 50 per cent containing some kind of outflows, and 2 per cent *absorbed*. The higher ratio of *disc* spectra for $H\beta$, specially when compared with the almost null ratio for $H\alpha$, shows that the extreme behaviour observed in the latter line is not present in $H\beta$. Indeed, the Balmer decrement variability reported in MS18 is driven by the more dramatic changes in $H\alpha$ (typically $EW \sim 20\text{--}300 \text{ \AA}$, but spiking up to $EW \sim 2000 \text{ \AA}$), while $H\beta$ is typically $EW \sim 10\text{--}50 \text{ \AA}$ (with certain epochs reaching up to $EW \sim 300 \text{ \AA}$). This view is fully consistent with the reported ratio of outflow types, being 32 per cent *blue-absorption*, 34 per cent *broad*, and 34 per cent *P-Cygni*. An overall agreement can be confirmed by inspecting Fig. 5, bottom left-hand panel, where the ATM classification populates the expected regions of the diagram. Finally, 2 per cent of the spectra are classified as *absorbed*. Absorption features have been seen before in a number of LMXBs (see Dubus et al. 2001 and references therein), and while they are predominantly reported in hydrogen lines, they have also been observed in He lines (e.g. Soria et al. 2000). A visual inspection of the affected lines in our data base shows that spectra classified as

