A&A, 691, A78 (2024)
https://doi.org/10.1051/0004-6361/202451043

Astronomy & Astrophysics

Characterisation of the stellar wind in Cyg X-1 via modelling of colour-colour diagrams

E. V. Lai[1,2,⋆], B. De Marco[3], Y. Cavecchi[3], I. El Mellah[4,5], M. Cinus[6], C. M. Diez[7], V. Grinberg[8], A. A. Zdziarski[2], P. Uttley[9], M. Bachetti[1], J. José[3], G. Sala[3,10], A. Różańska[2], and J. Wilms[11]

[1] INAF-Osservatorio Astronomico di Cagliari, Via della Scienza 5, I-09047 Selargius, CA, Italy
[2] Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences, Bartycka 18, PL-00-716 Warszawa, Poland
[3] Departament de Física, EEBE, Universitat Politècnica de Catalunya, c/Eduard Maristany 16, 08019 Barcelona, Spain
[4] Departamento de Física, Universidad de Santiago de Chile, Av. Victor Jara 3659, Santiago, Chile
[5] Center for Interdisciplinary Research in Astrophysics and Space Exploration (CIRAS), USACH, Santiago, Chile
[6] Institute of Astronomy, Faculty of Physics, Astronomy and Informatics, Nicolaus Copernicus University, Grudziądzka 5, 87-100 Toruń, Poland
[7] European Space Agency (ESA), European Space Astronomy Centre (ESAC), Camino Bajo del Castillo s/n, 28692 Villanueva de la Cañada, Madrid, Spain
[8] European Space Agency (ESA), European Space Research and Technology Centre (ESTEC), Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands
[9] Anton Pannekoek Institute, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands
[10] Institut d'Estudis Espacials de Catalunya, c/Gran Capità 2-4, Ed. Nexus-201, 08034 Barcelona, Spain
[11] Dr. Karl Remeis-Observatory, University of Erlangen-Nuremberg, Sternwartstr. 7, 96049 Bamberg, Germany

Received 8 June 2024 / Accepted 12 August 2024

ABSTRACT

Context. Cygnus X-1 (Cyg X-1) is a high-mass X-ray binary where accretion onto the black hole (BH) is mediated by the stellar wind from the blue supergiant companion star HDE 226868. Due to its inclination, the system is a perfect laboratory to study the not yet well-understood stellar wind structure. In fact, depending on the position of the BH along the orbit, X-ray observations can probe different layers of the stellar wind. Deeper wind layers can be investigated at superior conjunction (i.e. null orbital phases).
Aims. We aim to characterise the stellar wind in the Cyg X-1/HDE 226868 system, analysing one passage at superior conjunction covered by *XMM-Newton* during the 'Cyg X-1 Hard state Observations of a Complete Binary Orbit in X-rays' (CHOCBOX) campaign.
Methods. To analyse the properties of the stellar wind, we computed colour-colour diagrams. Since X-ray absorption is energy-dependent, colour indices provide information on the parameters of the stellar wind, such as the column density, $N_{\rm H,w}$, and the covering factor, $f_{\rm c}$. We fitted colour-colour diagrams with models that include both a continuum and a stellar wind component. We used the kernel density estimation method to infer the unknown probability distribution of the data points in the colour-colour diagram, and selected the model corresponding to the highest likelihood. In order to study the temporal evolution of the wind around superior conjunction, we extracted and fitted time-resolved colour-colour diagrams.
Results. We found that the model that best describes the shape of the colour-colour diagram of Cyg X-1 at superior conjunction requires the wind to be partially ionised. The shape of the colour-colour diagram strongly varies during the analysed observation, due to concurrent changes of the mean $N_{\rm H,w}$ and the $f_{\rm c}$ of the wind. Our results suggest the existence of a linear scaling between the rapid variability amplitude of $N_{\rm H,w}$ (on timescales between 10 s and 11 ks) and its long-term variations (on timescales >11 ks). Using the inferred best-fit values, we estimated the stellar mass loss rate to be $\sim 7 \times 10^{-6}\,{\rm M}_{\odot}\,{\rm yr}^{-1}$ and the clumps to have a characteristic mass of $\sim 10^{17}$ g.

Key words. black hole physics – stars: winds, outflows – X-rays: binaries – X-rays: individuals: Cyg X-1

1. Introduction

High-mass X-ray binaries (HMXBs) are a class of binary systems in which a neutron star or a black hole (BH) is fed by the stellar wind from a massive ($\gtrsim 10\,M_{\odot}$) OB companion or by the decretion disc surrounding a quickly spinning Be donor star. The stellar wind is line-driven (Castor et al. 1975) and can be accelerated up to terminal velocities of $v_{\infty} \sim 2500\,{\rm km\,s^{-1}}$ (e.g. Martínez-Núñez et al. 2017). Internal shocks lead to the formation of overdense regions, or clumps, making the wind inhomogeneous and highly structured (e.g. Owocki et al. 1988; Feldmeier et al. 1997; Oskinova et al. 2012; Sundqvist & Owocki 2013). While several stellar wind models have been developed (e.g. Oskinova et al. 2012; Sundqvist & Owocki 2013; Sundqvist et al. 2018; El Mellah et al. 2018, 2020), the physical properties of the clumps, such as their shape, density, and ionisation structure, are yet to be precisely constrained. Investigating the structure of the stellar winds produced by massive stars is crucial for various purposes, including constraining models of stellar and galactic evolution, and understanding the accretion mechanism in wind-fed HMXBs.

