

Time-resolved X-ray spectra of Proxima Centauri as seen by XMM-Newton

A. Damonte^{1,2,3,*}, I. Pillitteri³, A. Maggio³, A. García Muñoz¹, and G. Micela³

¹ Université Paris Cité, Université Paris-Saclay, CEA, CNRS, AIM, 91191 Gif-sur-Yvette, France

² Dip. di Fisica e Chimica, Università degli Studi di Palermo, Piazza del Parlamento 1, 90134 Palermo, Italy

³ INAF – Osservatorio Astronomico di Palermo, Piazza del Parlamento 1, 90134 Palermo, Italy

Received 5 September 2025 / Accepted 16 December 2025

ABSTRACT

Context. Stellar soft X-ray ([1, 100] Å) and extreme ultraviolet (also EUV, [100, 920] Å; jointly, XUV) radiation affect the evolution and chemistry of exoplanet atmospheres. It is, however, uncertain to what extent the radiation's short-term variability contributes to these effects. The answer might indeed depend on the atmospheric composition in ways that remain largely unexplored. We are interested in what this variability might imply for planets around M dwarf stars and focus on Proxima Centauri (Prox Cen) for three reasons: It is an active M dwarf with high levels of variability; it hosts a likely terrestrial exoplanet within its habitable zone (HZ) that will be a prime target for future direct imaging; and its proximity has led to extensive observations, yielding some of the best available X-ray data.

Aims. We set out to produce time-resolved XUV spectra of Prox Cen that will serve as input to atmospheric models and to characterize the star's intrinsic variability and uncertainties in the inferred spectra.

Methods. We analyzed the entire dataset of archival XMM-Newton observations for Prox Cen. To derive the time-resolved X-ray spectra, we implemented a new pile-up correction, a new adaptive time-binning algorithm, and a time-dependent plasma model selection. The estimated EUV spectrum is based on a published template that we scale with proposed relationships between X-ray and EUV fluxes.

Results. We produced spectra of Prox Cen from 1 to 920 Å over ~260 ks of observations with unprecedented time resolution. The instantaneous X-ray flux of Prox Cen varies between about 20 times and one-fifth of the average value over the available baseline, with significant differences between wavelengths. We further quantified how variability affects the estimated average flux when a limited number of snapshots (each typically of 30 ks exposure) are available, as is common in X-ray surveys. Future investigations of the planet atmospheres of Prox Cen should include the time variability and uncertainties described here.

Key words. planets and satellites: atmospheres – stars: activity – stars: flare – stars: late-type – stars: low-mass – planetary systems

1. Introduction

The evolution of planetary atmospheres is affected by the stellar environment. In particular, soft X-rays and extreme ultraviolet radiation (jointly, XUV) may lead to strong heating of the upper atmosphere, causing hydrodynamic escape (or photoevaporation) and erosion (Yelle 2004; García Muñoz 2007; Colombo et al. 2025). XUV-induced photoionization also affects the chemical balance of the atmosphere, occasionally resulting in detectable amounts of excited states of the H and He atoms, for example (Yan & Henning 2018; Czesla et al. 2022). Stellar emission at these wavelengths varies by several orders of magnitude over timescales from minutes to gigayears (Favata & Micela 2003; Güdel 2004; Kowalski 2024). The atmospheric response to XUV radiation is complex, and it remains unclear whether different temporal distributions of irradiation, even with the same total deposited energy, affect atmospheric evolution differently. It also remains unclear to what extent atmospheric composition dictates this response. Characterizing the XUV radiation over time and as a function of wavelength (or equivalently, energy) is critical for atmospheric modeling (Binder et al. 2024; Amaral et al. 2025; Chen et al. 2025; García Muñoz 2025; Rockcliffe et al. 2025).

Flares, which are responsible for most of the short-term (seconds to days) variability of a star's emission, are events in which the plasma in the corona, transition region, and chromosphere is impulsively heated by magnetic reconnection. They are stochastic phenomena, and it is generally agreed that their occurrence frequency follows a power-law distribution with energy, $dN/dE \propto E^{-\alpha}$ (Kowalski 2024). During flares, ionizing radiation increases by up to several orders of magnitude, and the spectral shape of the emission hardens, with the XUV component being particularly enhanced. This hardening is especially relevant when considering planetary atmospheres with high metallicity, as the photoionization cross sections are typically large for the heavy elements in the X-ray band. It has been shown that in some systems, the flaring component enhances the atmospheric evaporation rate and shortens atmospheric lifetimes (Lee et al. 2018; Amaral et al. 2022, 2025). Moreover, single high-energy flares drive rapid changes in the chemistry, which may alter the short-term climate (Segura et al. 2010; Konings et al. 2022; Ridgway et al. 2023; Chen et al. 2025). Lastly, flares are frequently associated with coronal mass ejections (CMEs), events involving the release of large amounts of plasma from the star (Yashiro & Gopalswamy 2009; Webb & Howard 2012; Reames 2021), which can also impact atmospheric escape rates (Cohen et al. 2011; Khodachenko et al. 2014; Hazra et al. 2025; Vidotto 2025). This highlights how important it is to know the

* Corresponding author: andrea@damonte.iit

flaring properties of the host star when studying the evolution of a planet.

Recently, the exoplanet community has been particularly interested in (low-mass, $0.08 < M_*/M_\odot < 0.6$) M dwarf stars. They are very abundant in our galaxy, and their small sizes enhance the planet-to-star size contrast. Because of their cooler effective temperatures, the habitable zone (HZ) (Kasting et al. 1993) lies considerably closer than the HZ for more massive stars. This makes them excellent targets for the detection and characterization of small HZ exoplanets (Dressing & Charbonneau 2013, 2015). HZ planets around M dwarfs are subject to XUV fluxes comparable to or greater than those received by HZ planets around more massive main-sequence stars. Furthermore, while the coronal activity of the latter usually decreases considerably after the first few hundred megayears, low-mass stars usually remain active much longer (Pizzolato et al. 2003; Preibisch & Feigelson 2005). Higher coronal activity results in both higher quiescent fluxes and an increased probability of large flares. For these reasons, the total amount of XUV energy absorbed by HZ planets around M dwarfs over their lifetimes is often considerably higher than for planets orbiting more massive stars (Johnstone et al. 2021).

This work is motivated by two poorly explored questions, namely, the extent that short-term stellar variability can influence atmospheric retention over gigayear timescales and how significantly it can affect the strength of certain atmospheric observables. Before addressing these questions through detailed atmospheric modeling, we first characterized the stellar XUV emission with temporal resolution matching flare timescales (minutes to days). That is one of the goals of the present work, but we also quantify how the observing cadence and different choices in data reduction and plasma modeling affect the derived spectra.

We focus on Prox Cen, the closest star to the Sun at a distance of 1.30 pc (Gaia Collaboration (DR3) 2020). Prox Cen is an M4.5–M7 dwarf (Maldonado et al. 2020; Bessell 1991; Gaidos et al. 2014) and one of the best-studied low-mass stars. It has a mass of $0.12 M_\odot$, a radius of $0.15 R_\odot$, and a rotation period of ~ 87 days (Boyajian et al. 2012; Kiraga & Stepien 2007). It hosts two small non-transiting planets, Prox Cen b and d, with minimum masses of ~ 1.1 and $\sim 0.3 M_\oplus$ and orbital periods of 11.2 and 5.1 days, respectively (Anglada-Escudé et al. 2016; Kipping et al. 2017; Mascareño et al. 2025). Planet b orbits within the host star’s HZ, making the system a key target for future direct imaging missions (Currie et al. 2023; Blind et al. 2024). Despite its estimated age of ~ 4.9 Gyr (Bazot et al. 2016), Prox Cen remains magnetically active. An activity cycle is suspected but remains unconfirmed, particularly in the X-rays. Wargelin et al. (2024) and Mascareño et al. (2025) report possible cycles with periods of ~ 8 and ~ 18 years, respectively, although the modulation is mainly observed in the optical and infrared. A detailed X-ray analysis by Ayres (2025) did not confirm any cycle, as often occurs in active stars near the saturated regime, characterized by high L_{bol}/L_X ratios. Prox Cen frequently flares over most of the electromagnetic spectrum (Pavlenko et al. 2019; Howard et al. 2022; Güdel et al. 2004). Due to its proximity, interstellar absorption of X-rays is minimal.

These characteristics have led to extensive observations of Prox Cen, and consequently its coronal X-ray emission is one of the best-characterized among M dwarfs, with broad multi-instrument coverage across X-ray and other wavelengths (Haisch et al. 1980, 1995; Güdel et al. 2004; Fuhrmeister et al. 2011; Ribas et al. 2017; Wargelin et al. 2024). Nevertheless, no time-series spectra of the star are currently available for atmospheric

Table 1. XMM-Newton observations used in this work.

OBSID	Date	Exposure (ks)	PI
0049350101	2001-08-12	67	M. Güdel
0551120301	2009-03-10	29	
0551120201	2009-03-12	31	C. Liefke
0551120401	2009-03-14	29	
0801880201	2017-07-27	22	
0801880301	2017-08-16	31	
0801880401	2017-09-07	27	B.J. Wargelin
0801880501	2018-03-11	25	
		Total: 261	

Notes. Observation 0049350101 is analyzed in, e.g., Güdel et al. (2002, 2004) and Reale et al. (2004). Observations 0551120301, 0551120201, and 0551120401 are analyzed in, e.g., Fuhrmeister et al. (2011). Observations 0801880201, 0801880301, 0801880401, and 0801880501 are analyzed in, e.g., Wargelin et al. (2024).

modeling. In this work, we produce such spectra in preparation for follow-up work that will investigate how various putative atmospheres for the Prox Cen planets respond to them. To that end, we used archival data from the XMM-Newton space telescope (Schartel et al. 2024).

The work presented in this paper is organized as follows. In Sect. 2, we present the observations of Prox Cen that we analyzed and the applied data reduction procedure. We give specific attention to the pile-up correction (Section 2.3), the time-binning procedure (Section 2.4), and the plasma modeling (Section 2.5). In Section 2.6, we present the spectral regression method and approach used to estimate uncertainties, while in Section 2.8 we apply literature-based scaling laws to derive EUV fluxes from the X-ray measurements. In Section 3 we discuss the variability across different energy bands, compare the obtained spectra with previous literature, and evaluate the impact of limited observational coverage on the inferred fluxes. Finally, Section 4 contains a summary of the main conclusions.

2. Observations and data reduction

2.1. Observations

We downloaded the Prox Cen data from the XMM-Newton (Schartel et al. 2024) Science Archive¹. The observations span the period 2000–2019, with a total exposure time of about 250 ks. Table 1 contains specific information on them. We used data from all three European Photon Imaging Camera (EPIC) instruments, namely MOS1, MOS2, and pn (Strüder et al. 2001; Turner et al. 2001) and from the Optical Monitor (OM) (Mason et al. 2001), which was equipped with the *U* filter in all observations.

2.2. Standard XMM-Newton pipeline

We used the current calibration files from release no. 406² (Ballet et al. 2024) and ran the standard pipeline of the XMM-Newton Science Analysis System (SAS³; ESA (2023), v. xmmxas_20230412_1735-21.0.0). X-ray raw data consist of a list of events, each corresponding to a detected photon with arrival

¹ [nxsa.esac.esa.int/nxsa-web/](https://nasa.esac.esa.int/nxsa-web/)

² cosmos.esa.int/web/xmm-newton/ccf-release-notes

³ cosmos.esa.int/web/xmm-newton/xsa

time, sensor position, energy (in instrument channels), and pattern. The “pattern” is an integer identifying the geometry of the charge distribution generated by the photon on the sensor. These detections are hereafter referred to as “counts” or “events.”

Following the SAS user guide⁴ (ESA 2023), we processed MOS, pn, and OM data with the *emchain*, *epchain*, and *omfchain* tasks, respectively. For MOS events, we also ran *emanom* to exclude anomalous CCDs (Charge-Coupled Devices) (SOC 2023). For each X-ray source, we identified its centroid and extracted the counts within a circular region. A 30 arcsec radius was used for pile-up correction (see Appendix A), while the final extraction radius was optimized for signal-to-noise via the *regionanalyze* task.

As in Güdel et al. (2004) (from now on G+04), we selected background regions for MOS1 and MOS2 from outer CCDs to avoid source-contaminated areas in the main CCD. An annulus covering active, non-anomalous CCDs was used. For pn, the background was extracted from the same CCD as the source. Prior to this, we excluded all source regions listed by the pipeline processing system (PPS)⁵. False positives near the target were removed manually.

We checked for flaring soft proton background (SPB, Kuntz & Snowden 2008) in each observation. This was done by comparing out-of-field count rates with user guide thresholds (SOC 2023) and by examining the output of the *espfilt* task. SPB levels were low or negligible and significantly weaker than the source; therefore, no good-time-interval correction was applied. Given the use of a wide background region, a standard background subtraction was deemed sufficient.

2.3. Pile-up correction

When the photon flux is high, the probability of multiple photons hitting the camera during the same frame exposure, and close enough on the detector to be indistinguishable, becomes non-negligible. This instrumental effect, known as pile-up, causes the onboard software to interpret two or more photons as a unique event, leading to flux loss. The resulting event has an energy equal to the sum of the true photon energies (causing spectral distortion) and, often, a higher pattern value, as more pixels than usual are typically involved. This latter effect, called pattern migration (Ballet 1999; Jethwa et al. 2015), contributes to both the flux loss and the spectral distortion. If uncorrected, pile-up causes the retrieved spectrum to appear both fainter and harder than the true stellar emission.

