Flare and Back Again:
The Analysis of High Energy Emission from Small Solar Flares

Kristopher Cooper
MSci (Hons)

Submitted in fulfilment of the requirements for the degree of Doctor of Philosophy

Astronomy and Astrophysics Group
School of Physics and Astronomy
University of Glasgow

December 2022

This thesis is my own composition except where indicated in the text. No part of this thesis has been submitted elsewhere for any other degree or qualification.

9th December 2022

All that is gold does not glitter,
Not all those who wander are lost;
The old that is strong does not wither,
Deep roots are not reached by the frost.

From the ashes a fire shall be woken,
A light from the shadows shall spring;
Renewed shall be blade that was broken,
The crownless again shall be king.

The Lord of the Rings: The Fellowship of the Ring
— J.R.R. Tolkien

Abstract

The solar atmosphere has a peculiar property where its outer layer, the corona, is hotter than its base, the photosphere; this issue is termed the *coronal heating problem*. One possible solution is that small solar flares—*microflares* and *nanoflares*—occur in such a frequency that they produce more net energy than their larger counterparts, that are insufficient in themselves, to heat the corona. To investigate this premise, microflare extreme ultraviolet (EUV) and X-ray emission must be studied to determine their production mechanisms and energetics.

In this thesis, we focus on the observations of, and analytical tools for, the study of high energy microflare emission with particular emphasis on X-ray spectroscopy. In Chapter 1, we introduce the necessary context for microflares with respect to the solar atmosphere as well as a description of their high energy EUV and X-ray emission and how this can be modelled mathematically. Chapter 2 provides an overview of the hard X-ray, soft X-ray, and EUV instruments used throughout this thesis with specific attention given to the Nuclear Spectroscopic Telescope ARray (NuSTAR).

NuSTAR is an astrophysical X-ray telescope capable of observing the Sun with direct imaging spectroscopy providing a unique sensitivity $>$2.5 keV. In Chapter 3, we investigate ten NuSTAR microflares that originated from active region AR12721 and occurred between 2018 September 9–10. The ten microflares are all weak sub-A GOES class and still reach temperatures up to $\sim$10 MK. One microflare shows direct evidence for non-thermal emission and eight of the ten show indications of photospheric magnetic flux cancellation in proximity to their footpoints. Chapter 4 then investigates the weakest microflare from Chapter 3 further, finding it to be the weakest active region X-ray microflare.

Chapter 5 provides an overview of X-ray spectral fitting approaches where the nuances of X-ray spectral data analysis, utilising different fitting and statistical methods, are discussed. We then describe two widely used X-ray spectral fitting programs, OSPEX and XSPEC, detailing their advantages and limitations before introducing a new Python X-ray spectral fitting tool called Sunxspex. Sunxspex is optimised for solar data products and aims to combine the capabilities of OSPEX and XSPEC into one program.

Using examples from NuSTAR, the Reuven Ramaty High-Energy Solar Spectroscopic Imager (RHESSI), and Solar Orbiter's Spectrometer/Telescope for Imaging X-rays (STIX), we showcase the abilities

of Sunxspex in Chapter 6. We re-analyse data from existing NuSTAR and RHESSI studies, finding good agreement between fitted model parameter values, while executing new analysis with the same data and STIX data that was not possible with software like OSPEX and XSPEC. Chapter 7 describes the nested sampling algorithm which is made available in Sunxspex and allows for quantitative comparisons between different models being fitted to the same data.

Finally, Chapter 8 presents the analysis of five microflares and a jet observed by NuSTAR. The first two microflares are observed during the 2020 January solar observation campaign while the last three microflares and the jet are from the 2021 November campaign. The microflares show evidence of temperatures $>$10 MK being present with several of the microflares showing potential direct evidence of non-thermal emission. The jet reaches quiescent, non-flaring active region temperatures of $\sim$4 MK while being located far from the closest active region.

Contents

List of Tables

List of Figures

Acknowledgements

First and foremost, I would like to thank my PhD supervisor, Iain Hannah, for providing me the opportunities I have had throughout the last four years, always being constructive and honest in our discussions and always pointing me in the right direction for my research. I can only look forward to more collaborations between the two of us.

There was a time where I was unaware of my affinity to solar physics and was, believe it or not, considering a PhD in another other field; my Master's project five years ago with Lyndsay Fletcher soon saw to that. As my second supervisor, you have been insightful and encouraging at every big step during my PhD. Thank you for setting me off on an amazing journey.

Thank you to the NuSTAR solar team for the science discussions in our bi-weekly meetings, even if they are at 5 pm on a Monday. Being able to present and discuss my own work as well as listen to everyone else's research has been enjoyable and invaluable.