As a matter of fact, stellar winds represent an important mechanism of stellar mass loss, with mass-loss rates up

⋆ Corresponding author; eleonora.lai@inaf.it

to $10^{-5}\,M_\odot\,yr^{-1}$ (Puls et al. 2008; Martínez-Núñez et al. 2017). This has an influence on the properties of the donor star (such as mass, luminosity, spin, and chemical composition) as well as its evolutionary timescales. Moreover, in binary systems interacting via mass transfer mediated by a stellar wind like most HMXBs, only a fraction of the expelled gas is accreted, while the rest escapes the system, causing a removal of angular momentum and a change in the orbital parameters. Whether this process leads to a shrinking or a widening of the orbit (e.g. Hoyle & Lyttleton 1939; Bondi & Hoyle 1944; Paczynski 1976; El Mellah et al. 2020) depends on key parameters, such as the wind terminal velocity, the initial orbital separation, and the accretion efficiency (e.g. Saladino et al. 2018). The evolutionary pathways of such interactions can have significant impact on the rate of merging events and guide us in the quest for merger progenitors (e.g. Bulik et al. 2011; Belczynski et al. 2011; Neijssel et al. 2021).

The well-known HMXB system Cygnus X-1 (Cyg X-1 hereafter) is a good target to constrain the stellar wind structure and infer information on the physical properties of wind clumps. The system hosts a BH with $M_{\rm BH} = 21.2 \pm 2.2\,M_\odot$ (Miller-Jones et al. 2021) in a ∼5.6 d quasi-circular orbit (Gies et al. 2003). The O9.7 Iab supergiant companion star HDE 226868, with $M_* = 40.6^{+7.7}_{-7.1}\,M_\odot$ and $R_* = 22.3 \pm 1.8\,R_\odot$ (Miller-Jones et al. 2021), launches fast stellar winds ($v_\infty = 2100\,km\,s^{-1}$, Herrero et al. 1995). The orbital inclination of the system ($i \sim 27°$, Miller-Jones et al. 2021) causes our line of sight (LOS) to intercept the clump-forming region, likely located close to the companion star's surface (e.g. El Mellah et al. 2020).

In this system, the presence of the wind significantly modifies the spectral (e.g. Nowak et al. 2011) and timing (Lai et al. 2022) properties of the X-ray source. Indeed, X-ray photons undergo different levels of absorption depending on the amount and the ionisation state of the wind material intercepting the LOS. Wind absorption mostly affects the softest energy bands ($E \lesssim 1$ keV), although its influence can be observed up to higher energies of $E \sim 8-10$ keV (e.g. Hirsch et al. 2019; Grinberg et al. 2015). The strength of the absorption events strongly varies as a function of the orbital phase (e.g. Bałucińska-Church et al. 2000; Poutanen et al. 2008; Hirsch et al. 2019; Grinberg et al. 2020; Lai et al. 2022). In particular, the X-ray light curves show more intense and frequent absorption dips when the X-ray source is at superior conjunction ($\phi_{\rm orb} = 0$), that is when the BH is the farthest and the companion star the closest from the observer along the orbit (see Li & Clark 1974 and references therein). In Lai et al. (2022), we showed that the absorption dips significantly contribute to the X-ray variability of Cyg X-1 by increasing its observed fractional variability on timescales longer than ∼1 s, and suggested that the motion of clumps of the size of $\sim 10^{-4}\,R_*$ may be responsible for the enhancement of X-ray variability on these timescales.

Radiative hydrodynamical simulations show that the size of the clumps should increase as they move away from the star (Sundqvist et al. 2018), producing dips of increasingly long duration. Therefore, to fully investigate the properties of the clumps, we need to access a broad range of timescales, down to the shortest timescales (∼1 min or shorter) characterising the smallest clumps formed close to the base of the wind. However, the limitations of current instruments preclude wind-induced spectral changes occurring on very short timescales from being investigated via canonical spectral analysis.

An alternative, powerful approach is to use colour-colour diagrams (e.g. Nowak et al. 2011; Grinberg et al. 2020). Detailed modelling of the colour-colour diagram of wind-fed systems can give constraints on the physical properties of the wind. Earlier studies made use of empirical functions or simple (neutral) stellar wind models to describe the characteristic shape of colour-colour tracks (e.g. Hanke et al. 2008; Nowak et al. 2011; Hirsch et al. 2019). Recently, Grinberg et al. (2020) explored more complex wind models, concluding that a proper description of the observed colour-colour tracks requires accounting for different levels of ionisation in the wind. In this paper, building on the work of Grinberg et al. (2020), we model the colour-colour diagrams of Cyg X-1 in order to constrain the parameters of the wind as a function of the orbital phase. This information can be used to infer physical properties such as the stellar mass-loss rate and the mass of the clumps (e.g. El Mellah et al. 2020).

In this paper, we first describe the data reduction procedure (Sect. 2), then we test a number of stellar wind models (Sect. 3) to determine the one that best describes the colour-colour diagram of Cyg X-1 in its hard state (the most affected by the wind, Nowak et al. 2011). In Sect. 4, we investigate the temporal evolution of the column density and covering factor of the wind by fitting time-resolved colour-colour diagrams. We discuss our results in Sect. 5 and present our conclusions in Sect. 6.