The available observations, except for observation 00493501, use a frame mode with a reduced active sensor area, enabling shorter readout times that reduce pile-up occurrence. This is enough to make pile-up negligible during quiescent periods, when the count rate stays well below the conservative thresholds of 3 s^{-1} and 1.5 s^{-1} for pn and MOS, respectively (Jethwa et al. 2015), but not during bright flares, when the count rate reaches up to 70 s^{-1} and 22 s^{-1} for pn and MOS respectively. We have devised a procedure to minimize the effects of pile-up that we briefly summarize in what follows and that is thoroughly described in Appendix A. As a takeaway message, we find that during the brightest flares, the flux loss (the relative number of counts missed when pile-up is uncorrected) can reach up to 22% for pn and 13% for MOS1 and MOS2 as shown in Figure A.1.

⁴ https://xmm-tools.cosmos.esa.int/external/xmm_user_support/documentation/sas_usg/USG/

⁵ cosmos.esa.int/web/xmm-newton/pipeline/

Since pile-up depends on the instrument mode, a consistent correction is required when combining instruments or analyzing datasets taken with different modes. In addition, pile-up is time-dependent, as it scales with the instantaneous flux, and its implications are potentially amplified in the data reduction process. These considerations pose various challenges in time-resolved investigations similar to ours.

There is no single prescription to address pile-up. It is common to exclude the affected time intervals or remove the affected instrument from the entire analysis. After several tests, we found that the standard approaches (see Appendix A) do not provide a suitable correction for our pn+MOS1+MOS2 time-resolved analysis. For them, the approach is not feasible across all instruments, or it reduces the number of counts too much, introducing uncertainties larger than the pile-up effect itself.

Our newly developed pile-up correction is based on ideas from Ballet (1999) and Molendi et al. (2003), which seek to enhance the single-pattern filtering strategy. The standard correction excludes all photon events that interact with more than one pixel on the detector. This strongly lowers the spectral distortion, as the probability of piled single-pattern events is typically negligible. A consistent amount of flux loss, however, remains even if the single-pattern filtering is accounted for in the response matrix, because of the pattern migration effect. Indeed, many events that were originally singles may have merged into doubles or higher patterns, which are now excluded. Accurately estimating how many of these excluded events should have been singles would require high signal-to-noise data and precise pattern-energy relations from the calibration. We simplified the correction by assuming that all flux loss is caused by singles merging into greater patterns (in principle there are also double plus single and double plus double events). This corresponds to the worst-case scenario of flux loss, which we consider sufficient for our goals.

To estimate the number of missing events, we used the empirical relation between count rate and count loss from Jethwa et al. (2015). The same number of events was generated using the energy distribution of the singles as a reference and these were added to the raw event list. Again, the full correction procedure, including the estimation of flux loss, generation of missing events, and results evaluation, is described in Appendix A.

2.4. Time binning

To obtain reliable X-ray fluxes model regression is required (Arnaud 1996). Instrumental counts are grouped into time intervals and an instrumental spectrum, i.e., the spectrum formed by the counts per energy channel, is obtained for each one of them. A plasma model is then selected and convolved with the response matrix. Finally, the model is fit to the spectrum to retrieve the physical quantities. When the source varies in flux and spectral shape, the time-binning strategy strongly affects the results. This is due not only to higher uncertainties at finer time resolution (there will be fewer counts per bin at a given count rate), but also to phase effects, i.e., where interval start times are placed. These effects are critical during flares, which can be smoothed and distorted by contamination from quiescent and cooling phases surrounding the peak.

We here propose an innovative time-binning algorithm to optimize the balance between time resolution and signal-to-noise ratio (S/N) while minimizing the impact of phase effects. The core of our strategy is a sliding window approach, where consecutive intervals overlap. Both the duration of the interval and the step size are variable functions of the count rate. This

dynamic adjustment ensures that S/N and time resolution remain within predefined thresholds. The algorithm operates by iteratively selecting the interval length and then determining the start of the subsequent interval. To select the interval length Δt , we follow the set of rules described in Equation (1). We start from the very first event of the event list. The interval length is selected as the maximum value between two: $\Delta t(C_{min})$, which is the interval length obtained by binning the first C_{min} counts after the start of the interval, and Δt_2 , which is obtained from the second equation of the system. Δt_2 is selected as the minimum between two values: δ_{max} , a maximum time resolution, and $\Delta t(C_{max})$, which is the length defined by the C_{max} counts after the start of the interval. We further comment on the reasons behind this algorithm below. To select the next interval start, i.e., the shift of the sliding window, we determine how many times the same counts are allowed to appear in different intervals. This is translated into a counts-percentage shift parameter, ζ , which defines the percentage of counts of the previously selected interval to be excluded from the event list. We then restarted from point (1) using the new event list.

$$\begin{cases} \Delta t = \max(\Delta t(C_{min}), \Delta t_2) \\ \Delta t_2 = \min(\delta_{max}, \Delta t(C_{max})) \end{cases} \quad (1)$$

Here, we comment on the reason for each parameter and the adopted values. The term C_{min} ensures a minimum S/N, and it limits time resolution during low-flux periods. It was fixed to 10^3 counts. The term δ_{max} prevents excessively fine (unnecessary for our goals) time resolution at high count rates. If the interval for C_{min} counts is shorter than δ_{max} , counts are added until reaching δ_{max} to improve S/N without sacrificing the desired resolution. It was fixed to 180 s. The term C_{max} limits the maximum counts per interval. If a spectrum has too many events, the dominant source of error changes from the low statistics to the instrumental plus model systematic, which renders comparisons less correct. It was set to 10^4 counts. The term ζ (step 2) represents a percentage defining how much of the previous interval's counts are excluded before starting the next interval. It was fixed at 33% to balance overlap and independence of adjacent intervals and their total number.

We applied the algorithm to background-subtracted, pile-up-corrected data within the [0.18, 12] keV energy range, combining counts from all three instruments. Consequently, time resolution also depends on the number of active instruments at each time.

A few consequences of our procedure are as follows. First, the time resolution is not uniform and depends on the S/N. Second, the superposition of consecutive intervals reduces the smoothing that naturally results from time binning. Third, consecutive intervals are not independent, as they share part of their counts. This does not affect the total fluxes, but it may need to be considered when performing certain time analyses.

From eight observations, totaling 260 ks, the algorithm produced about 3200 intervals with an average duration of 300 s, ranging from 150 to 2000 s. Because of varying time resolution and interval overlap, we linearly interpolated the results using *scipy.interpolate* to generate a uniformly binned time series with 30-second resolution.

2.5. Plasma modeling

The coronal plasma is usually considered an optically thin collisionally ionized plasma (thermal plasma). Güdel et al. (2004) confirmed that this assumption also holds for Prox Cen. The

emission is highly dependent on temperature, density and elemental abundances, properties that also vary with time and are non-uniformly distributed in space for most stars. Spectroscopic measurements necessarily average information both over time and over the visible surface of the star since for stars other than the Sun no spatial resolution is available. Flares further complicate the description of this plasma, being unpredictable in time, limited in space (over the star's surface), and spanning temperatures from hundreds of thousands to tens of millions of degrees.

The standard approach involves looking at the thermal structure of the plasma. The emission measure quantifies how much plasma contributes to the emission at a given temperature. Measuring how this quantity varies as a function of temperature yields the differential emission measure (DEM). It is possible to derive it from high-resolution spectroscopy for the brightest sources. This is done looking at single emission lines of different elements (sensitive to different temperatures) and comparing their intensities.

This analysis was performed by G+04 for Prox Cen by dividing the major flaring events of observation 0049350101 into different time intervals. Fuhrmeister et al. (2022) (from now on F+22) did the same but using Chandra X-ray Observatory and Hubble Space Telescope data. They divided the data into flaring and non-flaring intervals and analyzed them separately. Sanz-Forcada et al. (2025) (from now on SF+25) used all the available XMM-Newton observations as well as Space Telescope Imaging Spectrograph, Far Ultraviolet Spectroscopic Explorer, and Extreme Ultraviolet Explorer observations. They did not distinguish between flaring and non-flaring intervals.

When performing time-resolved analyses though, it is more common to use simpler models and low- to medium-resolution data, as we do in this paper. This is because the S/N is usually too low for a time-resolved high-resolution analysis. The usual strategy (Ciaravella et al. 1997; Güdel 2004; Pandey & Singh 2008; Osten et al. 2010; Pillitteri et al. 2022) is to approximate the temperature distribution with a few effective temperatures, either free or fixed on a grid. The standard model used to describe coronal plasmas at a fixed temperature is the Astrophysical Plasma Emission Code (APEC) (Smith et al. 2001). Depending on the S/N and on the goal, we can use the sum of one or more APEC components. The parameters of each component are the temperature, the elemental abundances and the norm (directly related to the emission measure). The abundances are usually given relative to a fixed reference table, usually the Solar one. In this work we use the abundance tables from Wilms et al. (2000). In the APEC model the elemental ratios are fixed to that table and only the overall metallicity can vary. A vAPEC model also exists where each elemental abundance can be adjusted separately.

Our model is described by the following expression: Model = TBabs * (APEC₁ + ... + APEC_N). The sum of multiple APEC components describes the multi-temperature plasma.

TBabs is a convolutional model that accounts for interstellar medium absorption (ISM). The hydrogen column value, i.e., the parameter of the TBabs model describing the ISM absorption, was fixed to 10^{18} cm² as done in Wargelin et al. (2017, 2024) for each of our models. This value is small enough to have a nearly negligible effect on the energy range we are considering, as expected for the closest star to us.

When modeling low-resolution spectra the information in the data is often insufficient to retrieve detailed information about the structure of the plasma, in particular for a time-resolved analysis. This also depends on the available prior information about the star and its plasma. For temperature and normalization

parameters, two approaches can be distinguished: (1) if the thermal structure is not known, a common approach is to represent the plasma with a few effective temperatures, usually between one and four. The temperatures are left free so that the dominant components are set by the data. The required number of components is the minimum that provides an acceptable fit, accounting for the S/N, luminosity, distance, absorption, and background. (2) Using either a previously derived DEM(T) or the temperatures obtained in point (1), a set of temperatures can be fixed while leaving only the normalizations free. This reduces the number of free parameters and allows more complex models to be fit in a stable and tractable way.

The plasma metallicity is the last parameter. In principle, it varies during flares, as detected in previous studies (Favata & Schmitt 1999; Güdel et al. 2001), but this issue is complex (Güdel & Nazé 2009) and beyond the scope of this work. Moreover, since our intervals contain few counts and have a low spectral resolution, we did not attempt to constrain elemental abundances, either as a function of time or globally. We excluded high-resolution RGS data because their time resolution cannot be matched to ours, and abundance estimates from longer integrations were already reported by G+04, SF+25, and F+22. RGS analyses also cannot constrain the H abundance and require adopting a reference ratio. The Fe/H ratio is usually taken as reference since Fe lines dominate the emission, ensuring that the main impact of abundances on the flux remains fixed even when other lines are constrained by the RGS measurements (assuming that the relative abundances found for the other elements are not excessively large; otherwise their contribution could be compared to the Fe one). To test whether a more detailed abundance calibration could improve the fits, we modeled the flare in observation 0049350101 using six vAPEC components selected following the RGS time-resolved analysis of the same flare in G+04. As shown in Figure 2, this model performs worse than the simpler one in terms of information criteria (AIC). We therefore deemed further tests with vAPEC components and fine-tuned abundances unnecessary.

In the APEC models the metallicity was fixed to 0.5, as done in G+04 and SF+25. Although other works (e.g., F+22) assumed solar metallicity, a subsolar value is closer to the average obtained when the abundances are allowed to vary in our retrieval. Indeed, when the metallicity is left free, its retrieved value ranges from 0.2 to 1.4 (depending on the number of temperature components), leading to large variations in the derived EMs while the fluxes are less affected. To avoid possible non-physical trends in the EMs, we fixed the metallicity to a value drawn from the literature, since allowing it to vary does not significantly affect the fluxes. We tested a subset of observations to evaluate the impact of metallicity as another possible source of uncertainty. Varying the metallicity between 0.2 and 1 changes the total flux by at most $\sim 20\%$, but the median difference remains below 2% and therefore within the flux uncertainties.

Additionally, the model choice may affect the final fluxes and, even more, the spectral shape, as shown in Appendix B. We used different models to test different hypotheses and the sensitivity of the retrieved synthetic spectra to the model choice:

1. 2 APEC with free temperatures: the structure of this model is described in detail in Table 2. No prior information was assumed except for the abundances, fixed to 0.5. It corresponds to description (1) above. We decided to use two temperature components. Our goal with this model is to get an idea of the temperature range of the coronal plasma and

Table 2. Model 1 parameter relationships.

Par	State	Val	Link	Delta	Min	Max
T1	Free	0.2	–	1	0.1	5
Z1	Fixed	0.5	–	–	–	–
norm1	Free	1e-3	–	1	1e-7	2
T2	Linked	–	T1+ ΔT	–	–	–
Z2	Fixed	0.5	–	–	–	–
norm2	Free	1e-3	–	1	1e-7	2
N_H	Fixed	10^{18}	–	–	–	–
ΔT	Free	0.5	–	1	0.1	5

Notes. Redshift is fixed to 0. The free parameter ΔT is defined as a custom additive model with norm fixed to 0. T_i are the temperatures in keV, Z_i are the relative abundances, and $norm_i$ the norm values. N_H stands for the hydrogen column density, in units of atoms/cm².

compare the results with the DEMs described in the literature. Specifically, we want to test which temperatures are reached for tuning the subsequent models.