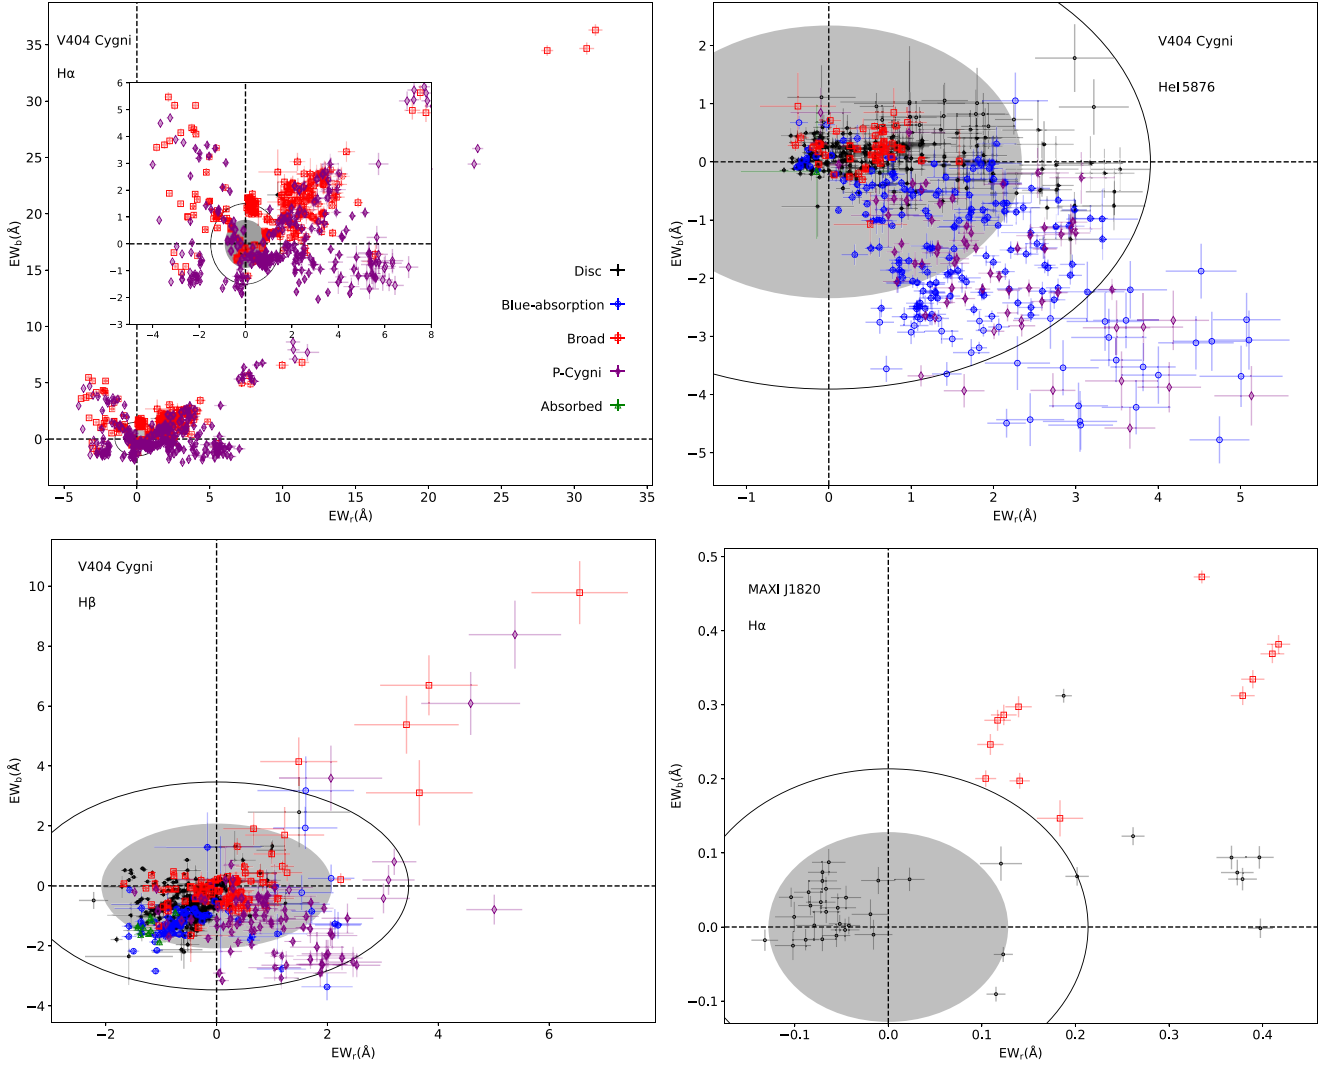


Figure 5. Excesses diagnostic diagram applied to the lines of interest in each of the individual spectra from V404 Cygni and J1820 observational data sets. We do not show He I 4686 for V404 Cygni neither H β for MAXI J1820 due to their line contaminants, which prevent a reliable application of the traditional method. The colour code matches that of Fig. 4, depicting the ATM classification for each of the spectra: *disc* (black dots), *blue-absorption* (blue empty circles), *broad* (red empty squares), *P-Cygni* (purple empty diamonds), and *absorbed* (green empty triangles). V404 Cygni correspond to the top panels, from left to right, H α (\pm)1000–4000 km s $^{-1}$; including an inset of the central region) and He I 5876 (-4000 to -500 km s $^{-1}$; 1250 to 4000 km s $^{-1}$; asymmetric due to the Na I interstellar doublet). Bottom left-hand panel also corresponds to V404 Cygni, but focused on the profile of H β (\pm)1000–2500 km s $^{-1}$). Bottom right-hand panel shows the MAXI J1820 excess diagram for the H α line (\pm)1000–4000 km s $^{-1}$). All diagrams also include the 3σ (shaded) and 5σ (solid line) ellipses defining the excess diagram detection limits.

absorbed type probably arise from the presence of *P-Cygni* profiles in the nearby He I 4922 line. This might be confused with a red-shifted component for the H β line and, when combined with a blue-shifted absorption, result in the aforementioned classification. Furthermore, the detections are all restricted to the same epoch (day 6 of the outburst, see Fig. 4), while typical broad absorption features are rather persistent, lasting from tens of minutes to days. For the above reasons, we believe that ATM classification for these spectra might not highlight an intrinsic feature of the system, and should be taken with caution.

4.1.4 He II 4686

ATM classification of this line profile resulted in 19 per cent *disc* profiles, 77 per cent with outflow components, and 4 per cent

absorbed. The large number of detected outflow profiles is surprising for a high-ionization line. If true, it would require explaining both the absence of previous outflow reports in this line, and the fact that high densities are required for this emission line to form, which hampers the visibility of outflow features (see Charles et al. 2019 for a discussion).

For this reason, we perform a visual inspection of the profiles and found that the detection of blue-shifted absorptions is heavily biased by the presence of the nearby Bowen blend. Indeed, this very contamination has hampered any attempts to produce an excesses diagnostic diagram for this line. This feature is a combination of various C, N, and O emission lines, with both narrow and broad components and blended together (McClintock, Canizares & Tarter 1975). They are thought to arise from both the companion star’s irradiated face through a fluorescence cascade (producing narrow emission line

components), or from regions in the accretion disc where similar conditions to form the lines occur (giving rise to broader components, due to the velocity distribution; see e.g. Steeghs & Casares 2002). For this reason, deriving a precise theoretical model of the Bowen blend is a challenging task. We injected into the simulated data base a broad Gaussian profile to try to emulate this contaminant, but in light of the results, we do not believe it is an acceptable approximation, specially given that the Bowen blend component intensity can be comparable to that of the line of interest. For this reason, we strongly believe that any outflow detection in He II 4686 is probably biased, and it should be taken with the utmost caution.