2. Observation and data reduction

We focussed our analysis on a single *XMM-Newton* observation of Cyg X-1 (obsID 0745250201, hereafter 201) in its hard state, which is part of the ∼1.5 orbital period-long monitoring from the multi wavelength campaign "Cyg X-1 Hard state Observations of a Complete Binary Orbit in X-rays" (CHOCBOX). The full CHOCBOX coverage is shown in Fig. 1. The monitoring consists of four pointings for a total duration of ∼7 days, and catches two consecutive passages of the X-ray source at superior conjunction.

In Lai et al. (2022), we investigated the changes driven by the stellar wind in the X-ray spectral-timing properties of the source throughout the entire monitoring. Among all the CHOCBOX observations, observation 201 turned out to be the most affected by wind absorption. It also showed the most prominent effects on the X-ray variability on short timescales (in the ∼0.1−10 s range) as compared to the other observations. Observation 201 covers orbital phases $\phi_{\rm orb} = 0.82-0.06$, corresponding to the first passage at superior conjunction (Lai et al. 2022). We focussed on observation 201 for the study of the stellar wind, although in Sect. 4 we also used observation ID 0745250501 (hereafter 501) to identify the region of the colour-colour diagram less affected by wind-absorption. Indeed, observation 501 covers the orbital phases $\phi_{\rm orb} = 0.17-0.46$. Being the closest to inferior conjunction with $\phi_{\rm orb} = 0.5$ it is also the least absorbed. For both observations we used EPIC-pn data (Strüder et al. 2001) in timing observing mode. We performed the data reduction using the *XMM-Newton* `Science Analysis Software` (`SAS`, version 20.0.0) and calibration files (CCF) as of 2022 March.

Diez et al. (2023) showed that the extraction of EPIC-pn spectra in timing mode using the default Rate Dependent PHA (RDPHA) correction led to energies higher than expected for line features, and suggest a non-standard calibration to mitigate the problem. We followed the same approach proposed by Diez et al. (2023). We ran `epproc` on the Observation Data Files (ODFs), turning off the RDPHA correction (`withrdpha='N'`) and applying the Rate Dependent CTI (RDCTI) correction using `epfast` (`runepfast='Y'`). Calibrated event files were screened for the presence of strong background flaring events using light curves with 1s time resolution in the 10−15 keV energy band. No significant flares were observed. Using the

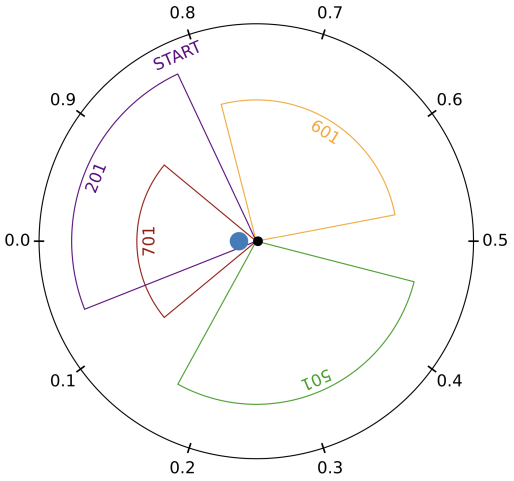

Fig. 1. Orbital phase coverage of the *XMM-Newton* CHOCBOX monitoring of Cyg X-1 in the hard state. The four arc labels correspond to the ObsIDs of each pointing (see Table 1 in Lai et al. 2022 for the full ObsID). The orbital phase, $\phi_{\rm orb} = 0$, indicates the passage at superior conjunction. Observation 201 (in purple) is the focus of the analyses presented in this paper. The starting orbital phase of the monitoring (corresponding to $\phi_{\rm orb} = 0.82$) is also indicated. HDE 226868 and Cyg X-1 are only schematically represented, no Roche lobe was graphically considered in the scheme.

SAS task `epatplot`, we found that a small fraction of pile-up is present in the data. We corrected for it by extracting source counts from a region that excludes the central pixels (RAWX = [36:39]); that is RAWX = [30:35; 40:46]. Background counts cannot be reliably estimated from timing mode observations of bright sources, as the entire CCD is illuminated by source photons. Since the PSF is energy-dependent, extracting the background from the outer columns is not recommended as this procedure modifies the intrinsic spectrum (Ng et al. 2010). Since the source is very bright, with $\gtrsim$220 counts s^{-1} in the 0.5−10 keV energy band, the background contribution is expected to be negligible. Therefore, we decided to not account for the background in our analysis. We made use of the SAS tasks `arfgen` and `rmfgen` to extract the Ancillary Response Files (ARF) and Redistribution Matrix Files (RMF). We produced two ARF files: one for the full region (RAWX = [30:46]) and one for the central excluded region (RAWX = [36:39]). Using the command `addarf`, we subtracted the latter from the former, generating the final ARF file (e.g. Wilkinson & Uttley 2009; Lai et al. 2022). The spectrum was rebinned to ensure a minimum of 20 counts per bin in order to apply chi-square statistics.