2. 3 APEC: we fixed the temperatures (three of them in total) and only looked for the norm parameters. The temperature values we used were 0.2 keV, 0.7 keV, and 1.2 keV. These values are based on the results from the 2 APEC model and the DEM(T) in G+04 and SF+25.
3. 5 APEC: in this case we added two hotter temperatures to the 3 APEC case that cannot reproduce the high-energy emission in the simpler model, in particular during flares. The values we used are 0.2 keV, 0.7 keV, 1.2 keV, 1.5 keV, and 2.7 keV.
4. 6 vAPEC: in this case, starting from the G+04 analysis of the flare in observation 0049350101, we assumed that the six DEMs derived for different time intervals (that they identify as quiescence, rise, flare peak, decay, secondary flare peak, following decay) may be used as a base to describe the general behavior of Proxima. We approximated each DEM with a vAPEC component with the effective temperature fixed to the DEM's highest value. In this model, when available from Güdel et al. (2004), we fixed the abundances to the reported ones; otherwise, we adopted solar abundances.

2.6. Regression and uncertainties

We simultaneously fit the MOS1, MOS2, and pn data. We used the SAS tasks *evselect*, *backscale*, *arfgen*, *rmfgen* to extract the source spectrum and its relative background, together with their response matrices, for each time interval identified with the binning algorithm. The task parameters were left unchanged from their default values except for the point spread function (PSF) energy that was set to 1 keV as Prox Cen's spectrum is usually centered around this value. We used the energy range within which all the instruments are calibrated, i.e., [0.18, 12.0] keV.

To perform the spectral regression, we used XSPEC (Arnaud 1996). We used the C-stat⁶ statistic, given the few counts per interval. As usually prescribed for the C-stat implementation in XSPEC, the energy bins were grouped to contain at least one count each. To maximize the likelihood and estimate the uncertainties of the parameters we used the XSPEC

⁶ Strictly speaking, XSPEC uses the B-stat when background is subtracted.

built-in Monte Carlo Markov chain (MCMC) method. To get the starting values for the chains we first fit the data with a Levenberg–Marquardt likelihood maximization algorithm. Comparison of the final results with those obtained using the initial Levenberg–Marquardt algorithm shows that, for all our models, the best fit does not change significantly. Using the MCMC speeds up and simplifies the error computation.

We generated synthetic unabsorbed spectra using the best-fit parameters and a resolution of 0.01 \AA . To consider uniform intervals and use a standard range for X-rays, the synthetic spectra are generated in the range $[0.12, 12.0] \text{ keV}$ corresponding to $[1, 100] \text{ \AA}$.

Uncertainties at the 1σ level were estimated for the model parameters using the XSPEC *error* command, which automatically uses the loaded MCMC chains to compute them. For the flux errors, using the standard *cflux* convolutional model would have needed an additional fit for each energy bin so we did not use this approach. We used the information contained in the chains instead. We selected the 68 percentile of vectors with the best statistic from the chains. From them, we selected 100 representative vectors and for each of them we computed the total flux. We took the maximum and minimum fluxes obtained in this way as upper and lower errors.

The 100 vectors were selected as follows: 98 of these vectors of parameters were randomly extracted. To these we added the vectors with the smallest and the biggest likelihood from the selection of vectors that are within 1σ . In this way we account for different shapes in the energy space as the *cflux* method does, but without the need of performing MCMC and error computation again. To estimate the errors for each energy bin of the synthetic spectrum we consider the maximum and minimum values occurring in each bin among the 100 spectra. This overestimates the errors because the energy bins are not independent but strongly related to the spectral shape, but we consider it reasonable as a conservative approach that also provides error estimates as a function of energy.

We highlight that 1σ refers to the criterion used to select the chain vectors. It is related to the errors of the fluxes and parameters, but the latter are not necessarily Gaussian and may be asymmetric, so when considering them we need to track the upper and the lower errors separately. From here on, when using the terms lower and upper error, we always refer to these errors or to errors propagated from them.

Errors on the total flux for the different time intervals are shown in Figure 1 as a function of the total flux. We identified three ranges by eye: fluxes smaller than $1.5 \times 10^{-11} \text{ erg s}^{-1} \text{ cm}^{-2}$, fluxes greater than $6 \times 10^{-11} \text{ erg s}^{-1} \text{ cm}^{-2}$ and intermediate fluxes. For each interval, we computed the average upper and lower uncertainties, represented as horizontal segments in the Figure. The spread is moderate, with uncertainties consistently below 20%. Nearly all flux values above the upper threshold originate from the peak of the flare in observation 0049350101.

Based on this and on the results of the pile-up correction presented in Appendix A, we applied the correction only when the estimated flux loss exceeds the uncertainty on the total flux, specifically for fluxes above $1.5 \times 10^{-11} \text{ erg s}^{-1} \text{ cm}^{-2}$, which correspond to the lower of the two thresholds defined above. This way we retained a better S/N ratio during the low-count time intervals, where pile-up is less significant. The total flux reference to perform this selection is obtained as the average between the fluxes obtained with the three APEC models for the two different datasets. We used the three-temperature model because we used it to define the thresholds.

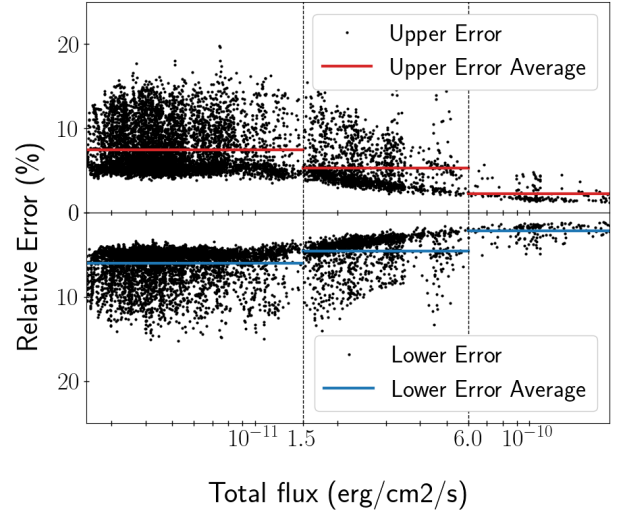


Fig. 1. Upper and lower relative errors on the total integrated X-ray flux as a function of total flux for the three-temperature model. Vertical dashed lines delimit the ranges selected by eye. Red and blue segments represent the average errors in the corresponding ranges.

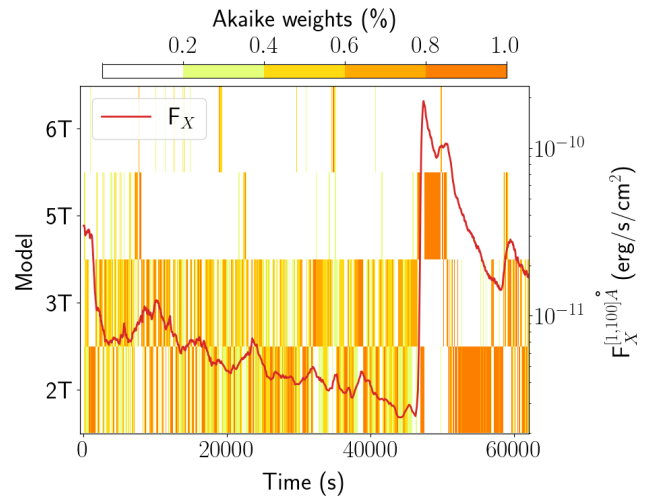


Fig. 2. Comparison of ω_i for the different intervals of observation 0049350101. The left axis represents the model; the right axis shows the wavelength-integrated flux light curve ($[1, 100] \text{ \AA}$). The values of the Akaike weights for each model as a function of time are color coded as indicated by the color bar.

2.7. Time-dependent model selection

Next, we selected the best-fitting model as a function of time by using the Akaike information criterion. We computed it for the four different models described in Section 2.5; each had a different number of temperatures (2, 3, 5, and 6) and settings (free temperatures, temperature grid, APEC, or vAPEC models). For each time interval, we selected between the dataset without pile-up correction (NOCORR; see Appendix A) and the one obtained with the new correction (NEWCORR). The selection is based on the total flux threshold defined in Section 2.6.

Assuming that these are the only possible models, we can use the Bayesian approach, to quantify the relative probability of each model being the one that generates the data. For this

purpose, we computed the Akaike weights, ω_i , defined as

$$\omega_i = \frac{\exp(-\frac{1}{2}(\text{AIC}_i - \text{AIC}_{\min}))}{\sum_{m=1}^M \exp(-\frac{1}{2}(\text{AIC}_m - \text{AIC}_{\min}))}. \quad (2)$$

In this equation, AIC_i is the Akaike information criterion as defined in [Burnham & Anderson \(1998\)](#). i is the index of the model, and the sum goes over the four models ($M = 4$). AIC_{\min} represents the minimum value obtained among the four models. In [Figure 2](#) we show the results of this exercise for observation 0049350101. For the four models we chose, we represent the weights as a function of time. For each interval we have a set of weights that add up to unity. We divide the space of the weights into five intervals, which we represent with a color bar. The ratio of the Akaike weights can be interpreted as the ratio between the probabilities that the corresponding models are the true ones. Accordingly, in our analysis we choose the model with the highest ω_i value at each time step. We also plot the wavelength-integrated flux light curve ($[1, 100] \text{ \AA}$), previously used to select the pile-up correction. We note that it is common for the best-fitting model weight to be smaller than 80%.

We highlight some common behaviors. As one may expect, the number of temperatures of the “best” model overall increases with the total flux. This is because a more complex model may be favored when more counts are available, and because higher fluxes are usually due to flares, which correspond to higher-temperature plasma requiring additional components. The best-fitting model tends to change rapidly in time, even when no evident variation is present in the light curve, and disregards the fact that adjacent intervals are not independent (at least in groups of three, as the shifting parameter is set to 33%) because of the binning procedure. This highlights that the data are rarely sufficient to strongly discriminate between the models. The six-temperature model is almost always discarded, even during the very same flare for which it was calibrated ([Figure 2](#)). This is probably due to its complexity relative to the data quality, and for this reason we exclude it from the analysis hereafter. The two-temperature model is the preferred one during the peaks of the two major events. We attribute this to the fact that at these higher temperatures, two components with free temperature parameters perform better than many fixed-temperature components, since the non-flaring emission is negligible and the temperature of the flaring plasma is changing rapidly and is hotter than the fixed components.

Hereafter, we exclude the six-temperature model. We repeat the selection but considering only the other three models recomputing the weights with just the three of them. From the selection we obtain a reference time series of spectra. Since the time intervals are not uniform, we perform a linear interpolation on the data to obtain evenly sampled time series with 30 s spacing. Further comments on model selection and how this affect the spectral shape, specifically in the end of the (X-ray) wavelength range, are presented in [Appendix B](#).

2.8. EUV scaling

Thus far we have focused on the X-ray energy range covered by the XMM-Newton data. The EUV range is as important in determining the planetary atmospheric evolution. Unfortunately, measurements in the EUV range are difficult due to severe absorption by the ISM and almost always unavailable. To obtain the EUV flux there are two possibilities. Both use information from the adjacent ranges of X-rays and FUV. The best approach is to model the stellar outer layers (chromosphere, transition




Fig. 3. Extreme ultraviolet spectra obtained using the linear (top panel) and quadratic (lower panel) relations from SF+25 scaled from the minimum, average, and maximum X-ray fluxes. The XUV-integrated fluxes are reported on the plot together with the EUV/X-ray ratio for maximum and minimum spectra.

region, and corona) and their plasma, deriving their DEM structure with high-resolution spectroscopy and plasma simulations. This method is beyond the scope of this work and difficult to apply in a time-resolved way. A more pragmatic way involves the use of scaling laws, which is the one used in this work. In this case many assumptions are made and results may change considerably from law to law, so it is important to take these results with a grain of salt. In the literature, various scaling laws are available to derive the EUV luminosity from the X-ray luminosity (e.g., [Linsky et al. 2014](#); [Johnstone et al. 2021](#); [Maggio et al. 2024](#); [Sanz-Forcada et al. 2025](#)). Depending on the adopted strategy and on the amount of data at our disposal, the reconstructed EUV flux may vary greatly. This adds significant uncertainty to the estimated XUV flux and, in turn, to the atmospheric model outputs. To give a quantitative idea of the uncertainties we can only compare the results from different scaling laws at different levels of emission. These differences dominate over the other sources of error. We comment on this in the text below, while discussing [Figure 3](#). Unfortunately, none of the available scaling laws takes into account flares, as they are calibrated with the average luminosity. The use of these scaling laws in time-resolved analyses adds uncertainty, but at the moment it remains the most practical approach. This is the approach we take here.

To gain insight into the uncertainties introduced by the scaling laws, we experiment with two of them and a fixed EUV spectral shape. Our reference work is SF+25, from which we take both the scaling laws and the average EUV spectrum of

Prox Cen. The first scaling law (Equations (2) and (3) of SF+25) is linear, adapted from [Sanz-Forcada et al. \(2011\)](#): the logarithm of the EUV luminosity is expressed as a linear function of the logarithm of the X-ray luminosity. The second scaling law (Equations (4) and (5) of SF+25) expresses the logarithm of the EUV luminosity, normalized to the bolometric luminosity, as a second-order polynomial function of the logarithm of the similarly normalized X-ray luminosity. This new relation is introduced in SF+25, where it is argued that it is more accurate. These relations are presented in SF+25 for two spectral ranges. We use the ones calibrated over the [100, 920] Å range (respectively Equations (2) and (4) of SF+25).

We proceeded as follows. First, we obtained an EUV-integrated flux time series by feeding the X-ray fluxes into the scaling laws. Then we used these values to scale the average fixed-shape EUV spectrum of Prox Cen. For each time interval, the resulting EUV spectrum was stitched to the X-ray spectrum to obtain a time series of XUV spectra.