4.2 MAXI J1820

While MAXI J1820 spectral sample is significantly smaller than that of V404 Cygni, it enables us to test ATM in a different system and using higher resolution spectra. This allows us to inspect how ATM behaves when applied to different scenarios as well as strengthen the conclusions derived from this novel approach when compared with traditional techniques. To aid this purpose, we show the excesses diagnostic diagram for the $H\alpha$ line of the J1820 data base in Fig. 5 (bottom-right-hand panel).

4.2.1 $H\alpha$

The resulting classification of the $H\alpha$ line shows a radically different situation to that of V404 Cygni: most of the spectra are identified as *disc* (75 per cent) and the remaining (25 per cent) correspond to the *broad* class. No other outflows containing blue-shifted absorptions are found, which is roughly consistent with the overall description of the data presented in MD19 and the excesses diagnostic diagram in Fig. 5 (bottom right-hand panel; constructed from the individual spectra instead of averaging in each epoch). A closer inspection of each particular spectrum reveals some discrepancies between ATM and the classification proposed in MD19, especially for those determined from visual inspection.

Most of the differences occur during the initial hard state of the outburst, where outflows have been traditionally observed. First, MD19 reports on a P-Cygni profile for epoch-3 and broad wings for epoch-4, while ATM classifies all of them as *disc*. We note that the original paper reports that the epoch-3 P-Cygni did not follow a traditional shape. A visual inspection confirms the profile departs from our simpler models, which would explain why ATM did not pick up the feature properly. On another vein, ATM classifications of epoch-5 to -10 reveal *broad* profiles, which match the results reported in MD19 except for epoch-7 (where they report a P-Cygni component on top of the broad wings). The profile of this latter epoch is clearly dominated by its emission component, one of the strongest of all the analysed spectra. Furthermore, the confirmation of the blue-shifted absorption component in MD19 relied on the simultaneous detection of a similar outflow feature in the nearby He I 6678 line; a piece of information ATM was not trained to identify. Together, they can explain the difference in the classification.

Epoch-11 is classified as having broad wings by MD19, but ATM suggests they can be reproduced as *disc* profiles. Indeed, visual inspection of the spectra shows that the broad wings present until the previous epoch are much less prominent at epoch-11, which might either favour the results from ATM or suggest it is not able to detect such subtle features. The remaining part of the data set corresponds to soft state spectra (except for the last three low-luminosity hard state epochs), and they are classified as *disc* profiles by both MD19

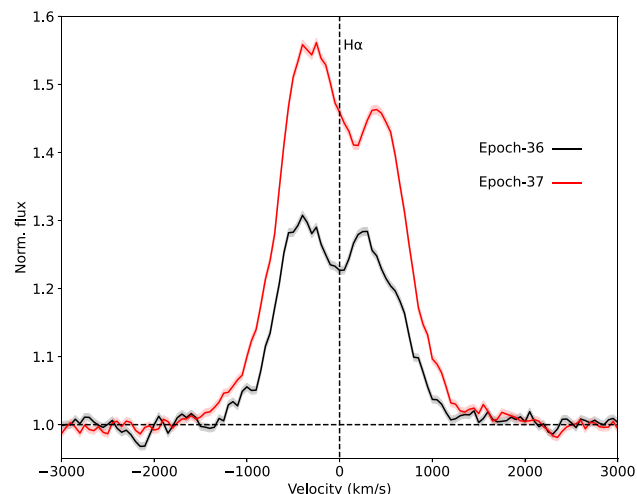


Figure 6. Observed spectra of MAXI J1820 during the low luminosity hard state, corresponding to epoch-36 and -37 as defined in MD19. Colours determine the ATM classification as in previous figures, marking them as *disc* and *broad*, respectively.

and ATM, providing further reassurance on the capabilities of the NN to not overestimate the presence of outflows in the data. This is consistent with the current picture of outflows not being detected during this accretion state in the optical regime (see, e.g. Sánchez-Sierras & Muñoz-Darias 2020). The only notable exception is the latest spectrum of the sample (epoch-37, hard state), where ATM reveals a *broad* emission component. It is worth remarking that this is the highest EW epoch of our sample, and the terminal velocities of the profile wings are similar to or larger than any of the other epochs classified as possessing broad wings. Nebular phases have been observed in a handful of systems, such as the canonical V404 Cygni (MD16, MS18), GX 339-4 (Rahoui, Coriat & Lee 2014), and more recently Aquila X-1 (Panizo-Espinar et al. 2021). They are characterized by the presence of a broad emission component typically observed in $H\alpha$, which becomes particularly prominent during the latest stage of the outburst decay (the so-called low hard state). Comparison with the closest spectrum in our sample (epoch-36, observed two weeks prior but still during the hard state decay) shows that the $H\alpha$ line in epoch-37 was remarkably stronger (see Fig. 6). The asymmetry on the line core (redshifted from the rest wavelength) also hints towards an underlying emission component. While the feature in J1820 is not as clear as in the aforementioned works (where the single-peak nature of the line eases their detection), we favour the ATM classification in this particular case.