Spectral fits were done using `XSPEC` v12.10.1 (Arnaud 1996). We used the Interactive Spectral Interpretation System (`ISIS`) version 1.6.2 and a custom code, written in Python 3.10, for the calculation and the modelling of the colour-colour diagrams. Time series were extracted using `stingray` v1.1.2 (Huppenkothen et al. 2019a,b; Bachetti et al. 2023).

3. Models for the colour-colour diagram of Cyg X-1

The passage of absorbing clumps across the LOS manifests as dips in the X-ray light curve of length ranging between several seconds to hours (e.g. Hirsch et al. 2019). To explore the presence and duration of dips in observation 201, we show its light curve in Fig. 2. In the upper and middle panels, we report the soft (0.5−1.5 keV) and the hard (3−10 keV) band light curves, respectively (for the light curves of the entire modelling, see Lai et al. 2022[1]).

The presence of dips in the soft band light curve, which are approximately three times less intense in the hard band light curve, produces sharp increases in the hardness ratios (as is shown in the bottom panel of Fig. 2, where we choose to compare two non-contiguous energy bands to maximise the intensity of the dips). The single dips are quite short ($\lesssim$1 ks). Therefore, given the swiftness of such events, their energy spectrum is difficult to study via time-resolved spectroscopy. An alternative approach is to investigate the broad band spectral variability caused by wind clumps using colour-colour diagrams (e.g. Nowak et al. 2011). These diagrams map the evolution of the ratio of the count rates in different energy bands. Specifically, a soft colour, defined as the ratio between count rates in a soft and an intermediate energy band, is plotted against a hard colour, defined as the ratio between count rates in the same intermediate band and a hard energy band. High values of both soft and hard colours characterise the least absorbed phases, and thus map the upper region of the diagram. As absorption along the LOS increases, the hard and soft colours change. In particular, for a partially covering absorber the resulting colour-colour tracks describe a 'pointy' or 'nose-like' shape as the source becomes more absorbed (i.e. in the lower region of the colour-colour diagram, see e.g. Hirsch et al. 2019). Following previous studies (Hanke et al. 2008; Nowak et al. 2011; Hirsch et al. 2019; Grinberg et al. 2020), we constructed the colour-colour diagram of observation 201 using light curve bins of 10 s, allowing us to detect also the spectral variability that could be caused by small clumps. The size of the clumps is directly related to the fly-by time across the LOS (El Mellah et al. 2020). With this time resolution, the minimum size of the clumps that can be tested is $\sim 2 \times 10^{-4}\,R_*$, assuming a wind terminal velocity of $v_\infty = 2400\,{\rm km\,s^{-1}}$ (Grinberg et al. 2015), where R_* is the radius of the companion star. The ratios were computed between the energy bands 0.5−1.5/1.5−3 keV (soft colour), and 1.5−3/3−10 keV (hard colour), see Fig. 3.

We then tested a number of physical scenarios in order to find a suitable stellar wind model able to describe the observed colour-colour diagram tracks in Cyg X-1. As our baseline model we chose a primary X-ray continuum modified by a partially covering absorber (Fürst et al. 2014; Fornasini et al. 2017; Grinberg et al. 2020). The adopted model is of the form:

$$\mathtt{abs}_{\rm ism} \times \mathtt{continuum} \times (f_{\rm c} \times \mathtt{abs}_{\rm wind} + (1 - f_{\rm c})) \quad (1)$$

where $\mathtt{abs}_{\rm ism}$ is the absorption of the interstellar medium (ISM), `continuum` refers to the unabsorbed primary emission from the X-ray source and $\mathtt{abs}_{\rm wind}$ is the absorption component associated with the stellar wind. This component blocks a fraction ($f_{\rm c}$, or "covering fraction") of the X-ray source photons, while only a percentage $1 - f_{\rm c}$ reaches the observer modified only by the interstellar medium.

We made use of standard `XSPEC` models (Arnaud 1996) for the continuum and for the stellar wind to simulate spectra within `ISIS` (Houck & Denicola 2000; Houck 2002; Noble & Nowak 2008), using the response matrices of observation 201 (see Sect. 2). We let the parameters of interest of the model vary and, for each combination of parameters, we calculate the corresponding soft and hard colours. In order to speed up the process, for each set of simulations, only one parameter is left free to

[1] Note that the correct starting time of the observation is 57535.9 MJD, instead of the 57535.0 MJD date reported in the caption of Figure 1 of Lai et al. (2022).