This procedure is not ideal for different reasons. The first reason is that the relation is based on the average EUV flux of the population. By using it to infer the time-varying EUV flux of an individual star, we implicitly assume that the variability of a single star follows the relation derived from average luminosities of different stars. Secondly, we are using this relation to scale the fluxes measured during short time intervals. These intervals might be dominated by flares or dimmings. This type of emission is not represented by the average, on which the scaling law is calibrated. Moreover, we use a fixed EUV spectral shape for all the time intervals, while at X-ray and presumably also at EUV wavelengths, the spectral shape changes considerably with time, specifically during flares. For these reasons, this procedure often leads to a discontinuity in the spectrum where the X-ray and EUV parts are stitched together. We argue that the discontinuity found for the average flux can be used as a criterion to compare different reconstructions. We stress the need for EUV spectra that describe the star in different flaring states. There is also a need for X-ray versus EUV relations during these conditions that can be used as simple prescriptions for future reconstructions.

In Figure 3, we show the reconstructed XUV spectra obtained for conditions of average, maximum and minimum X-ray flux, using both the linear (top panel) and quadratic (lower panel) scaling laws. We also report the XUV-integrated flux at 1 AU and the ratio between the EUV-integrated and the X-ray-integrated fluxes for maximum and minimum X-ray conditions. The two scaling laws yield ratios that differ by significant factors.

The quadratic relation is presumably the most accurate, as it minimizes the discontinuities in the reconstructed spectrum at 100 Å, according to our criterion. For these reasons, we focused on the corresponding reconstructed spectra. For the brightest flares, the X-ray fluxes exceed the reliability threshold defined in SF+25. We are aware of the limitations of this scaling law, but as we lack a viable alternative, we accepted them and decided to push the relation outside its recommended limits.

To quantify the uncertainty in the EUV flux (not reported in the final spectra), we compared values derived from two different relations in the three regimes: minimum, average, and maximum. For the average flux, the ratio of the total EUV fluxes produced by the two methods is 5.4. The linear relation yields $3.85 \text{ erg cm}^{-2} \text{ s}^{-1}$ (at 1 AU), while the quadratic relation yields $0.71 \text{ erg cm}^{-2} \text{ s}^{-1}$. The corresponding ratios are 3.6 for the minimum flux and 18 for the maximum. The discrepancies for the minimum and average fluxes lie within the expectations

for the relations reported in SF+25. The large discrepancy for the maximum flux arises because the flare-specific relation was applied outside the recommended parameter range in SF+25.

2.9. Data products

The EUV and X-ray spectra are stitched together to obtain the time series of XUV spectra over the 1–920 Å range. Each X-ray spectrum is associated with a set of plasma parameters retrieved from the regression. Light curves are obtained by integrating the spectral flux time series over different energy ranges. In the data products, we also include the light curves measured with the U filter, which cover longer wavelengths. Light curves and spectra are presented and discussed in the next section.

3. Results and discussion

Next, we discuss the variability in our data. We are particularly interested in the energy incident on Prox Cen b, both instantaneously and as a time average. The EUV range is included in the analysis where relevant.

3.1. Short-term variability and flare diversity

In Figure 4, we select three snapshots in time of the light curves. The top panel (second half of observation 0049350101) shows the strongest flare recorded in our dataset, justifying the use of a logarithmic scale. The middle panel (observation 0551120201) depicts one of the three observations in which the stellar emission always stays below average. The lower panel (observation 0801880301) shows the only observation during which the flux is always above the average.

Prox Cen is a frequently flaring star, and its flare strengths vary significantly over time. Many of them, usually the most energetic ones, are detected at both optical and X-ray wavelengths, with the optical peak preceding the X-ray peak, as described by the Neupert effect ([Neupert 1968](#)). Heating, and subsequent cooling, are also evident for many intermediate events, with peaks appearing first in the hard and later in the soft bands. This behavior is visible in the flare at 12 500 s in the middle panel of Figure 4. This variability shows itself also in the more quiescent periods, as noted by F+22, who report that the quiescent count rates vary by about a factor of two within a few days.

Many peaks appear for some energies but not for others, which is plausibly explained by heating acting different in different layers of the star. For example, the main event in the top panel of Figure 4, the two events at ~2500 s of the middle panel and the event at ~12 500 s in the same observation, probably involve a wide region from the chromosphere to the corona as the peak is clearly seen at all wavelengths (including the optical). The loss of resolution in the peaks of the two events at ~2500 s due to time-binning is noticeable.

Other events are visible only at certain wavelengths. For example, in the middle panel of Figure 4, at ~5000 s, a small optical flare does not seem to have an evident X-ray counterpart. At the start of the observation in the top panel and at ~15 000 s in the middle one, there are peaks at high-energy X-ray wavelengths (purple and red) that are undetected at both softer X-ray (green, blue) and optical wavelengths. This might be interpreted as the energy release due to magnetic reconnection happening only in the corona or because the foot-points of the loop generating the event are behind the limb and thus concealed from view. In a different type of event, as seen at ~7500 s in the middle

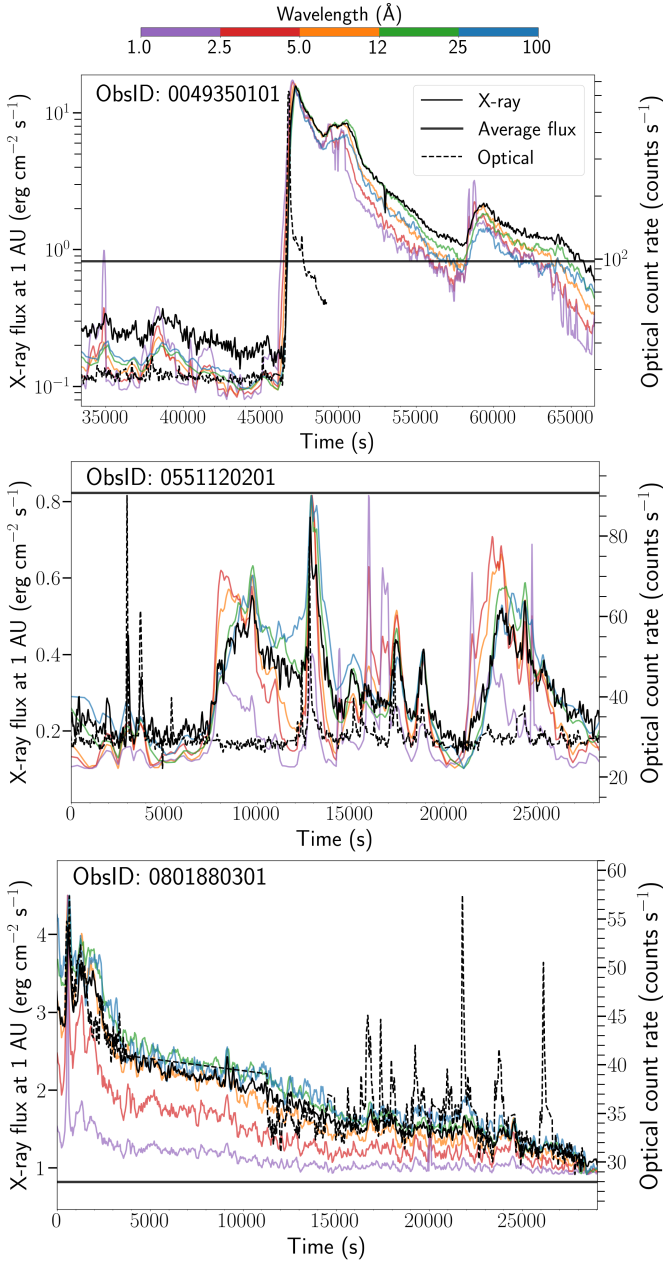


Fig. 4. X-ray flux (1–100 Å; count s⁻¹; solid black line) plotted for selected snapshot observations together with the X-ray average over all observations. The optical count rate is plotted as a dashed black line. Colored lines (see the color bar at the top panel) indicate light curves in different energy bands, obtained by integrating the spectral fluxes. No axis labels or ticks are displayed for these, as they are plotted only for qualitative comparison. Limits are set to (0.8 x min, 1.2 x max) separately for each band.

panel, the high-energy X-ray flux rises with seemingly no optical counterpart. The flux increases by a factor of three, which is comparable to medium-intensity flares. Its timescale is less than 5000 s, and therefore the event is relatively short.

Observation 0801880301 (lower panel of Figure 4) constantly remains above the average flux. Observations similar to this one could considerably bias our estimate of the star’s average behavior, specifically for more distant targets, for which the S/N does not allow time-resolved or spectral analysis. We also highlight how, toward the end of the observation, high variability is measured in the optical range, but not in the X-ray one. Similar

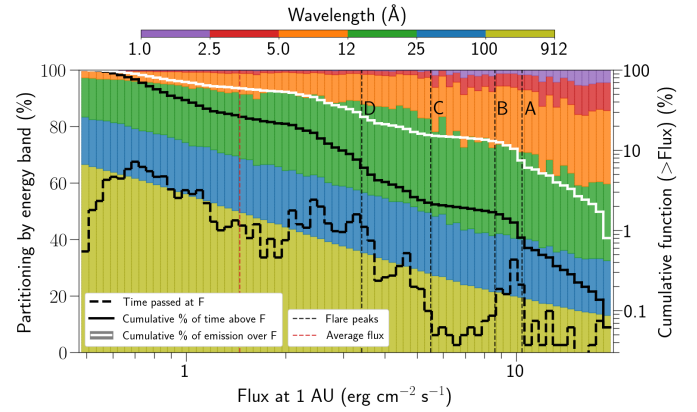


Fig. 5. Statistical properties of Proxima Centauri flux at 1 AU, binned by total flux level. The x -axis shows the flux of the star at 1 AU (integrated from 1 to 920 Å). The flux is logarithmically binned. The right-hand-side y -axis refers to the three bold black and white lines. The three lines share this axis because all quantities are expressed as percentages. The dashed black line represents the percentage of time Prox Cen is emitting at the corresponding x -axis flux bin. The continuous black line is the cumulative curve of the dashed black line. It is showing the percentage of time the star is emitting above the corresponding x -axis flux value. The continuous white line corresponds to the percentage of the total emitted energy of Prox Cen above the corresponding flux value. The vertical dashed black lines indicate important relative maxima in the observations, excluding the absolute maximum, which would have been plotted at the extreme right. From right to left, in order from the most to the least intense, they refer to: the secondary peak of the main flare of observation 0049350101 (A), the peak of the main flare of observation 0551120401 (B), the maximum of observation 0801880301 (C), and the peak at the start of observation 0049350101 (D). The vertical dashed red line is the average flux over all the observations. The left-hand-side y -axis describes the partitioning by wavelength range of each flux bin. They are color-coded by the top color bar. This partitioning is obtained by averaging all spectra in the corresponding flux bin.

patterns were already reported by F+22 looking at X-ray and FUV data.

The major event in the top panel increases the total flux relative to pre-flare values by two orders of magnitude. This rise is even more pronounced in the shortest wavelengths ranges. At the shortest available wavelength (purple curve), the intensity increases by four orders of magnitude.

The variability over timescales between 5 minutes and 10 hours is seen in all observations. For these observations, the emission intensities vary by up to a factor of 30.

3.2. Flux distribution

Next, we study how Prox Cen’s emission is distributed over different flux levels. Figure 5 shows three curves, all based on the total flux of Prox Cen (at 1 AU, integrated over 1–920 Å) on the x -axis. The total flux is binned logarithmically. The curves represent: (black-dashed) the percentage of time the star emits at the flux level of the corresponding x -axis bin; (black-solid) the percentage of time the emission is above that flux; (white-solid) the percentage of the total emitted energy above that flux. All the percentages are relative to the whole set of observations. The three curves share the right-hand-side y -axis.

For each flux bin, we plot (left-hand-side y -axis) stacked bars of different colors to show how the emitted energy is partitioned across six selected wavelength intervals. This partitioning

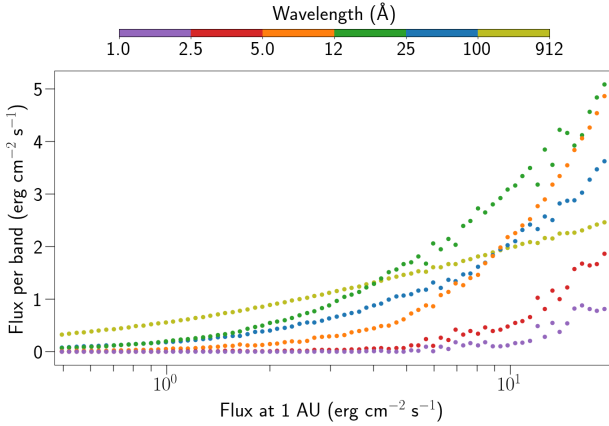


Fig. 6. Average flux over all intervals within each flux bin for different wavelength ranges. The x -axis and the color bar for the wavelength intervals are the same as in Fig. 5. As a reminder, the EUV fluxes are increasingly less accurate as the flux rises (during flares specifically).

is obtained by averaging all spectra falling in the corresponding flux bin. This conveys the spectral hardness of the emission. The different colors always sum up to 100 and the wavelength intervals are described by the top color bar.

Vertical dashed black lines indicate important relative maxima throughout the observations. The absolute maximum (over all the observations) reached during the main flare of observation 0049350101 is not included. From right to left, they refer to the secondary peak of the main flare of observation 0049350101 (A), the main flare of observation 0551120401 (B), the maximum of observation 0801880301 (C), and the peak at the start of observation 0049350101 (D). These events appear as clear changes in the shape of the cumulative distributions. This also highlights that single major events can significantly bias the time-averaged emission.