4.2.2 He I 5876

The classification of this line profile results in 58 per cent of *disc* and 40 per cent *blue absorption* (as well as a single *absorbed* spectrum that we believe is due to normalization issues). This seems consistent with the overall reported detection in the original paper, though the latter also reported two epochs of broad wing components classified as *blue-absorption* profiles by ATM. Unfortunately, the excesses diagnostic diagram for this line was not reliable due to the narrowness of the line profile (with wings extending up to $< 1000 \text{ km s}^{-1}$), which made the red wing cross the continuum within the NaI interstellar doublet (i.e. limiting our ability to measure EW_r values). Therefore, the original paper identification relied on visual inspection and

comparison with other epochs to define them as such. However, the agnostic approach of ATM suggests that similarly broad profiles can be generated using only *disc* models, as long as a hotspot contribution is included to account for the line asymmetry. We cannot confidently claim that broad wings are unambiguously detected.

Regarding the detection of *blue-absorption* profiles, they are mainly detected during the hard state, including all epochs reported in MD19. The potential detection of *blue absorption* by ATM in epoch-8 is interesting. Visual inspection of the data obtained for this epoch, however, does not show a clear blue absorption in the data. Instead, the profile for this particular date is remarkably different from those of the rest of the outburst: it consists of a single broad peak, clearly skewed towards blue wavelengths. ATM classification is actually not particularly conclusive, with a split probability of 70–30 between *blue-absorption* and *disc* profiles. We conclude that this reflects the limitations of the simulated data base to reproduce such profiles, rather than a definitive outflow feature. The remaining detections of blue absorptions by ATM correspond to ambiguous spectra obtained during the soft state. Previous detections of optical outflows have been reported during the bright soft intermediate state (see Mata Sánchez et al. 2021; Panizo-Espinar et al. 2022), but they have yet to be found during the longer-lived soft state. On the other hand, the detection of NIR outflows during most of the outbursts in this very system (see Sánchez-Sierras & Muñoz-Darias 2020) proves that they remain at all epochs. Indeed, visual inspection also suggests the presence of a blue absorption in our highlighted spectra. Nevertheless, the ATM probability associated with the classification is not particularly high (between 40–80 per cent), which combined with the lower SNR (due to a combination of lower exposure times and the brightness of the system) forbids us from reporting a definitive confirmation.

4.3 Additional tests

We performed a handful of tests to evaluate the performance and limitations of the ATM script defined above. A non-exhaustive list with the most remarkable among them follows:

(i) Gaussian *disc* profiles: we created an independent data base, where the *disc* profiles are constructed through the combination of two Gaussians. The Gaussian parameters were selected so as to produce double-peaked profiles comparable to those observed in LMXBs, and the ML algorithm was trained on this new data base. Regarding the performance on the separation between classes for the test data set, the results were fundamentally similar to those obtained with the disc models described in the paper. This further strengthens the case for the retrainability of the NN if new, more complex disc profiles are proposed in the future. Classification of the observed spectra produced different outcomes in the final classification for a number of them, while others remained the same (e.g. most of the V404 Cygni $H\alpha$ spectra are still classified as being *broad* or *P-Cygni* types). We still favour using the disc profiles over the Gaussian approach, as the former are based on a physical description of the system, even if simplified.

(ii) Universal grism approach: we explore the possibility of training the ML classifier on a single data base, which can then be applied to both R1000B (V404 Cygni) and R2500R (MAXI J1820) spectra. The effect on the spectral resolution parameter is straightforward to include, as we are already accounting for a range of broadening values due to the local broadening component (see Section 2.1.1). The different dispersion values for each grism implies they are covering the same range of velocities with a different step

size. We created a data base with the thinnest step size out of the two grism set-ups, and allowed for a high enough range of broadening values to cover both spectral resolutions. After training the NN on this data base, we applied it to the observed data sets. Application to MAXI J1820 produced fully consistent results with those presented in the previous section (using the R2500R set-up), as expected due to the almost identical nature of the data bases. On the other hand, application to V404 Cygni required interpolation of the observed spectrum (R1000B) to a thinner velocity grid (that of R2500R) before the ML script could be applied. The overall resulting classification is similar to that obtained with the NN trained with an adequate grism set-up. However, we found a larger ratio of spectra miss-classified as *absorbed* (for $H\beta$ and He I 5876). This could be an unwanted by-product of the necessary interpolation, which might create artificial shallow features out of the noise. Taking into account that this effect would be even more critical for thinner grids, we favour using a NN trained on the particular observing set-up.