I would also like to say thank you to my internal examiner, external examiner, and convenor who were John Veitch, Patrick Antolin, and Nicolas Labrosse, respectively. You made my viva a very relaxing and enjoyable experience, it was a pleasure to hear your comments and have those discussions which improved this thesis and has left me with no shortage of ideas on what could be done next.

To everyone in Room 604, thank you. I think the office is only just getting back to normal since the pandemic, over which I missed all the antics that being in the office brought. It is a place I will miss greatly because of the laughs and fun we all had. I would also like to thank everyone along the astronomy corridor for the really interesting, sometimes completely off-topic, conversations at coffee or in the office.

I would not have been able to accomplish what I have without the help and support from my friends, particularly George and Lisa who read over some of the material in this thesis which was no small feat. I owe you all a big thank you, this journey would not have been the same without any of you.

I must thank my family for always being there for me not just throughout my academic expedition but also in life. Thank you, Mum (and of course Tilly as well). If not for our early morning walks and you kindly listening to whatever science issue I had that week I would not have been able to keep myself

motivated. Gran & Grandad, I think it is safe to say that I could not have done any of this without your guidance, support, and understanding, thank you. I would also like to thank my Dad who has helped and been there at every step of my academic journey. Another thank you has to go to Andy and Moira for always providing me support and encouragement for which I am incredibly grateful.

Last, but by no means least, to my wife, Lindsey. Literally none of this would have been possible without your help, support, friendship, and love which has never wavered since we met. I would not be the person I am today, capable of doing any of this, without you. There is no other person I would have rather shared this adventure with and I cannot wait to see where life takes us next. To the next adventure, to the future!

1

Introduction

The Sun has been worshipped since the dawn of civilisation and has been scientifically investigated over centuries (see Chapter 1 of Judge 2020). Today, many solar concepts and processes that have been opaque are now well understood. Solar physics employs a vast number of physical domains. The construction and testing of instrumentation, the thermal properties and general behaviour of plasma, the diagnostics of spectral line and continuum emission, are but a few examples of where various areas of physics must work harmoniously to constantly improve and investigate our understanding of the Sun.

In this chapter, we introduce the structure of the solar atmosphere in Section 1.1 before detailing the solar flare process and their environment in Section 1.2. We then focus on the weakest of X-ray flares, *microflares*, in Section 1.3 and give an overview of microflares studied by the Nuclear Spectroscopic Telescope ARray (NuSTAR; Harrison et al. 2013, see Section 2.1) before discussing important concepts relating to the analysis of flare extreme-ultraviolet (EUV) and X-ray emission (Section 1.4).

1.1 The Solar Atmosphere

The Sun's atmosphere is studied extensively, along with many different types of events that occur within it. Much like the Earth, the Sun's atmosphere is split into layers characterised by their temperature and density properties. These layers can be seen clearly in different ultraviolet and extreme-ultraviolet wavelengths (Figure 1.1). These layers are called the *photosphere*, the *chromosphere*, and the *corona* (Carroll & Ostlie 2014, pp. 332–352).

The temperature at the centre of the Sun (10^7 K) drops to ~6,000 K at the photosphere, the visible surface of the Sun (Figure 1.1, top left panel). The temperature continues to decrease into the chromosphere (top right panel) to ~4,400 K—the chromospheric temperature minimum—before rising to 10^5 K over a distance of ~2,000 km. Beyond the chromosphere, into the corona (bottom

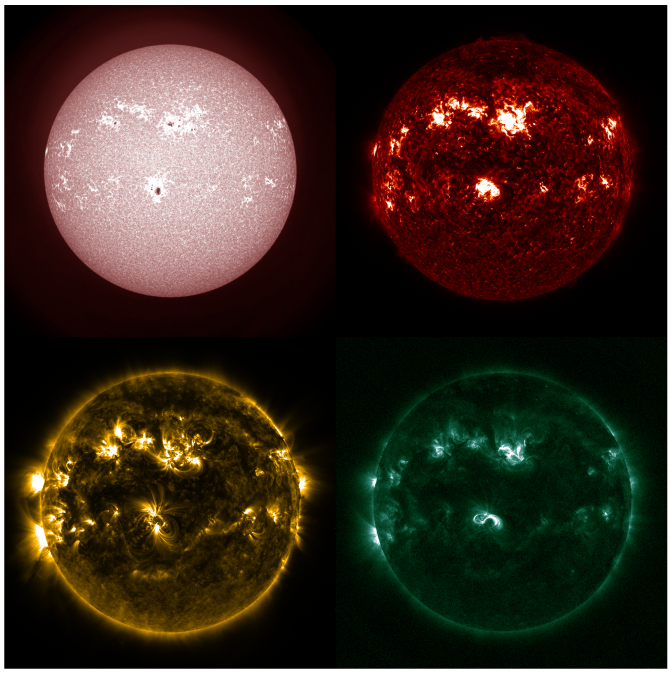

Figure 1.1: Layers of the solar atmosphere as seen in EUV by the Solar Dynamics Observatory's Atmospheric Imaging Assembly (Lemen et al. 2012, see Section 2.5.1). The solar photosphere is seen in the 1700 Å channel (top left panel), the chromosphere in the 304 Å channel (top right panel), with the quiet and flaring corona seen in the 171 Å and 94 Å channels (bottom left and right panels), respectively.