For different flux levels (x -axis values), the contribution of the EUV wavelengths ranges from nearly 70% to less than 20%. Correspondingly, the contribution of the 1–5 Å wavelengths is typically negligible. Exceptionally, when the stellar emission is strong, the very energetic wavelengths can contribute up to 10% to the total XUV flux. The fact that even rare events contribute significantly to the total emission makes a strong case for including flares in atmospheric models, ideally with the correct spectral shape.

Figure 6 shows the average flux for each wavelength interval within each flux bin. This figure is analogous to Figure 5, but uses absolute values in terms of flux at 1 AU. Wavelengths shorter than 5 Å contribute to the total flux only during the highest emission levels.

Again, we note that the EUV flux strongly depends on the adopted scaling relation, particularly for values far from the average. Therefore, these results should be regarded as one possible realization among several, but still, highlight how different could be the X-ray to EUV ratio as a function of time.

3.3. Reduction uncertainties and literature comparison

We assessed the effects that different reduction procedures and observational baselines have on the reconstructed spectra and the associated uncertainties. We also compared our fluxes with those from other analyses reported in the literature. These spectra are always averaged over multiple intervals, which leads to

significant smoothing of the peaks and to large differences in their relative maxima.

First, we summarized the literature on the XUV spectrum of Prox Cen. We downloaded the published spectra from the corresponding references, and for an easier comparison, we scaled them to 1 AU and binned them onto a common energy grid. Ribas et al. (2017) presented an average spectrum based on ROSAT and XMM-Newton observations. For the latter dataset, a correction factor of 0.83 was applied to the measured fluxes to account for the presumed high activity level of the star at the time of observation (Wargelin et al. 2017). Their spectrum covers the range [7, 100] Å. Loyd et al. (2018) presented two versions of an average spectrum derived from XMM-Newton data, differing in spectral range and resolution. The APEC version retains the native resolution of the XSPEC model, while the Multi version is smoothed and extended over a broader range. Both are based on a two-temperature model and cover the ranges [5, 100] Å and [2.5, 100] Å, respectively. SF+25 derived an average spectrum reconstructed from emission lines in high-resolution XMM-Newton data, covering [1, 100] Å. Binder et al. (2024) produced five reference spectra representative of quiescent, elevated, descending, rising, and flare conditions. Their fitting models include two or three temperature components and cover the range [1, 100] Å.

For comparison, we considered five spectra based on our work in the [1, 100] Å spectral range. We labeled them as follows:

- Minimum: spectrum with the minimum integrated flux in the time series.
- Average: average spectrum over all times.
- Median: spectrum corresponding to the 50th percentile in integrated flux.
- 80th percentile: spectrum corresponding to the 80th percentile in integrated flux.
- Maximum: spectrum with the maximum integrated flux in the time series.

In Figure 7, we plot the integrated fluxes over the available wavelength ranges. In Figure 8, we plot a selection of our spectra and others reported in the literature. Table 3 contains the ratios between the references selected above and the average integrated flux obtained with this work. In the first row, unlike in the other rows, we show the flux values (integrated over the corresponding wavelength range) for the time-averaged spectrum.

These values are computed both in the full spectral range and in different wavelength ranges. The third row of the table, unlike the others, reports the time-averaged flux at 1 AU from Prox Cen calculated in this work. Here are two examples illustrating how to read the table. In the [5, 10] Å band, the flaring flux reported in Binder et al. (2024) is 12^{+1}_{-1} times stronger than the average flux of this work. In the [10, 20] Å band, the flux from Ribas et al. (2017) is $2.7^{+0.3}_{-0.2}$ times smaller than our average. Errors include only uncertainties from this work. Dividing two values in the same column allows for comparison between different references; for example, in the [1, 100] Å range, the maximum flux is $5 \times 20 = 100$ times smaller than the minimum flux.

The [1, 100] Å column of Table 3 can be used to quantify the relations between the values in Figure 7. In Figure 7, we also show, for reference, the Sun’s activity cycle variability range and a flaring Sun value (in orange) from Peres et al. (2000), covering [4.1, 100] Å.

The average fluxes presented in Figure 8 (gray, pink, brown, and azure colors for different sources) are generally in good

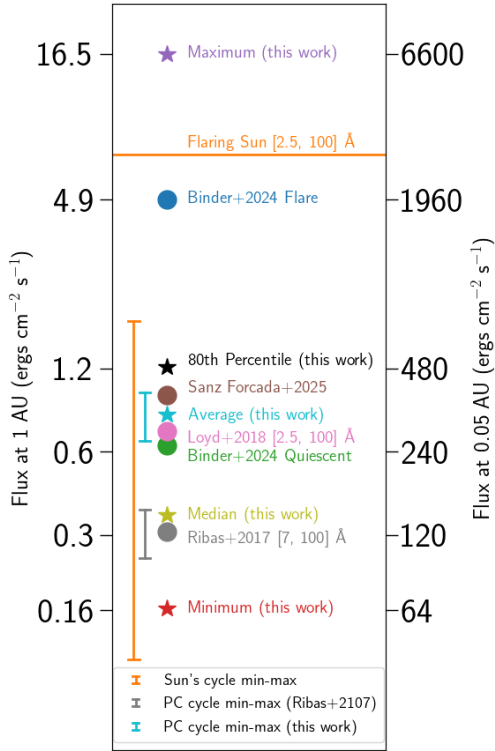


Fig. 7. Integrated fluxes at 1 AU (right-side axis) and at 0.05 AU (left-side axis), corresponding to the orbital distance of Prox Cen b). When not included in the figure, the integration range is $[1, 100]$ Å. Star-shaped markers indicate results from this work, while round markers indicate results from other works. For the Sun (integrated over $[4.1, 100]$ Å), the variability range over its activity cycle in quiescence and a flaring reference are taken from [Peres et al. \(2000\)](#). Assuming a cycle maximum-to-minimum ratio of 1.5 for Prox Cen ([Wargelin et al. 2024](#)), we show the star’s quiescent variability due to its activity cycle. To provide a reference for the amplitude of the activity cycle on the plot, we use two example spectra, the [Ribas et al. \(2017\)](#) spectrum (gray) and our average spectrum (azure), as mid-cycle quiescent flux references.

agreement. Integrated over $[1, 100]$ Å⁷ (see Table 3 and Figure 7), the time-averaged values range from 2.6 times lower to 1.2 times higher than the time-averaged value obtained in this work. This implies that the lowest and highest time-averaged values differ by a factor of ~ 3 . These differences arise from a combination of factors, namely the observational data and instruments used, the calibration and reduction procedures, the plasma modeling and assumed ISM absorption, and the marginally different integration ranges adopted in [Ribas et al. \(2017\)](#) and [Loyd et al. \(2018\)](#).

With a time resolution of 150 s during major events, we are able to generate spectra that capture the emission maxima and minima with little or no smoothing from time binning. This, for example, explains differences between our work and that of [Binder et al. \(2024\)](#). For higher fluxes, the pile-up effect also contributes to increasing the difference in the results.

Next, we assessed to what extent time-resolution affects the average spectrum. In principle, the average could differ depending on whether it is computed as the mean of the time-resolved spectra or from a single spectrum accumulated over all photon events. We test this by the following procedure. For each observation, we fit a three-temperature model to the spectrum

⁷ If the integration range differs, it is reported in the corresponding table or figure.

resulting from all the data from that observation. This spectrum is then compared to the average of the time-resolved spectra for the same observation. The difference in the integrated flux always stays below 10% for all observations.

[Wargelin et al. \(2024\)](#) report an 8-year X-ray activity cycle for Prox Cen, with a maximum-to-minimum amplitude of ~ 1.5 . In this respect, we note the following: 1) over the timescale of XMM-Newton observations (~ 8 days spanning 17 years), as well as additional observations from other telescopes, the flux differences reach a factor of ~ 3 , resulting solely from different datasets combination and data analysis procedures. We note that this factor is double the amplitude of the cycle; 2) On timescales of hours or less, the variability reaches two orders of magnitude, while in the $[2.5, 5]$ Å and $[1, 2.5]$ Å ranges, the variability reaches three and four orders of magnitude, respectively.

[Mascareño et al. \(2025\)](#) report a ~ 18 -year activity cycle, which differs significantly from the 8-year cycle reported by [Wargelin et al. \(2024\)](#). If the longer period is correct, it turns out that Prox Cen also flares significantly during the cycle minimum, as the XMM-Newton observations would not always coincide with the cycle maximum anymore.

The spectra we reconstructed and present in Figure 8 led us to question what we should define as quiescent conditions. In the higher energies the median (yellow-olive) spectrum is noticeably lower than the average (azure) spectrum, more sensitive to extreme values, indicating that flares have a substantial effect on the total emitted energy in this range. More importantly, the median obtained in this work is lower than the quiescence spectrum of [Binder et al. \(2024\)](#) that was using exactly the same dataset. We wonder about the significance of defining a quiescent flux when the median value is smaller.

It is interesting to assess the effects that may arise when observations are sparse, short, and widely separated in time. For most stars, there are at best a few observations, usually lasting no more than one day each, which may also be separated by years. If the star is variable over these timescales, the data may not represent the star’s average behavior. Indeed, over timescales of years, some G stars exhibit high-amplitude activity cycles that change their X-ray fluxes by more than an order of magnitude ([Aschwanden 1994](#); [Favata et al. 2008](#); [Robrade et al. 2012](#)), and the same may hold for M stars. In addition, flares can change the X-ray flux by large factors over hour- or day-long timescales. We tested how this could affect Prox Cen by performing a simple combinatorial exercise, comparing the average flux over all observations with that obtained from a subset.

We have eight observations, for a total exposure time of ~ 250 ks. We split the first observation into two segments, obtaining a total of nine observations with approximately equal durations. Using the average over all nine observations as a reference, we considered cases in which only a subset (1–8) of them is available. For each case, we consider all unique combinations without repetition and compute the corresponding averages. Results for integrated XUV fluxes ($[1, 100]$ Å) are shown in Table 4.

Each row of Table 4 contains a histogram showing how many combinations of observations produce a given ratio between the average flux over the full dataset and that obtained from the subset. The vertical line in each plot represents unity. Bins to the left of this line indicate that the subset average is less than the full-sample average, while bins to the right indicate an overestimate. The ratio of left-to-right combinations is shown as U/O (underestimated-overestimated), indicating whether under- or overestimation of the flux is more likely. Values near 1 (i.e., bins adjacent to the vertical line) are excluded from this analysis. The label “inf” indicates that there are no cases of flux

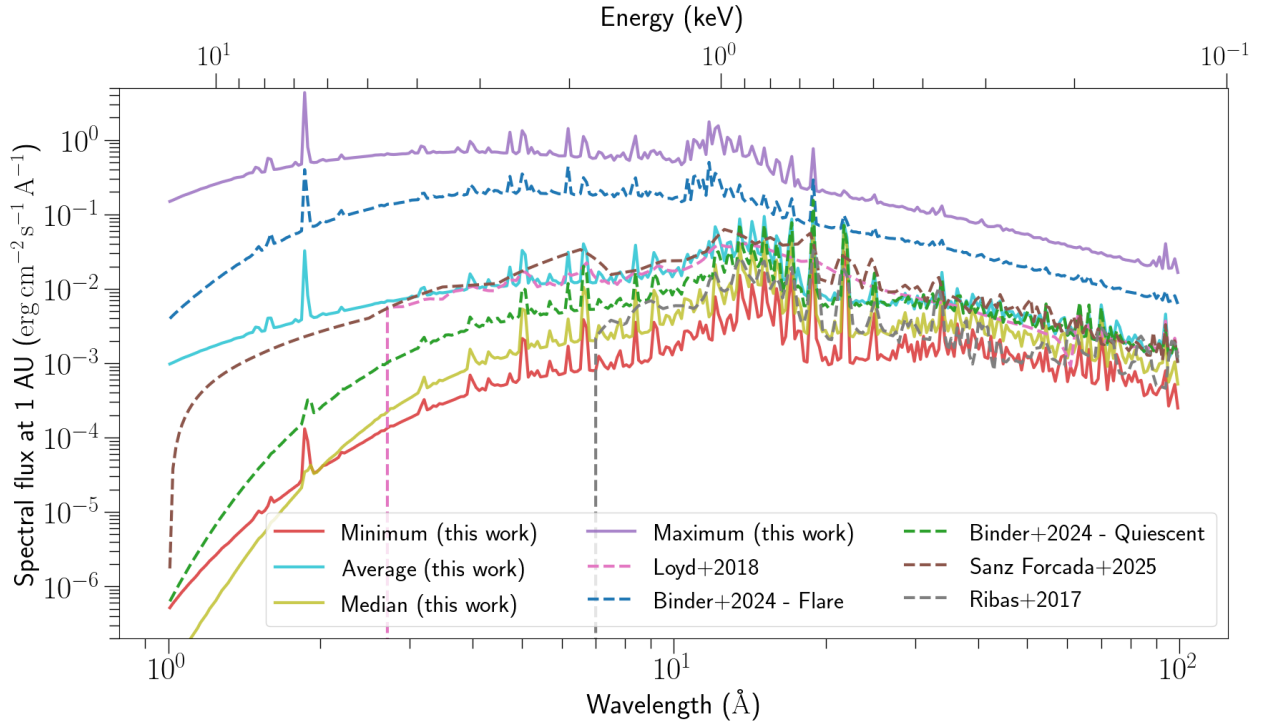


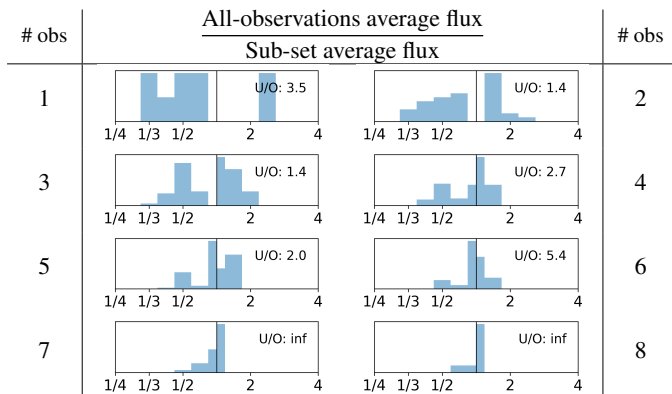
Fig. 8. Spectra fluxes at 1 AU. Results from this work (minimum, average and maximum total fluxes in the [1, 100] Å range) are compared with those from [Binder et al. \(2024\)](#) (quiescent and flaring, [1, 100] Å range); [Ribas et al. \(2017\)](#) ([7, 100] Å range); SF+25; and [Loyd et al. \(2018\)](#) ([1, 100] Å range).