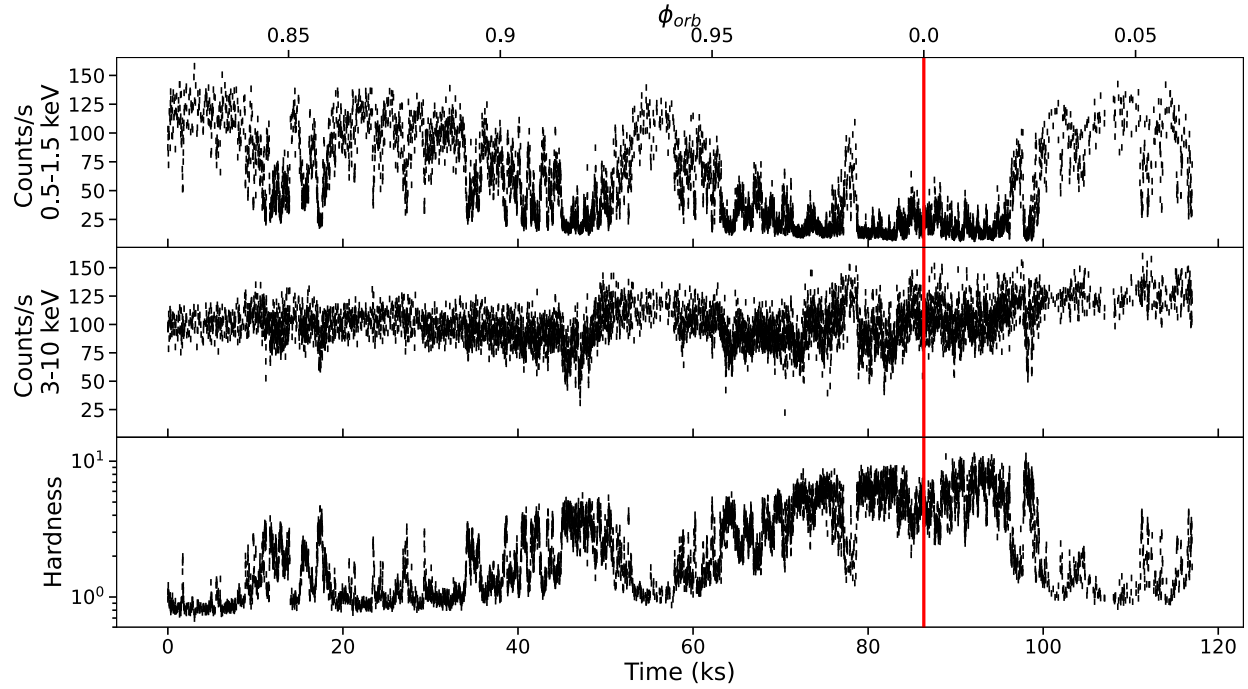

Fig. 2. *XMM–Newton* EPIC-pn light curves of the observation 201 of Cyg X-1 with time bins of 10 s. The upper and the middle panels show, respectively, the 0.5−1.5 keV and the 3−10 keV light curves. The bottom panel reports the hardness, i.e. the ratio between count rates in the 3−10 keV and 0.5−1.5 keV energy bands. The vertical red line indicates the passage at superior conjunction (i.e. $\phi_{\rm orb} = 0$).

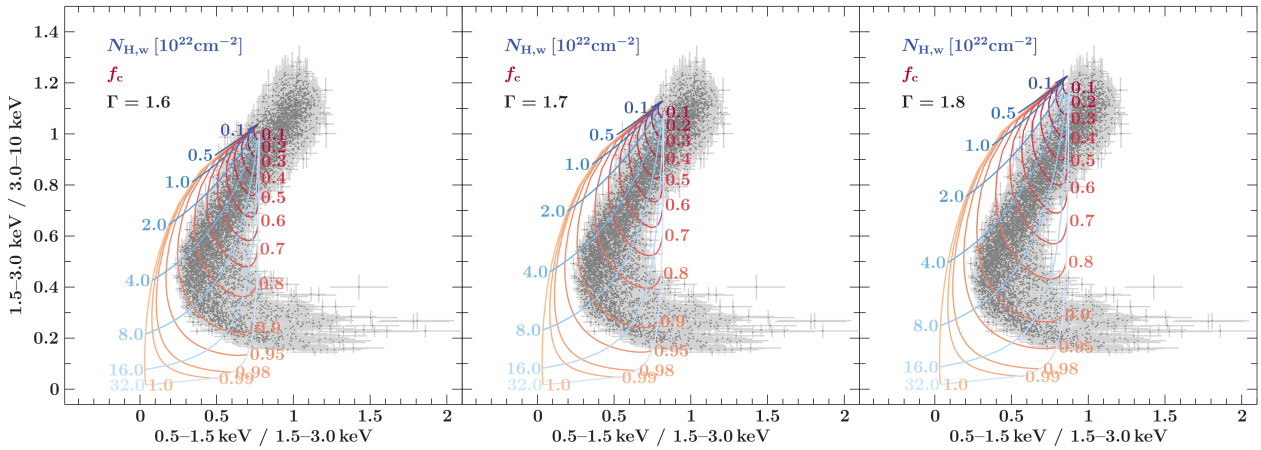

Fig. 3. Simulated tracks for a power law (with Γ fixed at 1.6, left panel, 1.7, middle panel, and 1.8, right panel) plus a neutral absorber, compared to the observed (grey points) colour-colour diagram tracks of observation 201 of Cyg X-1. The blue (red) tracks describe changes in the wind f_c ($N_{\rm H,w}$) for a fixed value of $N_{\rm H,w}$ (f_c).

vary, while the others are kept fixed. This procedure allows us to build simulated colour-colour tracks, which are then compared to the data. We tested different models as is described in the following sections. Throughout this work ISM absorption was modelled with `TBabs`[2] using the `wilm` abundances (Wilms et al. 2000) and the `vern` cross-sections (Verner et al. 1996). The column density of the ISM was kept fixed at the tabulated value of $\sim 0.7 \times 10^{22}$ cm^{-2} (HI4PI Collaboration 2016). We note that in this paper we are interested in testing for variability of the stellar wind component; therefore, in our simulations, we do not investigate intrinsic spectral variations of the X-ray source.