panels), the temperature reaches 10^6 K. The ~100 km region encapsulating the drastic transition between the chromosphere and the corona is commonly referred to as the transition region. The atmospheric density also changes greatly, decreasing by over seven orders of magnitude from the photosphere to the corona (Carroll & Ostlie 2014, pp. 332–341). A large motivation for much solar physics research is to better understand this drastic temperature increase from the photosphere to the corona and is termed the *coronal heating problem*.

There is a plethora of differently structured and evolving phenomena that have been observed about and above the visible surface of the Sun. One observable that has garnered much attention is the solar active region (AR). ARs are associated with strong magnetic fields, are related to sunspots, and are the location of solar flares. As these regions have a higher magnetic flux, their temperature becomes relatively cooler compared to the surrounding plasma, reaching down to ~4,000 K in some cases resulting in sunspots (Carroll & Ostlie 2014, pp. 353–357). Sunspots can be seen at optical wavelengths as dark spots—an umbra surrounded with filament structures termed the penumbra—on the Sun's photosphere, even without the aid of a telescope (Smith & Smith 1963, pp. 12–13, and Mossman 1989). Three separate sunspot groups can be seen in the top left panel of Figure 1.1. The strong magnetic field causes bright plasma-filled magnetic loops to be visible above the AR which are easily viewed

in the chromospheric and coronal channels (Figure 1.1, top right and bottom panels). The magnetic field strength can be measured through the Zeeman effect (Zirin 1966, pp. 105–106). This is because the local magnetic field splits a single atomic energy level into several, the extent of this splitting is indicative of the magnetic field strength. These conditions appear to help breed locations capable of producing multiple solar flares of a wide range of magnitudes in the corona.

To understand the significance of the magnetic field in relation to solar flares in the corona, we make use of a convenient expression called the plasma-β given by the ratio of the gas p_{gas} and magnetic pressure p_{mag} as

$$\beta = \frac{p_{gas}}{p_{mag}} = \frac{8\pi N k_B T}{B^2}, \tag{1.1}$$

where N is the total particle number density, k_B is the Boltzmann constant, T is the isothermal plasma temperature, and B is the magnetic field strength (Phillips et al. 2012, pp. 22). When $\beta > 1$, the convective motions of the surrounding plasma dominate each magnetic element; however, when $\beta < 1$, the magnetic field is dominant. We generally find in the photosphere $\beta > 1$ but $\beta < 1$ in the corona (Tandberg-Hanssen & Emslie 2009, pp. 137); therefore, the magnetic field plays a significant role for plasma motion in the corona which is where solar flares are observed.

1.2 Solar Flares

The first solar flare to be scientifically documented was observed independently by Richard Carrington and Richard Hogdson in 1859 (Carrington 1859; Stewart 1861). The *Carrington flare* is associated with an intense geomagnetic storm which disrupted telegraph lines for hours in several countries with the aurorae being visible far from the poles for several nights, reaching as far as the equator (Moreno Cárdenas et al. 2016). With the advent of modern instrumentation, solar flares of a weaker scale are now observed regularly.

Solar flares are often described as impulsive events producing increased emission across a wide range of wavelengths and are predominantly observed to occur in active regions. The flaring process is thought to produce various spatial and temporal emission signatures throughout its evolution. The driving force behind this impulsive release of magnetic energy is thought to be through magnetic reconnection of the coronal loop magnetic field (Figure 1.2). The flaring process then rapidly converts the free magnetic energy into particle acceleration, mass flows, and plasma heating (Benz 2017).