Table 3. Ratios between the reference spectra described in the text and the time-averaged flux of this work.

E. range Spectrum	[1, 100] Å	[1, 2.5] Å	[2.5, 5] Å	[5, 10] Å	[10, 20] Å	[20, 50] Å	[50, 100] Å
(this work) Average (erg cm ⁻² s ⁻¹) at 1 AU	0.81 ^{+0.10} _{-0.08}	0.006 ^{+0.001} _{-0.001}	0.024 ^{+0.004} _{-0.003}	0.083 ^{+0.009} _{-0.007}	0.34 ^{+0.03} _{-0.03}	0.23 ^{+0.02} _{-0.02}	0.13 ^{+0.01} _{-0.01}
DA+25 Minimum	1/5.0 ^{+0.7} _{-0.5}	1/126 ⁺¹¹⁵ ₋₈₆	1/27 ⁺¹³ ₋₈	1/13 ⁺⁴ ₋₂	1/6.3 ^{+1.0} _{-0.6}	1/3.6 ^{+0.4} _{-0.4}	1/3.4 ^{+0.5} _{-0.3}
DA+25 Median	1/2.3 ^{+0.4} _{-0.3}	1/121 ⁺⁹⁰⁸ ₋₃₀	1/12 ⁺¹⁵ ₋₂	1/5.1 ^{+1.7} _{-0.6}	1/2.6 ^{+0.4} _{-0.4}	1/1.8 ^{+0.3} _{-0.2}	1/1.7 ^{+0.3} _{-0.2}
DA+25 80th percentile	1.5 ^{+0.2} _{-0.2}	1/18 ⁺³⁶ ₋₈	1/2.4 ⁺¹ _{-0.4}	1.0 ^{+0.2} _{-0.1}	1.5 ^{+0.2} _{-0.2}	1.6 ^{+0.2} _{-0.2}	1.6 ^{+0.2} _{-0.2}
DA+25 Maximum	20.1 ^{+2.1} _{-1.8}	123 ⁺²⁶ ₋₃₂	71 ⁺¹² ₋₁₁	39 ⁺⁴ ₋₄	17 ⁺² ₋₂	14 ⁺¹ ₋₁	11 ⁺¹ ₋₁
BI+24 Quiescent	1/1.3 ^{+0.1} _{-0.1}	1/20 ⁺⁴ ₋₅	1/3.7 ^{+0.6} _{-0.5}	1/2.2 ^{+0.2} _{-0.2}	1/1.3 ^{+0.1} _{-0.1}	1/1.1 ^{+0.1} _{-0.1}	1.0 ^{+0.1} _{-0.1}
BI+24 Elevated	1.1 ^{+0.1} _{-0.1}	1/16 ^{+3.3} _{-3.7}	1/2.6 ^{+0.4} _{-0.3}	1/1.4 ^{+0.2} _{-0.1}	1.1 ^{+0.1} _{-0.1}	1.2 ^{+0.1} _{-0.1}	1.3 ^{+0.1} _{-0.1}
BI+24 Rising	1.3 ^{+0.1} _{-0.1}	1/4.9 ^{+1.0} _{-1.1}	1/1.2 ^{+0.2} _{-0.2}	1.1 ^{+0.1} _{-0.1}	1.3 ^{+0.1} _{-0.1}	1.4 ^{+0.1} _{-0.1}	1.4 ^{+0.2} _{-0.1}
BI+24 Descending	1.5 ^{+0.2} _{-0.1}	1/15 ⁺³ ₋₃	1/1.8 ^{+0.3} _{-0.25}	1.2 ^{+0.1} _{-0.1}	1.8 ^{+0.2} _{-0.2}	1.4 ^{+0.1} _{-0.1}	1.5 ^{+0.2} _{-0.1}
BI+24 Flare	6.0 ^{+0.6} _{-0.5}	15 ⁺³ ₋₄	18 ⁺³ ₋₃	12 ⁺¹ ₋₁	4.8 ^{+0.4} _{-0.4}	4.9 ^{+0.4} _{-0.4}	4.2 ^{+0.4} _{-0.3}
SF+25 Average	1.2 ^{+0.1} _{-0.1}	1/2.1 ^{+0.4} _{-0.5}	1.0 ^{+0.2} _{-0.1}	1.3 ^{+0.1} _{-0.1}	1.2 ^{+0.1} _{-0.1}	1.3 ^{+0.1} _{-0.1}	1.0 ^{+0.1} _{-0.1}
RI+17 Average	1/2.6 ^{+0.3} _{-0.2}	–	–	1/5.6 ^{+0.6} _{-0.5} (★)	1/2.7 ^{+0.3} _{-0.2}	1/2.4 ^{+0.3} _{-0.2}	1/1.8 ^{+0.2} _{-0.1}
LO+18 MUSCLES APEC	1/1.2 ^{+0.1} _{-0.1}	–	–	1/1.1 ^{+0.2} _{-0.1}	1/1.1 ^{+0.1} _{-0.1}	1/1.1 ^{+0.1} _{-0.1}	1/1.1 ^{+0.1} _{-0.1}
LO+18 MUSCLES Multi	1/1.3 ^{+0.1} _{-0.1}	–	1/1.1 ^{+0.2} _{-0.2}	1/1.1 ^{+0.2} _{-0.1}	1/1.2 ^{+0.1} _{-0.1}	1/1.1 ^{+0.1} _{-0.1}	1/5.6 ^{+0.6} _{-0.4}

Notes. DA+25: this work; BI+24: [Binder et al. \(2024\)](#); SF+25: [Sanz-Forcada et al. \(2025\)](#); RI+17: [Ribas et al. \(2017\)](#); LO+18: [Loyd et al. \(2018\)](#). The first row shows the total flux values in the different wavelength ranges for the average. Absolute 1σ errors are propagated considering only the uncertainties of the references from this work. For example, the value $1/1.1^{+0.1}_{-0.2}$ indicates that in that band the referenced spectrum is $1.1^{+0.1}_{-0.2}$ times fainter than our time-averaged spectrum (first row, same column), whereas the value 18^{+3}_{-3} indicates that the referenced spectrum is 18^{+3}_{-3} times brighter than our time-averaged spectrum (first row, same column). (★) The reference spectrum for this value is integrated over [7, 100] Å instead of [5, 100] Å.

Table 4. Ratio between the average flux over all eight observations (the first one is split in two) and that obtained by randomly sampling a given number of observations (considering all possible combinations without repetition).



Notes. The vertical line at the center represents unity. Values to its left indicate underestimation, while values to its right indicate overestimation. The ratio of underestimating to overestimating cases is reported as U/O. Cases falling within the -1 to 1 bins are excluded from this ratio.

overestimation. Obviously, increasing the number of observations reduces the deviation from the full-sample average. We are therefore primarily interested in the first rows, which correspond to cases with fewer observations. Regardless of the number of observations, the probability of underestimating the average always exceeds that of overestimating it. Using only one to three observations, the flux can be underestimated by up to a factor of three and overestimated by a factor of two to three. This effect likely contributed to the factor of three discrepancy observed between averages reported by different authors. Variability is amplified at higher energies (wavelengths shorter than 5 \AA). Even when using half of the observations, the flux can be underestimated by up to an order of magnitude in the $[1, 2.5] \text{ \AA}$ band and by up to a factor of five in the $[2.5, 5] \text{ \AA}$ band.

4. Conclusions

The main goal of this work was to provide high time-resolution X-ray fluxes from Prox Cen that can be used to model the response of planetary atmospheres to short-term stellar variability. We analyzed all available XMM-Newton X-ray observations of Prox Cen and produced a time series of spectra covering the $1\text{--}100 \text{ \AA}$ range and spanning nearly three days of data collected over 17 years. In addition to this, an average Prox Cen EUV spectrum and the X-ray to EUV scaling law from SF+25 were used to extend the time series of spectra up to 920 \AA .

Time resolution is important in the study of exoplanets for three reasons: (1) Atmospheric responses are non-linear, so the timing of energy release can be as important as the total energy. (2) Observations of exoplanet atmospheres are instantaneous snapshots, which may occur shortly after flaring events or at different phases of the activity cycle, making knowledge of the full flux variability crucial for accurate interpretation. (3) Averaging the stellar spectrum before or after the model regression can lead to different results, although in our case this effect alters the integrated flux by less than 10%, a value comparable with the uncertainties.

Taking advantage of the high number of photon counts that can be obtained from Prox Cen due to its proximity, we

implemented a novel time-binning strategy, achieving an average time-resolution of 300 s . The high number of counts, though, also causes a non-negligible pile-up effect. To solve this, a new correction for pile-up was applied. We showed that for Prox Cen, during flares, this effect can reduce the measured fluxes by up to 30% if left uncorrected.

We tested thermal plasma models with different numbers of components, and the best-fitting model was selected for each time interval. In many cases, multiple models provided comparably good fits. We compared our fluxes with those reported in previous studies and discussed how methodological choices affect them. Finally, we examined the impact of selecting only a subset of the available observations on the inferred average flux and applied two scaling laws to reconstruct the $100\text{--}920 \text{ \AA}$ flux from the $1\text{--}100 \text{ \AA}$ range.

The intrinsic variability of Prox Cen is characterized across different wavelengths. In the $1\text{--}100 \text{ \AA}$ band, the flux varies by about two orders of magnitude between minimum and maximum, with fluctuations of factors from two to ten occurring commonly on timescales of minutes to hours. Variability is even stronger in the $[1\text{--}20] \text{ \AA}$ range, where changes span three to four orders of magnitude. Continuous variability, largely attributable to flares over a broad intensity range, precludes the precise definition of a stellar quiescent state. Even the lowest observed flux levels may correspond to phases affected by coronal dimming following previous flares and thus cannot be reliably used as a quiescent reference. We note that if a cycle with an amplitude of approximately 1.5 were confirmed, its effects on planetary atmospheres would likely be negligible compared to those of flares and short-term variability.

Over the full observational baseline, we studied the average, median, minimum, maximum, and 80th-percentile X-ray fluxes and compared them with flux estimates from previous works. We find that time-resolved analyses are essential to properly capture peaks and minima. Labeling intervals as flaring or quiescent and then averaging them, as commonly done to improve the S/N, may lead to differences of up to a few factors, which should be considered when studying the response of planetary atmospheres to transient events.

The average flux is strongly influenced by stellar variability and significantly differs from the quiescent level, unlike the case of less active stars. Specifically, the average flux can vary by up to a factor of about three when computed over different combinations of observations, implying that single or sparse observations of active targets may not be representative of the long-term average. For this reason, an additional uncertainty should be considered beyond the statistical ones, with its amplitude depending on the activity level and total observing time.

With the study of the specific case of Prox Cen, we identified different sources of uncertainties. Some of the uncertainties are common to all the sources, while some others are specific to such bright sources as Prox Cen. The data reduction choices (e.g., the time-binning strategy and the target time resolution); observational setup (e.g., instruments used, their sensitivity, and observing conditions); and assumptions on the hydrogen column density, the metallicity (see, e.g., the up to 20% differences in the flux reported in Section 2.5), and plasma models are generally common sources in most of the analyses. For bright sources, pile-up can also introduce further uncertainties.

When looking at total fluxes, these uncertainties are likely negligible compared to the reported factor of about three related to the limited observational baseline. However, they may dominate in specific situations, such as during high-flux flares or in

high-energy wavelength intervals. In these cases, uncertainties increase due to modeling limitations, as multiple components can overlap during high-flux intervals or the S/N becomes very low at high energies, making the hot plasma components poorly constrained. In addition, when not corrected, pile-up can be significant during flares.

Furthermore, using two different X-ray-to-EUV scaling relations from the same authors introduces another possible factor of about three variation in the average flux. This further emphasizes the need to account for flux differences of several factors when modeling atmospheric responses.

These results emphasize that when modeling exoplanetary atmospheres, it is important to test a range of fluxes that may be considerably broader than implied by the “standard” uncertainties alone, even for a well-studied target such as Prox Cen. Our analysis further highlights the need for long-term monitoring to accurately characterize a star’s X-ray emission and, in turn, better constrain the response of its exoplanets.

Data availability

Data products from this work are available at the CDS via <https://cdsarc.cds.unistra.fr/viz-bin/cat/J/A+A/706/A349>. Additional data will be provided upon reasonable request to andrea.damonte@inaf.it or andrea.damonte@cea.fr.

Acknowledgements. The work by A.D. has been funded through the VINCI 2023 Programme of the Université Franco-Italienne. This work was supported by the Programme Nationale de Planétologie (PNP) – Action Thématique Exosystèmes de CNRS/INSU, co-funded by CEA and CNES. I.P. acknowledges support from Bando per il Finanziamento della Ricerca Fondamentale 2024 dell’Istituto Nazionale di Astrofisica (INAF). A.Ma. and I.P. acknowledge financial contribution from the INAF grant 2023 for data analysis. G.M. acknowledges the support of the ASI-INAF agreement 2021-5-HH.0 and its addendum. Based on observations obtained with XMM-Newton, an ESA science mission with instruments and contributions directly funded by ESA Member States and NASA.