[2] https://pulsar.sternwarte.uni-erlangen.de/wilms/research/tbabs/

3.1. Model 1: Single power law plus neutral stellar wind

We first simulated theoretical colour-colour tracks assuming a power law for the underlying continuum and a partially covering neutral absorber. For this simple model, the shape of the colour-colour tracks mainly depends on three parameters: the power law photon index Γ of the continuum, the wind covering fraction f_c and the wind column density $N_{\rm H,w}$.

Fig. 3 shows the resulting simulated tracks overplotted on the colour-colour diagram of observation 201. For fixed values of Γ and $N_{\rm H,w}$, a variable f_c describes the curvature of the blue-cyan tracks shown in Fig. 3. Within each track the f_c parameter samples the range of values $f_c = 0.1-1$ (at steps of 0.1). The different tracks from deep blue to light cyan are obtained by increasing the value of $N_{\rm H,w}$ from a minimum value of $N_{\rm H,w} = 0.1 \times 10^{22}$

to a maximum value of $N_{\rm H,w} = 32 \times 10^{22}\,{\rm cm}^{-2}$ (for higher values of $N_{\rm H,w}$ hard colours start to increase again, producing an upward tail that is not observed in the data). On the other hand, for fixed values of Γ and $f_{\rm c}$, a variable $N_{\rm H,w}$, within the range $N_{\rm H,w} = 0.1{-}32 \times 10^{22}{\rm cm}^{-2}$, describes the curvature of the red-orange tracks shown in Fig. 3. The different tracks (from deep red to light orange) are obtained by increasing the value of $f_{\rm c}$ (from 0.1 to 1). We repeated the same simulations for $\Gamma = 1.6$, 1.7, and 1.8 (panels from left to right in Fig. 3), which match typical values for the hard state (e.g. Joinet et al. 2008; Motta et al. 2009; Gilfanov 2010; Basak et al. 2017; Zhou et al. 2022).

We confirm the results of previous studies (e.g. Hirsch et al. 2019; Grinberg et al. 2020), which showed that this simple model fails to describe the observed colour-colour diagram. In particular, a variable $N_{\rm H,w}$, while approximating the curvature of the observed track, does not properly reproduce its pointy shape at low hard colours (giving the track a nose-like shape). We observe that a varying Γ shifts the tracks vertically. Therefore, in line with previous studies (e.g. Grinberg et al. 2020), we infer that the data require a more complex modelling of the stellar wind and/or the intrinsic X-ray continuum. Indeed, several X-ray spectroscopic studies highlight the need for a more complex primary continuum for Cyg X-1 (e.g. Basak et al. 2017; Tomsick et al. 2018) as well as a structured stellar wind made of differently ionised material (Hirsch et al. 2019).

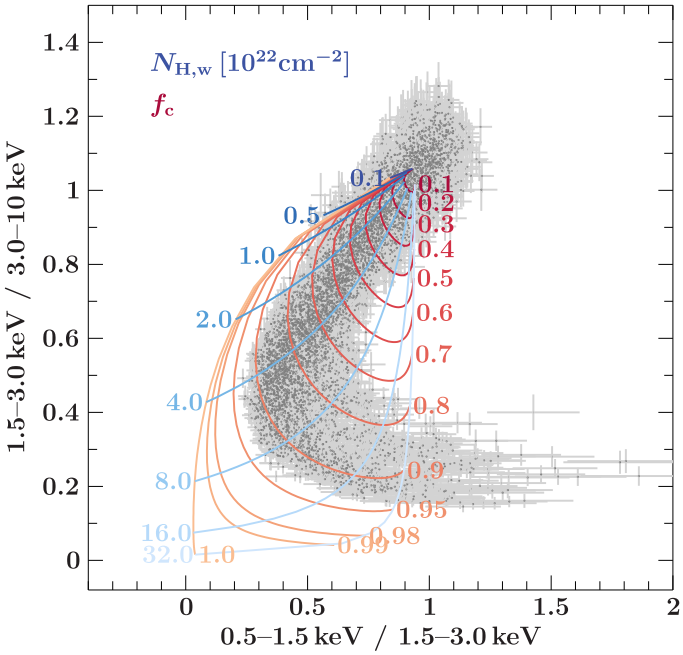

Fig. 4. Simulated tracks for the complex continuum absorbed by a neutral wind model, as compared to the colour-colour diagram of Cyg X-1.

3.2. Model 2: Structured continuum plus ionised stellar wind

We then changed our assumptions on the underlying primary continuum and tested a more complex model. To this aim, we assumed the model presented in Lai et al. (2022), which includes a soft and a hard Comptonisation component, as well as direct thermal and reflected emission from the disc. The employed XSPEC model is: `TBabs` × `[diskbb + nthComp + relxillCp]`. In order to properly assess the time-averaged parameters of the continuum we fit a subset of good time intervals (GTIs), selected so as to be the least affected by the intervening stellar wind. To this aim, we adopted the same criterion defined in Lai et al. (2022); that is, we selected data with hard colour ≥0.95 and soft colour ≥0.7, so as to single out GTIs from the upper, least absorbed region of the colour-colour diagram. Details about the continuum model used in these simulations are reported in Appendix A.