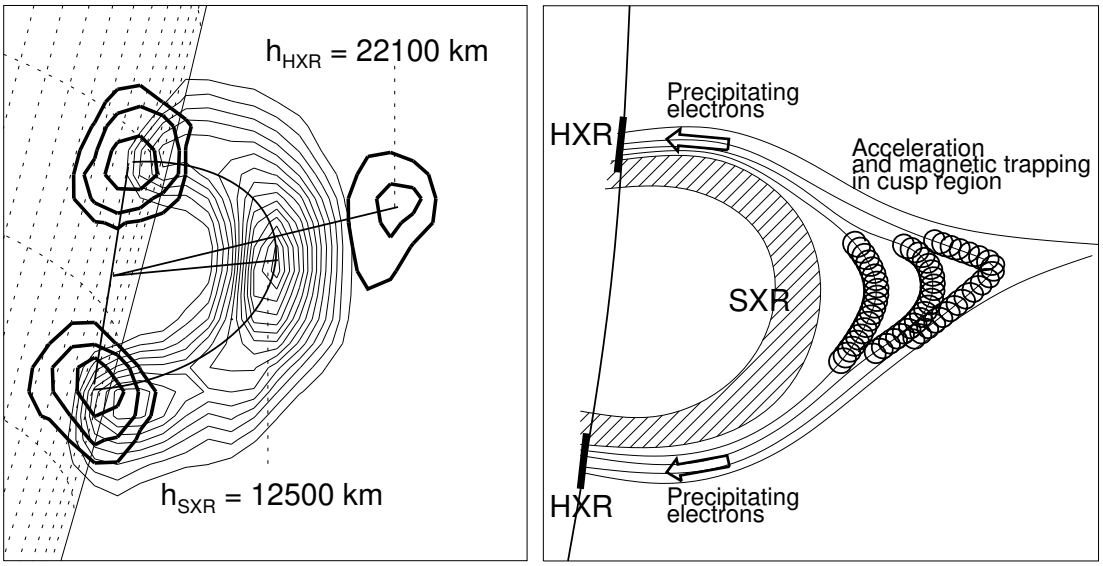

Figure 1.2: A diagram describing reconnecting magnetic field lines (right panel) with corresponding geometry to an observed flare (left panel). The hard X-ray (HXR) footpoints are produced by decelerating electrons in the chromosphere, heating the plasma that fills the loop and emits soft X-rays (SXRs). As reconnection occurs in successive field lines, the footpoints appear to separate. This figure is taken from Aschwanden (2004).

1.2.1 Flare Evolution

In a 'standard' flare, as a coronal loop stretches it becomes possible for the structure to impulsively transform into a more preferable lower energy state, whereby the oppositely directed magnetic fields from each side of the loop reconnect instead of extending out into space. Electrons in the reconnection site are accelerated down the loops anchored in the chromosphere producing microwaves and other non-thermal emission (Priest & Forbes 2002). The non-thermal electrons then smash into the denser chromosphere producing emission, via non-thermal bremsstrahlung (Section 1.4.5), in the form of hard X-rays (HXRs) and heating the footpoint plasma producing a thermal signature (Section 1.4.4) in soft X-rays (SXRs)[1]. This process happens consecutively during the reconnection process producing, what appears to be, physically moving footpoints (Aschwanden 2004). In addition to electrons being accelerated from the reconnection site and propagating down to the chromosphere, it has been suggested that energy is also transported down to the chromosphere via transverse waves in the magnetic field (Alfvén waves) which accelerates the electrons locally at the footpoints (Emslie & Sturrock 1982; Fletcher & Hudson 2008).

The heated chromospheric plasma at the footpoints *evaporates*, filling the loop, and emits thermal

[1] SXRs and HXRs are usually defined to be photon energies between 0.1–10 keV and >10 keV, respectively (Fletcher et al. 2011).

SXRs. The plasma upflows have been found to be supersonic with speeds $>200\,\mathrm{kms}^{-1}$ in EUV and SXR (Zarro & Lemen 1988; Milligan et al. 2006). In addition, it has been observed that SXR and HXR emission are related such that the SXR emission roughly follows the integral of HXR emission with respect to time (the Neupart Effect; Neupert 1968). More generally, the solar flare process is split into three phases: the pre-flare phase, the impulsive phase, and the decay/gradual phase (Fletcher et al. 2011; Benz 2017). Although it can be difficult to precisely break down all flares into the same steps due to their orientation to the observer or the number of flaring loops, we present three general flare phases often observed and indicate the physical processes thought to be occurring during them.

The Pre-flare Phase

During the pre-flare stage, the ultraviolet (UV) and SXR emission starts to increase and the loop starts to rise (Priest & Forbes 2002). The pre-flare phase can be the most difficult to observe of the three stages as it is easily hidden by other activity (Fletcher et al. 2011). This is especially true for instruments that observe the Sun as a star by integrating over the full solar disk like the Geostationary Operational Environmental Satellite (GOES, see Section 2.2). Treating the Sun as a star also makes finding clear signatures of incredibly small solar flares more challenging. The importance of investigating small solar flares will be discussed in later sections (Section 1.3).