References

Amaral, L. N. R., Barnes, R., Segura, A., & Luger, R. 2022, *ApJ*, **928**, 12
 Amaral, L. N. R. d., Shkolnik, E. L., Loyd, R. O. P., & Peacock, S. 2025, *ApJ*, **985**, 100
 Anglada-Escudé, G., Amado, P. J., Barnes, J., et al. 2016, *Nature*, **536**, 437
 Arnaud, K. A. 1996, *ADASS*, **101**, 17
 Aschwanden, M. J. 1994, *Sol. Phys.*, **152**, 53
 Ayres, T. 2025, *AJ*, **169**, 281
 Ballet, J. 1999, *A&AS*, **135**, 371
 Ballet, J., Motch, C., & Batalha, R. 2024, XMM-CCF-REL-0406
 Bazot, M., Christensen-Dalsgaard, J., Gizon, L., & Benomar, O. 2016, *MNRAS*, **460**, 1254
 Bessell, M. S. 1991, *AJ*, **101**, 662
 Binder, B. A., Peacock, S., Schwieterman, E. W., et al. 2024, *ApJS*, **275**, 1
 Blind, N., Shinde, M., Dinis, I., et al. 2024, in *Adaptive Optics Systems IX*, 13097 (SPIE), 1622
 Boyajian, T. S., von Braun, K., van Belle, G., et al. 2012, *ApJ*, **757**, 112
 Burnham, K. P., & Anderson, D. R. 1998, *Model Selection and Inference* (Springer)
 Chen, H., De Luca, P., Hochman, A., & Komacek, T. D. 2025, *AJ*, **170**, 40
 Ciaravella, A., Maggio, A., & Peres, G. 1997, *A&A*, **320**, 945
 Cohen, O., Kashyap, V. L., Drake, J. J., Sokolov, I. V., & Gombosi, T. I. 2011, *AJ*, **738**, 166
 Colombo, S., Locci, D., Spinelli, R., et al. 2025, in *Handbook of Exoplanets*, eds. H. J. Deeg, & J. A. Belmonte (Springer Nature Switzerland), 1
 Currie, T., Biller, B., Lagrange, A., et al. 2023, *PP7*, **534**, 799
 Czesla, S., Lampón, M., Sanz-Forcada, J., et al. 2022, *A&A*, **657**, 26
 Dressing, C. D., & Charbonneau, D. 2013, *ApJ*, **767**, 95
 Dressing, C. D., & Charbonneau, D. 2015, *ApJ*, **807**, 45
 ESA. 2023, XMM-Newton SAS v.xmmsas_20230412_1735-21.0.0
 Favata, F., & Schmitt, J. H. M. M. 1999, *A&A*, **350**, 900
 Favata, F., & Micela, G. 2003, *Space Sci. Rev.*, **108**, 577

Favata, F., Micela, G., Orlando, S., et al. 2008, *A&A*, **490**, 1121
 Fuhrmeister, B., Lalitha, S., Poppenhaeger, K., et al. 2011, *A&A*, **534**, A133
 Fuhrmeister, B., Zisik, A., Schneider, P. C., et al. 2022, *A&A*, **663**, A119
 Gaia Collaboration (DR3). 2020, *VizieR*, **1350**, I/350
 Gaidos, E., Mann, A. W., Lépine, S., et al. 2014, *MNRAS*, **443**, 2561
 García Muñoz, A. 2007, *P&SS*, **55**, 1426
 García Muñoz, A. 2025, *A&A*, **698**, A199
 Güdel, M. 2004, *A&AR*, **12**, 71
 Güdel, M., & Nazé, Y. 2009, *A&AR*, **17**, 309
 Güdel, M., Audard, M., Briggs, K., et al. 2001, *A&A*, **365**, L336
 Güdel, M., Audard, M., Skinner, S. L., & Horvath, M. I. 2002, *ApJ*, **580**, L73
 Güdel, M., Audard, M., Reale, F., Skinner, S. L., & Linsky, J. L. 2004, *A&A*, **416**, 713
 Haisch, B. M., Harnden, F. R., Seward, F. D., et al. 1980, *ApJ*, **242**, L99
 Haisch, B., Antunes, A., & Schmitt, J. H. M. M. 1995, *Science*, **268**, 1327
 Hazra, G., Vidotto, A. A., Carolan, S., Villarreal D’Angelo, C., & Ó Fionnagáin, D. 2025, *MNRAS*, **536**, 1089
 Howard, W. S., MacGregor, M. A., Osten, R., et al. 2022, *ApJ*, **938**, 103
 Jethwa, P., Saxton, R., Guainazzi, M., Rodriguez-Pascual, P., & Stuhlinger, M. 2015, *A&A*, **581**, A104
 Johnstone, C. P., Bartel, M., & Güdel, M. 2021, *A&A*, **649**, A96
 Kasting, J. F., Whitmire, D. P., & Reynolds, R. T. 1993, *Icarus*, **101**, 108
 Khodachenko, M. L., Sasunov, Y., Arkhypov, O. V., et al. 2014, *IAU Symp.*, **300**, 335
 Kipping, D. M., Cameron, C., Hartman, J. D., et al. 2017, *AJ*, **153**, 93
 Kiraga, M., & Stepien, K. 2007, *Acta Astron.*, **57**, 149
 Konings, T., Baeyens, R., & Decin, L. 2022, *A&A*, **667**, A15
 Kowalski, A. F. 2024, *Liv. Rev. Sol. Phys.*, **21**, 1
 Kuntz, K. D., & Snowden, S. L. 2008, *A&A*, **478**, 575
 Lee, Y., Dong, C., Pawlowski, D., et al. 2018, *Geophys. Res. Lett.*, **45**, 6814
 Linsky, J. L., Fontenla, J., & France, K. 2014, *ApJ*, **780**, 61
 Loyd, R. O. P., France, K., Youngblood, A., et al. 2018, *ApJ*, **867**, 71
 Maggio, A., Pillitteri, I., Argiroffi, C., et al. 2024, *A&A*, **690**, A383
 Maldonado, J., Micela, G., Baratella, M., et al. 2020, *A&A*, **644**, 23
 Mascareño, A. S., Artigau, E., Mignon, L., et al. 2025, *A&A*, **700**, A11
 Mason, K. O., Breeveld, A., Much, R., et al. 2001, *A&A*, **365**, L36
 Molendi, S., Bassini, V., & Sembay, S. 2003, XMM-SOC-CAL-TN-0036
 Neupert, W. M. 1968, *ApJ*, **153**, L59
 Osten, R. A., Godet, O., Drake, S., et al. 2010, *ApJ*, **721**, 785
 Pandey, J. C., & Singh, K. P. 2008, *MNRAS*, **387**, 1627
 Pavlenko, Y. V., Suárez Mascareño, A., Zapatero Osorio, M. R., et al. 2019, *A&A*, **626**, A111
 Peres, G., Orlando, S., Reale, F., Rosner, R., & Hudson, H. 2000, *ApJ*, **528**, 537
 Pillitteri, I., Argiroffi, C., Maggio, A., et al. 2022, *A&A*, **666**, A198
 Pizzolato, N., Maggio, A., Micela, G., Sciortino, S., & Ventura, P. 2003, *A&A*, **397**, 147
 Preibisch, T., & Feigelson, E. D. 2005, *ApJS*, **160**, 390
 Reale, F., Güdel, M., Peres, G., & Audard, M. 2004, *A&A*, **416**, 733
 Reames, D. V. 2021, *Lect. Notes Phys.*, **978**, Solar Energetic Particles: A Modern Primer on Understanding Sources, Acceleration and Propagation (Springer International Publishing)
 Ribas, I., Gregg, M. D., Boyajian, T. S., & Bolmont, E. 2017, *A&A*, **603**, 12
 Ridgway, R. J., Zamyatina, M., Mayne, N. J., et al. 2023, *MNRAS*, **518**, 2472
 Robrade, J., Schmitt, J. H. M. M., & Favata, F. 2012, *A&A*, **543**, A84
 Rockcliffe, K. E., Newton, E. R., Youngblood, A., et al. 2025, *AJ*, **169**, 321
 Sanz-Forcada, J., Micela, G., Ribas, I., et al. 2011, *A&A*, **532**, A6
 Sanz-Forcada, J., López-Puertas, M., Lampón, M., et al. 2025, *A&A*, **693**, A285
 Schartel, N., González-Riestra, R., Kretschmar, P., et al. 2024, in *Handbook of X-ray and Gamma-ray Astrophysics*, eds. C. Bambi, & A. Santangelo (Springer Nature), 1501
 Schellenberger, G., Kraft, R., Nulsen, P., et al. 2025, *JATIS*, **11**, 018006
 Segura, A., Walkowicz, L. M., Meadows, V., Kasting, J., & Hawley, S. 2010, *Astrobiology*, **10**, 751
 Smith, R. K., Brickhouse, N. S., Liedahl, D. A., & Raymond, J. C. 2001, *ApJ*, **556**, L91
 SOC, E. X.-N. 2023, *Users Guide to the XMM-Newton Science Analysis System*, issue 18.0
 Strüder, L., Briel, U., Dennerl, K., et al. 2001, *A&A*, **365**, L18
 Turner, M. J. L., Abbey, A., Arnaud, M., et al. 2001, *A&A*, **365**, L27
 Vidotto, A. A. 2025, *ARA&A*, **63**, 299
 Wargelin, B. J., Saar, S. H., Irving, Z. A., et al. 2024, *ApJ*, **977**, 144
 Wargelin, B. J., Saar, S. H., Pojmański, G., Drake, J. J., & Kashyap, V. L. 2017, *MNRAS*, **464**, 3281
 Webb, D. F., & Howard, T. A. 2012, *Liv. Rev. Sol. Phys.*, **9**, 3
 Wilms, J., Allen, A., & McCray, R. 2000, *ApJ*, **542**, 914
 Yan, F., & Henning, T. 2018, *Nat. Astron.*, **2**, 714
 Yoshiro, S., & Gopalswamy, N. 2009, *IAU Symp.*, **257**, 233
 Yelle, R. V. 2004, *Icarus*, **170**, 167

Appendix A: Dealing with pile-up

Pile-up is an instrumental effect in CCD photon-counting detectors that reduces the measured counts and hardens the spectrum. It distorts spectral shapes and can compromise model fitting.

When a photon hits a CCD photon-counting detector, it generates an electronic cloud within a pixel. This cloud often spreads in adjacent pixels, and the resulting pixel configuration is called “pattern” in XMM-Newton nomenclature (or also “grade” in other instruments, e.g., Chandra). Higher-energy photons tends to produce broader charge distributions, resulting in larger pattern values. Spurious events, such as those from cosmic rays or charged particles, usually yield the highest pattern values and are filtered automatically by the onboard software. Valid patterns (e.g., 0–12 for pn, 0–4 for MOS) correspond to events most likely caused by X-rays. These are often grouped into singles (pattern = 0) and doubles or greater ($0 \leq \text{pattern} \leq 12$ or 4).

When two or more photons hit the same or neighboring pixels within a single readout frame, they are recorded as a single, more energetic event, a phenomenon known as pile-up. This effect lowers the total count rate, distorts the spectral energy distribution, and shifts the pattern distribution toward higher values. It may also produce non-standard patterns (e.g., diagonal ones). For point sources, most photons concentrate in the core of the (PSF), making the pile-up rate dependent not only on source brightness but also on radial distance from the PSF center. Many high-pattern and high-energy events are flagged as spurious and excluded by onboard filtering. This leads to additional count loss beyond that strictly caused by pile-up itself, often visible as a depletion in the core of the PSF. Both PSF and pattern distribution are strongly energy-dependent, and complex corrections are applied to take this into account. Therefore, pile-up distortion of both the spectral and pattern distributions introduces complex and non-linear effects. Because of these complications, there is currently no universal method to correct for pile-up. Different correction strategies are described in the literature.

1. PSF core excising: When the pile-up fraction is high, the most effective strategy is to excise the core of the PSF. This works because the count rate per frame per pixel is higher at the core, whereas the wings of the PSF are usually unaffected. After excluding the core, the total flux can be reconstructed using knowledge of the PSF shape using the wings, where the pile-up fraction is negligible. The method can be applied even when core depletion is not visually apparent, although its accuracy depends on the quality of the PSF calibration and the number of events available in the wings. Two main challenges arise: (a) determining the appropriate excision radius, and (b) reconstructing the excluded counts to recover the total flux. There is no universal criterion for selecting the excision radius, and during periods of high variability it may be necessary to vary it over time.
2. Single-pattern filtering: Single-pixel events are much less affected by pile-up for geometric and probabilistic reasons. A common strategy is therefore to retain only single-pixel patterns⁸. In general, the probability that single-pixel events are affected by pile-up is negligible (Ballet 1999; Jethwa et al. 2015). As with the PSF excision method, calibration files are needed to correctly generate a response matrix

⁸ Using single-pixel events also helps reject soft proton contamination that may have been missed by onboard filtering, particularly in pn small window mode (as in some of our observations), since such invalid events rarely appear as singles but may appear as doubles (Schellenberger et al. 2025).

that takes into account the single-pixel event filtering. However, this correction is automatically applied by SAS without accounting for the pattern migration caused by pile-up, which reduces the fraction of single-pixel events (Ballet 1999). Because the standard pattern calibration assumes a fixed singles-to-doubles ratio, it does not correct for this shift, leading to an underestimation of the flux in already pile-up-affected data. Consequently, this method corrects spectral distortions but does not recover the true flux. This limitation is not always explicitly stated. Even without pile-up, the single-pattern filtering method inherently biases against high-energy photons, which are more likely to produce broader charge distributions and higher pattern values. Although this effect is accounted for in the calibration, the filtering still reduces the photon statistics at high energies.