Fig. 4 shows the colour-colour tracks computed for a neutral gas partially absorbing the adopted structured continuum. As opposed to simulations shown in Fig. 3, the new tracks all start right at the centre of the observed distribution of data points (upper part of the diagram), meaning that the new continuum model describes the unabsorbed time-averaged spectrum much better than a simple power law. Therefore, hereafter we use this more complex model to describe the continuum in our simulations.

We then considered the effects of an ionised wind on the simulated colour-colour tracks, substituting the $\texttt{abs}_{\rm wind}$ component in Eq. (1) with the `warmabs` (v2.31) analytic photoionisation model (Kallman et al. 2009). As in Grinberg et al. (2020), we used the standard population files delivered with `warmabs`, with densities of $10^{12}\,{\rm cm}^{-3}$, typical of stellar wind densities close to the compact object (Lomaeva et al. 2020). We first verified how different levels of ionisation modify the simulated tracks. To this aim, we considered the ionisation parameter $\xi = L_{\rm X}/nr^2$ (Tarter et al. 1969), where $L_{\rm X}$ is the ionising luminosity above 13.6 eV, n is the absorbing gas density, and r its distance from the ionising X-ray source. Values of $\log\xi$ between −4 and 2 were considered. For each fixed value of $\log\xi$, we let the $N_{\rm H,w}$ vary between 0.1 and $32 \times 10^{22}\,{\rm cm}^{-2}$, leading to the tracks shown in Fig. 5. We reproduced the same simulations for three different values of covering factor ($f_{\rm c} = 0.8$, 0.87, 0.95, Fig. 5 from left to right).

By increasing the ionisation parameter, the simulated tracks become more pointy, thus more closely resemble the observed tracks. The effect of increasing the covering factor is mostly to increase the range of encompassed hard and soft colours. In particular, for sufficiently high $N_{\rm H,w}$, a higher $f_{\rm c}$ produces more absorption at low-to-intermediate energies, causing the absorption dips to be more spectrally hard (lower values of hard and soft colours during the most absorbed stages). Among the simulated curves, we observe that the one better resembling the observed tracks corresponds to $f_{\rm c} = 0.87$, $\log\xi = 1.6$, and $N_{\rm H,w}$ varying between 0.1 and $32 \times 10^{22}\,{\rm cm}^{-2}$, as is highlighted in black in the middle panel of Fig. 5, suggesting mild ionisation for the partially covering wind.

3.3. Model 3: Stellar wind with variable ionisation

The simulations shown in Sect. 3.2 assumed a constant ionisation parameter throughout the different stages of the absorption dips. Nonetheless, Hirsch et al. (2019) reported the appearance of less ionised species in the most absorbed stages corresponding to the lower part of the colour-colour diagram. This suggests the presence of structured clumps, made of differently ionised material. Such results highlight the need to consider changes in the ionisation parameter (Grinberg et al. 2020). Therefore, we now account for variations in the stellar wind ionisation parameter as a function of $N_{\rm H,w}$.

Following Grinberg et al. (2020), we test two empirical functions to describe the dependence of the ionisation parameter on the column density of the absorbing material. The first function is defined as:

$$\log\xi = \log(A/[N_{\rm H,w}/10^{22}\,{\rm cm}^{-2}]), \quad (2)$$

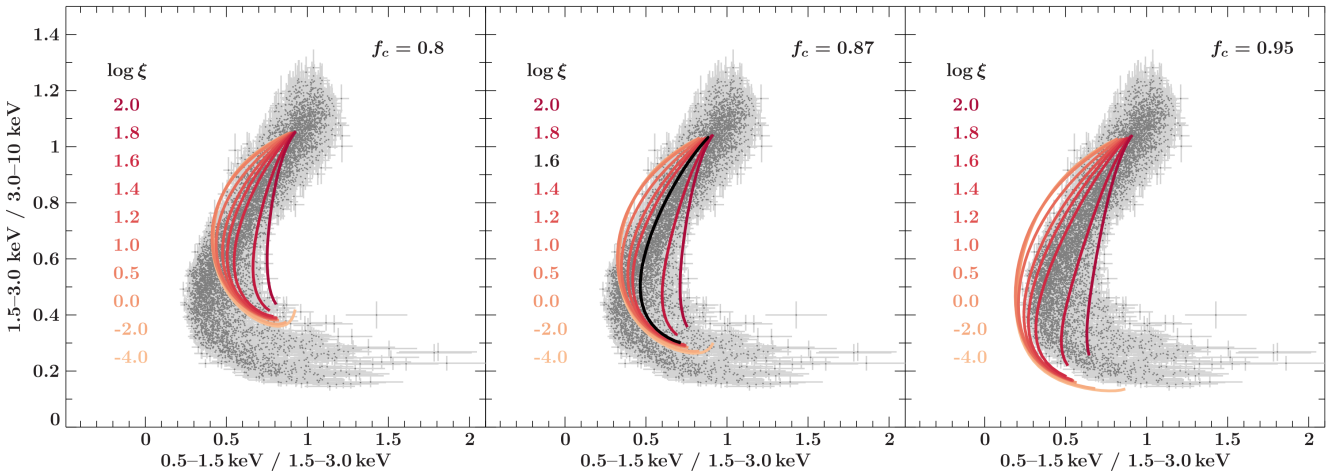

Fig. 5. Simulated tracks for a homogeneously ionised absorber (with $\log \xi$ values ranging between −4 and 2) compared to the colour-colour diagram of observation 201. The different panels show the simulated tracks computed for different values of f_c: 0.8 (left), 0.87 (middle), 0.95 (right). The black curve in the middle panel represents the one that best describes the data.