The Impulsive Phase

The impulsive phase of a flare is the most explosive stage, heating plasma to high temperatures, and is characterised by a sharp rise in emission. It is at this stage that the non-thermal HXR footpoints in the chromosphere become visible. Therefore, this is taken as evidence for accelerated particle emission localised in the flare's footpoints (Fletcher et al. 2011). The HXR emission peaks and decays quickly during this phase while the SXRs peak at the end of this phase; in keeping with the Neupert Effect as discussed (Neupert 1968; Benz 2017). The footpoint plasma, meanwhile, gets heated and evaporates up into the loop while emitting SXRs. There is also evidence of super-hot plasma being present from the early stages of flares (Caspi & Lin 2010).

The Gradual Phase

Once there is no further explosive energy input from the magnetic reconnection process, i.e. at the end of the impulsive phase, the gradual phase begins. It has been observed that the post-flare loops can produce temperatures comparable to the flaring loops (Fletcher et al. 2011) and that the decay can last longer than the cooling timescale, suggesting further energy input into the loops (MacCombie &

Rust 1979). In addition, it has also been observed that post-flare loops release a significant fraction of the total energy of the flare (Kuhar et al. 2017). Other potential post-flare phenomena, such as solar moss, may also suggest further energy input after a flare's impulsive phase (Berger et al. 1999).

From this flare description obtained from numerous observations and studies, a common charactieristic seen during the evolution of a flare is that the high energy HXR emission peaks earlier and is more impulive than the lower energy EUV and SXR emission. This helps define the flaring process and should be present across all flare scales.

1.2.2 Flare Location, Classification, and Distribution

It is difficult to predict the occurrence of individual solar flares, as is the case with many solar phenomena; however, their long term bulk behaviour appears to follow certain trends. One trend is the 11 year solar cycle where the Sun's activity peaks and troughs (Phillips et al. 2012, pp. 16–20). This cycle can be seen in the number of ARs, sunspots, and flares that appear over several years. The number and scale of these phenomena is at its most during solar maximum and its least at solar minimum.

In addition to collective temporal behaviour there is also spatial behaviour. At the beginning of a cycle, sunspots, and so flares, occur furthest away from the equator (~30° north and south) then progressively occur closer to the equator until the end of the cycle (Charbonneau 2010 and Phillips et al. 2012, pp. 17). Both the spatial and temporal distribution of solar flares, and other measures of solar activity, is well documented with *butterfly diagrams*, an example of which is shown in Figure 1.3.

In order to compare different scaled flares they are usually categorised in terms of their thermal energy release (Table 1.1). In addition, they are also classified through their SXR 1–8 Å emission detected by GOES. The GOES flare classification scale separates flare SXR emission into five bands; X, M, C, B, and A. An X-class flare is the largest GOES classification starting at a SXR flux of $>10^{-4}$ W m^{-2}, and drops an order of magnitude with each band until the A category. Table 1.1 shows the comparison between the GOES SXR flux classification and the broader instantaneous thermal energy categories. Table 1.1 helps us classify flares across decades of energies allowing the flare frequency distribution to be investigated.

The solar flare frequency distribution is consistent with a power-law magnitude distribution with lower energy flares occurring more frequently (Crosby et al. 1993). Even with the individual massive energy releases of the larger sized flares, they do not occur often or steadily enough throughout the solar cycle to offer a solution to the coronal heating problem (Hudson 1991); therefore, weaker energy flares are of particular interest as they may contribute more net energy compared to their