3. SAS response matrix correction: XMM-Newton guidelines recommend a correction (implemented in SAS) that incorporates pile-up effects into the response matrix. However, this method is implemented only for the pn instrument. Moreover, when performing time-resolved analyses, the response matrix should be recomputed for each time interval individually, which would be extremely time-consuming, and we suspect it would also be affected by the low statistics. For this reasons, we do not consider this option.

We tested method 1), applying a variable excision radius determined from the count rate and using the *epatplot* output as a reference. However, we found that this correction introduced uncertainties larger than those caused by pile-up itself, likely because of the low number of counts in the PSF wings during most of the affected intervals. Consequently, we discarded this approach.

We adopted method 2), developing a new method to correct the effects of pattern migration. Our correction builds on the relations presented in Jethwa et al. (2015), linking count rate, flux loss, and spectral distortion for XMM-Newton (see their Figure 5). This represents an original adaptation of their framework. The flux loss, F , is defined as

$$F = 1 - \frac{M}{\Lambda}, \quad (\text{A.1})$$

where Λ is the expected number of events and M is the measured number of events. We can invert the relation to obtain

$$\Lambda = \frac{M}{1 - F}. \quad (\text{A.2})$$

Then the number of event counts lost due to pile-up will be P :

$$P = \frac{MF}{1 - F}. \quad (\text{A.3})$$

We also know the expected ratio between different pattern events as a function of their energies. For example, from the calibration files, we can extract the fraction of expected single-pattern events as a function of energy, $\alpha(E)$. Assuming that the number of piled-up singles is negligible, the expected number of single-pixel events can be written as

$$\Lambda_s(E) = \alpha(E)\Lambda(E). \quad (\text{A.4})$$

We do not know $\Lambda(E)$ but we know the spectrum of the measured singles $M_s(E)$, and because we assume that singles are not pile-up, when the number of events is reasonably high, we can assume

$$\Lambda_s(E) = \text{constant} \cdot M_s(E). \quad (\text{A.5})$$

If we do so, we can write

$$\int \Lambda(E)dE = \Lambda = \int \text{constant} \cdot \alpha^{-1}(E)M_s(E)dE \quad (\text{A.6})$$

and solve for the constant. This is also suggested in a different way in Ballet (1999). Given the complexity of this approach, before testing it, we deem more reasonable to test the extreme limit in which we assume

$$\Lambda_s - M_s = \Lambda - M = P. \quad (\text{A.7})$$

That is, we consider the extreme case in which all the lost counts are considered as single counts. We do not know the impact of this correction on the results and it could be negligible compared to our error, only introducing more unpredictable errors, so we evaluated the worst case scenario.

We applied this correction directly on the raw data in the following way:

- we divide the observation into 30 seconds time intervals, then we proceed for each interval;
- we use data from Figure 5 of Jethwa et al. (2015) and equation A.3 to estimate the number of missing events due to pile-up;
- we construct a cumulative distribution from $M_s(E)$, the measured singles, and we extract P events energy randomly;
- we assign a random time within the interval to the events energies and we get a set of events;
- we add those events to the original event list to obtain a half-synthetic event list corrected for single pattern events pile-up loss.

A limitation of this correction, beyond it representing an upper limit, is that single-pattern filtering typically excludes the most energetic photons, which are generally doubles or greater patterns. Although this effect should be accounted for by SAS in the calibration steps, it nonetheless reduces the number of high-energy counts and therefore their statistical significance and our ability to constraint the plasma higher-temperature components. Moreover, since pile-up occurs primarily during flares, which produce harder spectra, this effect may particularly increase uncertainties at the high-energy end of the reconstructed spectra during such events when filtering for single-pattern events.

We ended up with three different source event lists for each observation and each instrument:

1. NOCORR: All valid-pattern events are included; no correction is applied. This list suffers from pile-up during bright intervals and serves as a lower bound as well as the uncorrected reference.
2. STDCORR: Only single-pattern events are used. This is the simplest correction, designed to correct spectral distortion only, though it is sometimes used as a general correction. Results from this list help estimate flux loss due to pattern migration.
3. NEWCORR: Our new method, which applies single-pattern filtering combined with generated events. Results from this list represent an upper limit because we are assuming that all the events affected by pile-up would have resulted in single pattern events.

In Figure A.1, we show the estimated count-loss percentage for the different instruments during the two major events in our sample (observations 0049350101 and 0551120401). Pile-up count loss reaches up to 22% for the pn. This occurs during the less intense of the two events because of different frame

mode used by XMM-Newton during observation 0049350101. It directly reflects the correction dependence on the frame time coming from the relations of Jethwa et al. (2015). In contrast, MOS1 and MOS2 were operated with the same setup throughout all observations and can therefore serve as a reference. For these instruments, pile-up reach its maximum (13%) during the greatest flare (observation 0049350101).

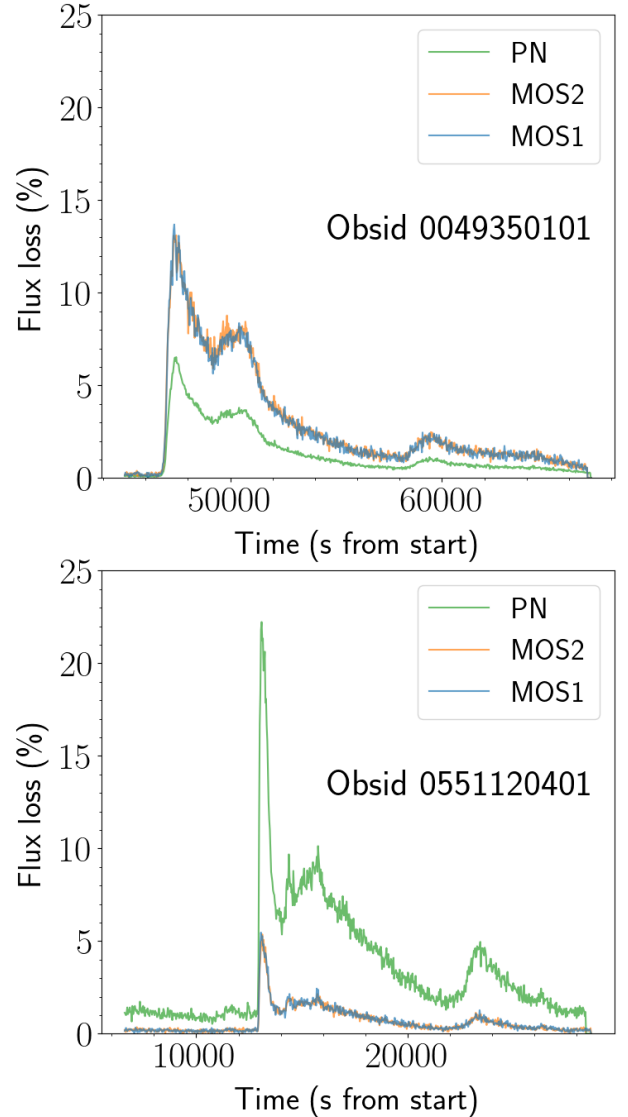


Fig. A.1: Comparison between the pile-up flux-loss inferred for the different instruments during observations 0049350101 and 0551120401. The flux loss is defined as in Jethwa et al. (2015).

Next, we examine how the subsequent steps in the reduction process affect the total flux, defining the effects of pile-up at the end of the analysis chain as a function of time and flux. We adopted the three-temperature model as a reference. We use uppercase abbreviations to refer to the total fluxes derived from each correction method. As a metric for pile-up-induced flux loss, we define the quantity $1 - \frac{\text{NOCORR}}{\text{XXXCORR}}$, where XXX = STD or NEW. In Figures A.2 and A.3, we plot this quantity for the STDCORR and NEWCORR cases, respectively.

Figure A.2 shows that, when comparing fluxes computed using all valid patterns (NOCORR) with those using only

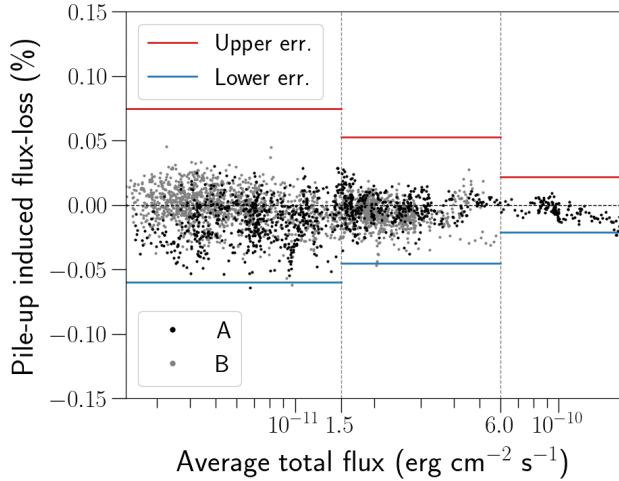


Fig. A.2: Comparison between fluxes obtained without correcting for pile-up and correcting by selecting only single patterns. Pile-up induced flux-loss defined as $(1 - \text{NOCORR}/\text{STDCORR})$ is plotted with black and gray dots over the total flux. Black dots represent intervals of the observation with id 0049350101, which was observing with pn in the small-frame mode. The gray dots are used for all the other observations observed with the standard frame time. Relative upper and lower total flux 1σ errors from the reduction procedure are depicted as red and blue segments for the same flux ranges as in Figure 1.

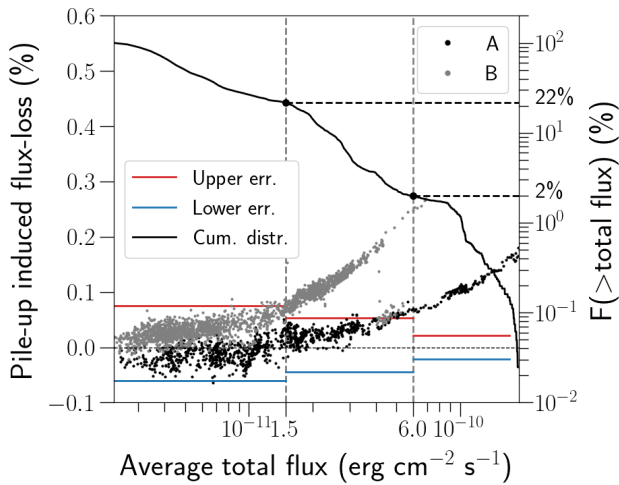


Fig. A.3: Same as Figure A.2, but here the flux loss is defined as $(1 - \text{NOCORR}/\text{NEWCORR})$ using the fluxes obtained with the new correction method. The right y-axis shows the inverse cumulative distribution of the total fluxes.

single-pixel events (STDCORR), the differences induced by pile-up remain well within the flux uncertainties. This indicates that STDCORR does not effectively correct for flux loss: The retrieved closely matches the NOCORR one, implying that applying the correction has negligible impact on the total flux.

Figure A.3 clearly shows flux loss when using the NEWCORR correction. Two distinct populations of time intervals, labeled (A) and (B) in black and gray, correspond to different frame modes of the pn detector. This reflects both the frame-mode-dependent nature of the correction model and the dominant contribution of pn relative to MOS1 and MOS2 in the total signal.




Fig. B.1: Different models and AIC-selected flux density averages over various total-flux-percentiles are compared. The color bar represents the percentile interval.

When the pile-up fraction is low, the correction may introduce uncertainties larger than the flux loss itself. Moreover, at low flux levels, pile-up-induced losses are generally smaller than the statistical uncertainties. For this reason, we identify the time intervals for which the correction is meaningful and omit it when unnecessary.

Our strategy is as follows: We apply the pile-up correction only when the estimated flux loss exceeds the total flux error. This allowed us to preserve the higher S/N of the full-pattern dataset in low-count intervals. We find that the pile-up effect can be neglected in the low-flux, eye-selected range, where it is consistently smaller than the average instrumental error and approximately symmetric around zero. For intervals in this range, we use the NOCORR reduction. For total fluxes above $1.5 \times 10^{-11} \text{ erg s}^{-1} \text{ cm}^{-2}$, corresponding to roughly 22% of the observing time, we apply the correction.

In principle, separate thresholds could be defined for each detector configuration (i.e., combinations of pn, MOS1, and MOS2 modes). However, based on our tests, a single threshold across all configurations provides sufficiently accurate correction for our purposes.

Appendix B: Further comments on model selection

In Figure B.1, we compare results obtained by applying a single model to all time intervals versus selecting the best model for each time bin. To show the wide range covered and the frequency of occurrence, we divide the spectra according to total flux percentiles. For each percentile interval we plot the average spectrum. For wavelengths shorter than $\sim 5 \text{ \AA}$, different models yield spectra that differ by many orders of magnitude. This is due to the low signal we have at those frequencies and to their sensitivity to the high temperatures, which changes in different models. Fluxes at these energies are better constrained only during flares.

Figure B.2 shows the AIC-selected spectra only. Averages are plotted with continuous lines. Shaded areas, delimited by dotted lines, represent the extremes reached in each energy for the specific percentile interval. To avoid confusion, we only plot three of the five percentile intervals, but similar behaviors are




Fig. B.2: Information criteria-selected flux density averages over various total-flux-percentiles (plotted with a continuous line). The shaded area delimited by dotted lines represents maximum and minimum values reached in that percentile interval at that wavelength. For clarity, only three percentile intervals are reported.

found also for the others. We note that the red set reaches higher flux density values than the purple one in the short-wavelength bins. This highlights the wide variety of spectral shapes, and shows that a higher total flux does not necessarily correspond to a higher high-energy partial flux. We warn against interpreting dotted lines as spectra. They are not, but they represent extreme values reached independently in each wavelength bin.