(iii) Observational uncertainties: when dealing with observational data sets, it is reasonable to assume the inherent uncertainties will limit our outflow detection ability, as shallow features are bound to smear within the noise. This is indeed a limitation intrinsic to the observed spectra that the ML technique cannot fully overcome. We accounted for its effect by including the uncertainties into the simulated spectral flux through the SNR parameter. Nevertheless, ATM remains agnostic to the observed uncertainties when classifying a particular spectrum. We attempt to evaluate this effect by generating a set of 100 spectra using a Monte Carlo approach, employing the observed spectrum as a seed and the uncertainties on each individual flux value as the standard deviation (effectively constructing a normal distribution). The resulting spectra were then classified using the trained ML algorithm. The classification rate for each type was reported, effectively providing an uncertainty in the classification due to the observation noise. We applied this method to some of the spectra where the SNR might have contributed to a more uncertain classification, and found evenly distributed probabilities among a few classes, as expected. When applied to the remaining spectra, the effect was negligible, and the results were consistent with the original classification. Nevertheless, given that the NN was already trained with noisy simulated spectra and that the observed spectra have the noise information indirectly engraved in their flux array, we fear double counting of the noise effect might be at play on the final classification, leading to a too conservative approach. Additionally, the NN itself provides us with a classification probability value (determined by the final softmax layer). Properly accounting for this effect requires the implementation of a Bayesian approach. Nevertheless, given that the classification results were not significantly impacted (except for a handful of low SNR spectra), we decided not to pursue this approach further and limit this work to the default analysis presented in the previous sections.

5 DISCUSSION

The pilot study presented in this paper aims to explore the capabilities of ML techniques in identifying the elusive outflow features in the optical spectra of LMXBs. The trained ATM algorithm produced overall consistent results with those derived via traditional methods when applied to real spectra, covering different line profiles and obtained during distinct outburst events of different LMXBs. In addition, this promising ML approach possesses a number of advantages compared with the traditional excess diagnostic diagram technique. The latter assumes Gaussian-like wings for the accretion disc line profile. Furthermore, it is dependent on the selection of the core mask, which,

due to the wide variety of profiles, ultimately relies on the observer's subjective decision and is bound to differ between systems. Finally, because of this very mask, information about the core of the line is neglected therefore rendering it extremely difficult to detect low-velocity outflow features. On the other hand, ATM can be applied to virtually any LMXB event, given that the observational set-up is properly accounted for during the construction of the simulated data base. It inspects the full emission line profile, and it accounts for known contaminants in order to minimize false detections. Extension to any given instrument and set-up configuration is straightforward, as only the spectral resolution and the dispersion of the instrument are required, which enables universal application of the script. ATM has been designed in a modular way, such that it is easy to explore different disc profiles than those presented here. That includes departures from the Keplerian distribution of velocities, such as advection-dominated accretion flows (ADAFs; see e.g. Narayan & Yi 1995), but also a more physical modelling of hotspot profiles in the disc. On the other hand, the code can be easily tweaked to include different outflow profiles (e.g. based on physical models; see Koljonen et al. 2023), allowing us to simulate more complex features, such as boxy profiles (e.g. Sánchez-Sierras & Muñoz-Darias 2020) or even inflows (e.g. Cúneo et al. 2020). Finally, ATM can be easily extended to any line in other wavelength ranges (such as the near-infrared or the ultraviolet), as long as any contaminants are properly accounted for. Application to other accreting objects exhibiting outflows, such as cataclysmic variables or even quasi-stellar objects, is also feasible, though a revision of the theoretical models would be required to accommodate each case.

At this point, it is worth highlighting a number of limitations of the ML approach that this pilot study revealed. First, a small number of spectra were miss-classified as belonging to the *absorbed* class. Exploring these particular spectra revealed the origin of this effect in a combination of normalization issues and contamination by nearby lines, which produced signatures apparently similar to broad absorptions. The most extreme example is illustrated by our attempts to classify the He II 4686 emission line, which, due to the presence of the nearby Bowen blend, resulted in a biased classification. Additionally, it is worth remarking that the classification of a particular spectrum is ultimately conditioned by the data base employed to train the algorithm. As such, any deviations not accounted for in the profile simulation might end up generating an incorrect classification of the observed spectra. On this regard, domain adaptation techniques have proven to be able to mitigate this effect (Ćiprijanović et al. 2021, 2023; Huertas-Company et al. 2023). Finally, detection of outliers to flag observed spectra too distinct from the simulated data sets would require a complementary approach using different architectures better suited for the task (e.g. autoencoders).