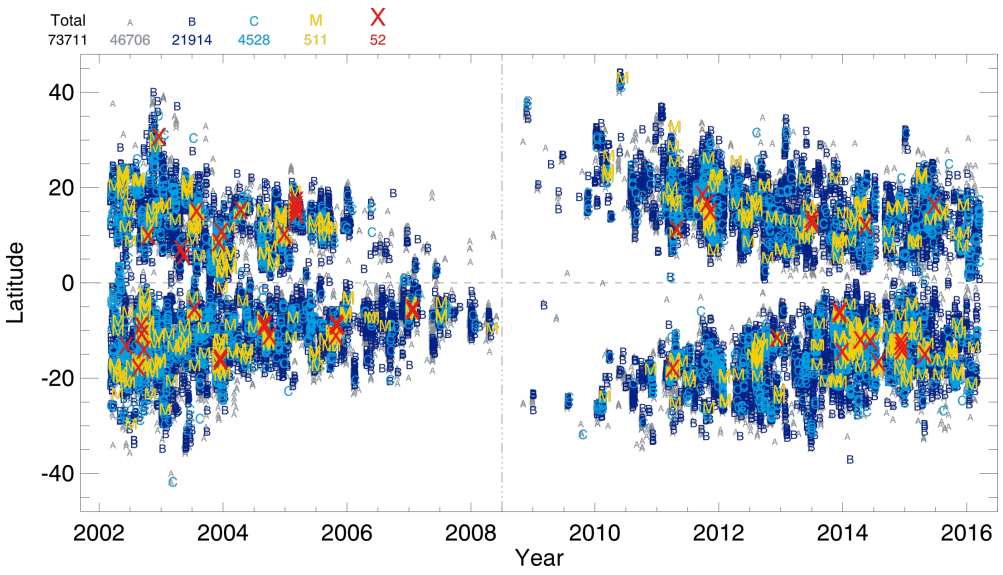

Figure 1.3: A flare butterfly diagram showing the latitude of RHESSI solar flares with their corresponding GOES classification (see Section 2.2) over a fourteen year period. This figure is taken from Loumou et al. (2018).

less frequent, higher energy counterparts. Microflares are flares with energies $\sim 10^{26-28}$ erg and are labelled as B, A, and sub-A GOES class flares (Lin et al. 1984; Hannah et al. 2011; Fletcher et al. 2011). GOES does not accurately detect sub-A class microflares; however, sub-A class events have been identified in the quiet Sun (Sylwester et al. 2012; Vadawale et al. 2021a) and in ARs (Gburek et al. 2011; Vadawale et al. 2021b) using more sensitive X-ray spectrometers like the Solar X-ray Monitor (XSM; Vadawale et al. 2014; Shanmugam et al. 2020). It has also been proposed that there may be even smaller events (*nanoflares*), with energies $\sim 10^{24}$ erg, that occur constantly over the solar disk, not just in ARs, and still output a greater total energy compared to the larger and less frequent flares (Parker 1988).

Events observed in EUV have been estimated to have energies of the same order as nanoflares, and are sometimes termed as such (Parnell & Jupp 2000; Aschwanden et al. 2000; Chitta et al. 2021); however, flares of this scale have not been observed in X-rays. The weakest X-ray flares observed have only been calculated down to $\sim 10^{26}$ erg, two orders of magnitude greater than the events seen in EUV (Vadawale et al. 2021b).

The distribution of solar flares, with regards to flare frequency, must have a power law index >2 if the small flare population is sufficient to heat the corona (Hudson 1991; Crosby et al. 1993). Therefore, it is important to investigate the physics operating in weaker flares to meaningfully compare the distribution of solar flares at different scales. In other words, are the weakest observed X-ray micro-

Table 1.1: Two scales on which flares are measured and compared; the instantaneous thermal energy of an event and the integrated flux ranges for GOES flare classification. This information is taken from Hannah et al. (2011).

Flare Size	Energy [erg]	GOES Classification	Flux [W m^{-2}]
Normal	$\leqslant 10^{33}$	X	$10^{-4} \leqslant F$
		M	$10^{-5} \leqslant F < 10^{-4}$
		C	$10^{-6} \leqslant F < 10^{-5}$
Micro	$\sim 10^{27}$	B	$10^{-7} \leqslant F < 10^{-6}$
		A	$10^{-8} \leqslant F < 10^{-7}$
Nano	$\sim 10^{24}$	$\ll$A	$\ll 10^{-8}$

flares, and weaker EUV nanoflares, scaled versions of their brighter counterparts—similar temporal, spatial, and spectral behaviour with thermal and non-thermal emission being present—or is there a point where the weaker events being observed become fundamentally different types of phenomena compared to larger flares?

1.3 Microflares

Studies conducted with the Reuven Ramaty High-Energy Solar Spectroscopic Imager (RHESSI; Lin et al. 2002, see Section 2.3) and NuSTAR (Harrison et al. 2013) have investigated solar X-ray microflares in great detail. Evidence of HXR non-thermal emission has been found to be produced during the evolution of microflares which suggests that similar processes to larger flares are indeed occurring at these weak flare scales; however, the energy release does not appear to scale with physical size at the microflare level (Christe et al. 2008; Hannah et al. 2008; Glesener et al. 2020). We also find that the higher energy X-ray microflare emission peaks before the lower energies which is sign of non-thermal emission or indicative of hotter plasma earlier in the flare phases (Duncan et al. 2021).

Figure 1.4 shows a comparison of temperatures and emission measures (Section 1.4.4) of RHESSI, STIX, FOXSI, and NuSTAR microflares from X-class flares (Warmuth & Mann 2016) to sub-A class microflares (Duncan et al. 2021). In Figure 1.4, it is shown that the hotter temperatures are only present in the flares with the largest emission measures. Battaglia et al. (2021) suggests that for flares to reach higher temperatures, more material must also be heated and, thus, more energy must be released. As a consequence, sensitive HXR observations are crucial to further our understanding of the physical processes that are present in microflares and to investigate the hottest temperatures that are being reached.