with $A > 0$. It is based on the definition of the ionisation parameter and assumes a linear decrease in ξ as a function of $N_{\rm H,w}$. The second function is defined as:

$$\log \xi = \log \frac{B + [N_{\rm H,w}/10^{22}\,{\rm cm}^{-2}]}{[N_{\rm H,w}/10^{22}\,{\rm cm}^{-2}]} + C, \quad (3)$$

with B, $C > 0$. It assumes that the ionisation parameter deviates from a linearly decreasing trend at high column densities, staying relatively high, as expected if an additional source of ionisation plays an important role in denser environments (e.g. Feldmeier et al. 1997). We arbitrarily chose the values of the constant parameters to be $A = 100$, $B = 10$, and $C = 1$. This choice of parameters is such that the two functions match at low column densities $N_{\rm H,w}/10^{22}\,{\rm cm}^{-2} << 10$, and start deviating significantly at higher densities, $N_{\rm H,w}/10^{22}\,{\rm cm}^{-2} \gtrsim 1$. We let $N_{\rm H,w}$ vary between 0.01 and $32 \times 10^{22}\,{\rm cm}^{-2}$. Sampling down to values of $N_{\rm H,w} = 0.01 \times 10^{22}\,{\rm cm}^{-2}$ allows us to test values of ionisation parameter as high as the highest value allowed by the `warmabs` model. Therefore, the first function spans the range of $\log \xi \sim 0.5$–4, while the second function spans the range of $\log \xi \sim 1.12$–4. The two empirical functions are shown in Fig. 6.

We built models comprising the complex continuum defined in Sect. 3.2 plus each of the two empirical functions describing variable ionisation. Since the data in the colour-colour diagrams are not normally distributed, a simple χ^2 minimisation method cannot be used to select the best-fit model. In order to find the best-fit model without any prior assumption on the distribution of the data we employed a non-parametric method.

We used the kernel density estimation (KDE) method to infer the probability density function (PDF) of the data in the colour-colour diagram (e.g. Hill 1985; Cavecchi & Patruno 2022). The KDE is a non-parametric estimator, and its functional form is obtained by combining as many building blocks – kernels – as the number of data points. In the KDE a single type of kernel K is used. The choice of the kernel is arbitrary, but for large datasets this choice does not have any significant effect on the final output. Therefore, we used a simple kernel with a Gaussian shape. The kernel K_i centred on the data point $\boldsymbol{x}_i$ is defined as

$$K_i(\boldsymbol{x}) \propto \exp\left(-\frac{(\boldsymbol{x} - \boldsymbol{x}_i)^2}{2h^2}\right) \quad (4)$$

where $\boldsymbol{x}$ is the point in the colour-colour diagram at which the kernel function is estimated. The width of the kernel, h, controls the smoothing. The normalised PDF is then computed as:

$$\mathrm{PDF}(\boldsymbol{x}) = \frac{1}{n}\sum_{i}^{n} K_i(\boldsymbol{x}), \quad (5)$$

which sums the contribution of the kernels centred on each data point, with n being the total number of data points. We implemented the method using the SciPy function `gaussian_kde`, with the optimal kernel width automatically determined by the function (default parameter `bw_method='scott'`[3]). Using the KDE method we estimated the PDF of the data. Then, for each simulated model we computed the combined likelihood as the product of the values of the probability density at each point of the model. The best-fit model is the one characterised by the highest combined likelihood.

In Fig. 6, we show the best-fit track obtained using the functions $\log \xi = \log(100/[N_{\rm H,w}/10^{22}\,{\rm cm}^{-2}])$ (middle panel) and $\log \xi = \log \frac{10+[N_{\rm H,w}/10^{22}\,{\rm cm}^{-2}]}{[N_{\rm H,w}/10^{22}\,{\rm cm}^{-2}]} + 1$ (right panel) to model the ionised absorber. The best-fit track is obtained in both cases by fitting for the covering factor f_c (between 0.7 and 0.95, at steps of 0.01) parameter while $N_{\rm H,w}$ spans the entire range 0.01–$32 \times 10^{22}\,{\rm cm}^{-2}$ for each fit value of the covering factor. In Fig. 6, we also show the tracks corresponding to $f_c = 0.7$ and $f_c = 0.95$ and note that outside these limiting values the simulated curves do not intersect the data in the lower part of the diagram.

During the least absorbed stages (i.e. at hard colours $\gtrsim$0.8), the simulated tracks are not very sensitive to the covering factor, and they all resemble fairly well the observed shape of the track in this part of the diagram. However, at lower values of hard colours, corresponding to the most absorbed phases, the dependence on the covering factor becomes significant. The two variable ionisation models considered here describe the data well, with a mild 5% difference between their inferred probabilities, which does not allow us to prefer either one of the two models. Notably, both functions resemble the pointy shape of the data much better than the homogeneously ionised absorber considered in Sect. 3.2, giving more support to the hypothesis of a structured absorber. Finally, we observe that the data all lay within the simulated tracks for $f_c = 0.7$ and $f_c = 0.95$, implying that the scatter of data points in this part of the diagram might be

[3] https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.gaussian_kde.html