The pilot study presented in this paper has proven the success and capabilities of ML techniques when applied to our particular case. Our intent is to make ATM open-source to be exploited by the community. Future plans for ATM before its public release include the development of a user-friendly interface, further development to apply it to any instrumental set-up configuration (e.g. creating a grid of classifiers), and the addition of a utility for outlier detection (to highlight profiles not reproduced by the simulated data base).

6 CONCLUSIONS

We present a pilot study aiming to automatically identify outflow features in LMXB spectra using a new method based on supervised ML techniques. In order to train and test the performance of the ML algorithm, we build a simulated data base containing a million spectra

with line profiles based on theoretical disc models. In a fraction of these, we also injected different types of outflow features in a fraction of them. After application of three different architectures to four separate line profiles (H α , He I 5876, H β , and He II 4686), we conclude that the best performance is attained with the ResNet classifier set-up (which we name as ATM), showing an accuracy and recall over 93 per cent for all the relevant profiles. Comparison of the outflow detection and classification provided by this new technique with that attained via the traditional excesses diagnostic diagram showed overall consistent results, but with a number of associated advantages. First, we remove the dependence on the subjective determination of a mask for the analysis, as well as enable the detection of low-velocity outflow features. Secondly, the underlying accretion disc model is now physically motivated, and the algorithm has been designed to ease future implementation of more complex and precise models both for the accretion disc and outflow features. We also discuss the current limitations of this technique, such as its biased performance on heavily contaminated lines or its intrinsic dependence on the theoretical models themselves.

The notable success of ATM at classifying optical spectra for the two analysed LMXB outbursts in this pilot study (V404 Cygni and MAXI J1820) paves the way for its application to larger and more heterogeneous data sets. This tool immediately enables automatic detection and classification of outflows in future outburst events, as well as it sets the foundations for studies of the whole known population.

ACKNOWLEDGEMENTS

DMS acknowledges support from the Consejería de Economía, Conocimiento y Empleo del Gobierno de Canarias, and the European Regional Development Fund (ERDF) under grant with reference ProID2021010132 (ACCISI/FEDER, UE); as well as from the Spanish Ministry of Science and Innovation via an Europa Excelencia grant (EUR2021-122010). TMD acknowledges support via Ramón y Cajal Fellowship RYC-2015-18148. This work has been supported in part by the Spanish Ministry of Science and Technology under grants AYA2017-83216-P, PID2020-120323GB-I00, and EUR2021-122010. We thank Tom Marsh for the use of MOLLY software. We are thankful to the GTC staff for their prompt and efficient response at triggering the time-of-opportunity program at the source of the spectroscopy presented in this work. Based on data available at the GTC Archive at CAB (CSIC-INTA). The GTC Archive is part of the Spanish Virtual Observatory project funded by MCIN/AEI/10.13039/501100011033 through grant PID2020-112949GB-I00.

DATA AVAILABILITY

All data employed in this study are available at the GTC archive <https://gtc.sdc.cab.inta-csic.es/gtc/>. Data reduction is discussed in the original papers dedicated to the relevant sources. The simulated data base employed for the training and testing of the ATM algorithm will be shared upon reasonable request to the corresponding author.

REFERENCES

- Abadi M. et al., 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, <https://www.tensorflow.org/>
- Atri P. et al., 2019, *MNRAS*, 489, 3116
- Ball N. M., Brunner R. J., 2010, *Int. J. Mod. Phys. D*, 19, 1049