[2] https://sprg.ssl.berkeley.edu/~tohban/wiki/index.php/STIX,_the_Hard_X-Ray_Telescope_on_board_Solar_Orbiter

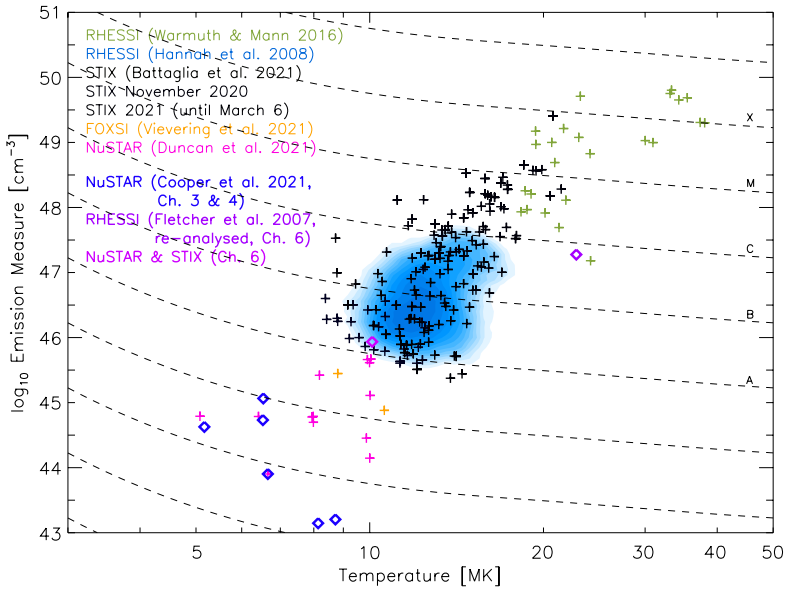

Figure 1.4: Isothermal temperature and emission measure comparison between RHESSI, STIX, FOXSI, and NuSTAR flares (Warmuth & Mann 2016; Hannah et al. 2008; Battaglia et al. 2021; Vievering et al. 2021; Duncan et al. 2021) with the dashed contours showing the GOES class. This figure is adapted from Figure 6 in Battaglia et al. (2021) with additional STIX observations[2] and microflares from Chapter 3, 4, and 6.

Note that the weakest microflares present in Figure 1.4 are observed with NuSTAR. Therefore, in Section 1.3.1, we provide an overview of published NuSTAR solar campaign events.

1.3.1 NuSTAR Solar Microflare Observation Overview

NuSTAR was launched in 2012 and observed the Sun for the first time in 2014 (Grefenstette et al. 2016; Hannah et al. 2016; Kuhar et al. 2017). A great deal of insight has been gained from several solar events that NuSTAR has observed, with sub-A class microflares within ARs reaching temperatures greater than their surroundings (>5 MK) with energies of $\sim 10^{26-28}$ erg (Glesener et al. 2017, 2020; Wright et al. 2017; Hannah et al. 2019) and small brightenings outside of ARs with temperatures 3–5 MK with energies of $\sim 10^{26}$ erg (Kuhar et al. 2018) as solar activity decreased towards solar minimum between cycle 24 and 25. NuSTAR has also been able to provide the first focusing optics imaging spectroscopy observation of non-thermal emission from an A5.7 class microflare (Glesener et al. 2020) and observe one of the weakest AR microflares ever studied (Cooper et al. 2020, see Chapter 4).

As well as studies that focus on one to a few events, there are two NuSTAR published works which analyse approximately ten solar microflares. Duncan et al. (2021) analyses eleven A-class microflares, from two different ARs (AR12671 and AR12712), with energies of $\sim 10^{28}$ erg. The microflare X-ray

emission generally all display more impulsive, and earlier peaking, HXR time profiles compared to lower energy EUV emission, properties often associated with larger solar flares. A further ten microflares that occurred in AR12721 are investigated in Cooper et al. (2021, see Chapter 3). These microflares are all sub-A class flares of energy $\sim 10^{26-28}$ erg, with some reaching temperatures up to 10 MK and one showing compelling evidence for non-thermal emission. The majority of the ten microflares also show corresponding photospheric flux cancellation at, or close to, the microflaring loop footpoints.

Since its first solar campaign, there have been thirteen NuSTAR solar publications with $>$30 solar observations. These studies and observations help build a picture of a 'standard' low A-class X-ray microflare and contribute to the sample size of NuSTAR solar microflares available for future statistical studies. NuSTAR allows us to investigate the Sun in a unique manner, providing insight into how small flares behave and how they are distributed compared to their larger counterparts[3]. NuSTAR will be discussed further in Section 2.1.