- Barthelmy S. D., D’Ai A., D’Avanzo P., Krimm H. A., Lien A. Y., Marshall F. E., Maselli A., Siegel M. H., 2015, *GCN Circ.*, 17929, 1
- Bianchini A., della Valle M., Masetti N., Margoni R., 1997, *A&A*, 321, 477
- Brink H., Richards J. W., Poznanski D., Bloom J. S., Rice J., Negahban S., Wainwright M., 2013, *MNRAS*, 435, 1047
- Callanan P. J. et al., 1995, *ApJ*, 441, 786
- Casares J. et al., 2022, *MNRAS*, 516, 2023
- Casares J., Charles P. A., Naylor T., 1992, *Nature*, 355, 614
- Casares J., Marsh T. R., Charles P. A., Martin A. C., Martin E. L., Harlaftis E. T., Pavlenko E. P., Wagner R. M., 1995, *MNRAS*, 274, 565
- Casares J., Muñoz-Darias T., Mata Sánchez D., Charles P. A., Torres M. A. P., Armas Padilla M., Fender R. P., García-Rojas J., 2019, *MNRAS*, 488, 1356
- Charles P., Matthews J. H., Buckley D. A. H., Gandhi P., Kotze E., Paice J., 2019, *MNRAS*, 489, L47
- Chollet F. et al., 2015, *Keras*, <https://github.com/fchollet/keras>
- Čiprijanović A. et al., 2021, *MNRAS*, 506, 677
- Čiprijanović A., Lewis A., Pedro K., Madireddy S., Nord B., Perdue G. N., Wild S. M., 2023, *Mach. Learn.: Sci. Tech.*, 4, 025013
- Corral-Santana J. M., Casares J., Muñoz-Darias T., Bauer F. E., Martínez-Pais I. G., Russell D. M., 2016, *A&A*, 587, A61
- Corral-Santana J. M., Casares J., Muñoz-Darias T., Rodríguez-Gil P., Shahbaz T., Torres M. A. P., Zurita C., Tyndall A. A., 2013, *Science*, 339, 1048
- Cúneo V. A. et al., 2020, *MNRAS*, 498, 25
- de Beurs Z. L., Islam N., Gopalan G., Vrtilik S. D., 2022, *ApJ*, 933, 116
- della Valle M., Benetti S., Cappellaro E., Wheeler C., 1997, *A&A*, 318, 179
- della Valle M., Jarvis B. J., West R. M., 1991, *Nature*, 353, 50
- Díaz Trigo M., Boirin L., 2016, *Astron. Nachr.*, 337, 368
- Dubus G., Kim R. S. J., Menou K., Szkody P., Bowen D. V., 2001, *ApJ*, 553, 307
- Fender R. P., Belloni T. M., Gallo E., 2004, *MNRAS*, 355, 1105
- Fender R., Muñoz-Darias T., 2016, *Astrophysical Black Holes, Lecture Notes in Physics, Volume 905*, Springer International Publishing, Switzerland, p. 65
- Giacconi R., Gursky H., Paolini F. R., Rossi B. B., 1962, *Phys. Rev. Lett.*, 9, 439
- Horne K., Marsh T. R., 1986, *MNRAS*, 218, 761
- Huertas-Company M. et al., 2023, preprint ([arXiv:2305.02478](https://arxiv.org/abs/2305.02478))
- Huertas-Company M., Lanusse F., 2023, *PASA*, 40, e001
- Iijima T., Esenoglu H. H., 2003, *A&A*, 404, 997
- Ioffe S., Szegedy C., 2015, in Bach F., Blei D. eds, *Proc. Mach. Learn. Res. Vol. 37, Proceedings of the 32nd International Conference on Machine Learning*. PMLR, Lille, France, p. 448, <https://proceedings.mlr.press/v37/loffe15.html>
- Ismail Fawaz H., Forestier G., Weber J., Idoumghar L., Muller P.-A., 2018, *Data Min. Knowl. Discov.*, 33, 3116
- Jiménez-Ibarra F., Muñoz-Darias T., Casares J., Armas Padilla M., Corral-Santana J. M., 2019, *MNRAS*, 489, 3420
- Killestein T. L. et al., 2021, *MNRAS*, 503, 4838
- Kimura M. et al., 2016, *Nature*, 529, 54
- King A. L., Miller J. M., Raymond J., Reynolds M. T., Morningstar W., 2015, *ApJ*, 813, L37
- Koljonen K. I. I., Long K. S., Matthews J. H., Knigge C., 2023, *MNRAS*, 521, 4190
- Krizhevsky A., Sutskever I., Hinton G. E., 2012, in Pereira F., Burges C., Bottou L., Weinberger K. eds, *Advances in Neural Information Processing Systems*, Vol. 25, Curran Associates, Inc., <https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf>
- Liu Q. Z., van Paradijs J., van den Heuvel E. P. J., 2007, *A&A*, 469, 807
- Masetti N., Bianchini A., della Valle M., 1997, *A&A*, 317, 769
- Mata Sánchez D. et al., 2018, *MNRAS*, 481, 2646
- Mata Sánchez D. et al., 2022, *ApJ*, 926, L10
- Mata Sánchez D., Muñoz-Darias T., Casares J., Corral-Santana J. M., Shahbaz T., 2015, *MNRAS*, 454, 2199
- Mata Sánchez D., Rau A., Álvarez Hernández A., van Grunsven T. F. J., Torres M. A. P., Jonker P. G., 2021, *MNRAS*, 506, 581
- McClintock J. E., Canizares C. R., Tarter C. B., 1975, *ApJ*, 198, 641
- Mikolov T., Sutskever I., Chen K., Corrado G. S., Dean J., 2013, in Burges C., Bottou L., Welling M., Ghahramani Z., Weinberger K. eds, *Vol. 26, Advances in Neural Information Processing Systems (NIPS 2013)*. Curran Associates, Inc., <https://proceedings.neurips.cc/paper/2013/file/9aa42b31882ec039965f3c4923ce901b-Paper.pdf>
- Motta S. E., Kajava J. J. E., Sánchez-Fernández C., Giustini M., Kuulkers E., 2017, *MNRAS*, 468, 981
- Muñoz-Darias T. et al., 2016, *Nature*, 534, 75
- Muñoz-Darias T. et al., 2019, *ApJ*, 879, L4
- Muñoz-Darias T., Ponti G., 2022, *A&A*, 664, A104
- Narayan R., Yi I., 1995, *ApJ*, 452, 710
- Neilsen J., Lee J. C., 2009, *Nature*, 458, 481
- Orosz J. A., Bailyn C. D., 1995, *ApJ*, 446, L59
- Orosz J. A., Bailyn C. D., Remillard R. A., McClintock J. E., Foltz C. B., 1994, *ApJ*, 436, 848
- Orwat-Kapola J. K., Bird A. J., Hill A. B., Altamirano D., Huppenkothen D., 2022, *MNRAS*, 509, 1269
- Panizo-Espinar G. et al., 2022, *A&A*, 664, A100
- Panizo-Espinar G., Muñoz-Darias T., Armas Padilla M., Jiménez-Ibarra F., Casares J., Mata Sánchez D., 2021, *A&A*, 650, A135
- Pattnaik R., Sharma K., Alabarta K., Altamirano D., Chakraborty M., Kembhavi A., Méndez M., Orwat-Kapola J. K., 2021, *MNRAS*, 501, 3457
- Ponti G., Bianchi S., Muñoz-Darias T., De K., Fender R., Merloni A., 2016, *Astron. Nachr.*, 337, 512
- Ponti G., Fender R. P., Begelman M. C., Dunn R. J. H., Neilsen J., Coriat M., 2012, *MNRAS*, 422, L11
- Rahoui F., Coriat M., Lee J. C., 2014, *MNRAS*, 442, 1610
- Ricketts B. J., Steiner J. F., Garraffo C., Remillard R. A., Huppenkothen D., 2023, *MNRAS*, 523, 1946
- Sainath T. N., Mohamed A.-r., Kingsbury B., Ramabhadran B., 2013, in 2013 IEEE International Conference on Acoustics, Speech and Signal Processing. Institute of Electrical and Electronic Engineers, p. 8614
- Sánchez-Sierras J., Muñoz-Darias T., 2020, *A&A*, 640, L3
- Shrader C. R., Wagner R. M., Charles P. A., Harlaftis E. T., Naylor T., 1997, *ApJ*, 487, 858
- Smith N., Hartigan P., 2006, *ApJ*, 638, 1045
- Smith N., Mauerhan J. C., Prieto J. L., 2014, *MNRAS*, 438, 1191
- Soria R., Wu K., Hunstead R. W., 2000, *ApJ*, 539, 445
- Srivastava N., Hinton G., Krizhevsky A., Sutskever I., Salakhutdinov R., 2014, *JMLR*, 15, 1929
- Steeghs D., Casares J., 2002, *ApJ*, 568, 273
- Szegedy C. et al., 2015, in 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Institute of Electrical and Electronic Engineers, p. 1
- Tetarenko B. E., Sivakoff G. R., Heinke C. O., Gladstone J. C., 2016, *ApJS*, 222, 15
- Thöne C. C. et al., 2017, *A&A*, 599, A129
- Torres M. A. P., Casares J., Jiménez-Ibarra F., Álvarez-Hernández A., Muñoz-Darias T., Armas Padilla M., Jonker P. G., Heida M., 2020, *ApJ*, 893, L37
- Torres M. A. P., Casares J., Jiménez-Ibarra F., Muñoz-Darias T., Armas Padilla M., Jonker P. G., Heida M., 2019, *ApJ*, 882, L21
- Torres M. A. P., Jonker P. G., Miller-Jones J. C. A., Steeghs D., Repetto S., Wu J., 2015, *MNRAS*, 450, 4292
- Tucker M. A. et al., 2018, *ApJ*, 867, L9
- Wang Z., Yan W., Oates T., 2017, in 2017 International Joint Conference on Neural Networks (IJCNN). Institute of Electrical and Electronic Engineers, p. 1578
- Wright D. E. et al., 2015, *MNRAS*, 449, 451

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.