1.4 Flare EUV and X-ray Emission Analysis

Since the solar atmosphere has relatively high temperatures and low densities compared to its surface, we find that its emissions are best observed in the EUV and X-ray ranges (Section 1.3 of Phillips et al. 2012). Therefore, the EUV and X-ray wavelengths provide a crucial insight into phenomena that occur in the solar atmosphere, such as flares, through imaging and spectroscopic measurements. Observations in these wavelength ranges can then be used to gain information on the temporal, spatial, and spectral evolution of flares and other solar events.

In this section, we will describe the Richardson-Lucy method for image deconvolution which can be used to uncover feature geometry in images, this is especially useful when comparing X-ray to EUV images as the former generally have a much lower spatial resolution (Section 1.4.1). In Section 1.4.2, we describe EUV and X-ray atomic lines and continuum seen in flare spectra with the X-ray bremsstrahlung continuum discussed in Section 1.4.3.

We describe the concept of *emission measure* in Section 1.4.4 which is a fundamental plasma parameter. We then present the thermal spectral model created through the combination of spectral lines and continuum emission from an isothermal plasma. The spectral model can then be used to *forward-fit* observed spectra with a variety of programs to determine the emitting plasma's properties (Chapter 5 and 6). Section 1.4.4 then describes how the thermal model parameters can help compare emission observed by different instruments under an isothermal assumption. Finally, in Section 1.4.5, we

[3]An overview of NuSTAR solar observations is available at `https://ianan.github.io/nsigh_all/`.

describe several other spectral parametric models able to be used in fitting analysis that are created by assuming the presence of a non-thermal electron distribution.

1.4.1 Image Deconvolution

Images produced from observations, especially lower resolution images such as those produced from X-ray instruments, can appear blurred due to the effects of the instrumental optics. We can, however, attempt to improve the constructed images by using a deconvolution method. The method we employ throughout this thesis is the Richardson-Lucy method (Richardson 1972) which attempts to remove the instrument point spread function's (PSF) effects from the observed image via a Bayesian iterative method. Although more computationally expensive than other methods, the iterative process allows some control over artifacts that are introduced due to noise sources in the image.

For this method of deconvolution, following the derivation by Richardson (1972), Bayes' theorem can take the form

$$P(W_l|D_k) = \frac{P(D_k|W_l)P(W_l)}{\sum_j P(D_k|W_j)P(W_j)}, \tag{1.2}$$

to find the probability of an event occurring at pixel l of the true image, W, given an event in the observed map, D, at some position k due to a point spread function. Here we have explicitly expressed the evidence term, $P(D_k)$, on the denominator as the marginalisation over all true image pixels, j, within the PSF. However, we aim to find an expression for the estimated true image probability; therefore, we obtain

$$P(W_l) = \sum_k P(W_l, D_k) = \sum_k \frac{P(D_k|W_l)P(W_l)}{\sum_j P(D_k|W_j)P(W_j)} \cdot P(D_k) \tag{1.3}$$

via the marginalised probability of the measured image from D. For a given pixel, we can associate the probability of an event occurring with the ratio of that pixel value to the whole integrated array. We can then replace each probability and likelihood-like expressions with

$$P(W_l) = \frac{W_l}{\sum_l W_l} \qquad P(D_k) = \frac{D_k}{\sum_l W_l} \qquad P(D_k|W_l) = P(A_{l,k}) = \frac{A_{l,k}}{\sum_{l,k} A_{l,k}},$$

where we rely on the fact that the deconvolution process is conservative (i.e., $\sum_l W_l = \sum_k D_k$) and A describes the mapping between D and W. We also interpret Equation 1.3 to have all prior information on the right hand side, therefore, allowing us to ideally use estimates of the true image to iteratively obtain a better result. Equation 1.3 then takes the form

$$W_{l,i+1} = W_{l,i} \sum_k \frac{A_{l,k} D_k}{\sum_j A_{j,k} W_{j,i}}. \tag{1.4}$$

where i refers to the iteration number. Equation 1.4 is of a general form where the indices l, k, and j can then be specified to apply to finite array sizes of n-dimensions.

In practice, we assume that only the optics affect the final recorded image and that there are no other sources of noise. For the first iteration, a uniform array is often assumed. This creates an image of what would have been detected if a perfect flat field was viewed by the instrument in question. The image created from the data is then divided by this result giving a measure of the relative blur. The relative blur is then convolved with the transpose of the PSF and multiplied by the estimated underlying image. Equation 1.4 then takes the form

$$W_{i+1} = W_i \times \left(\frac{D}{W_i * A} * A^{\mathrm{T}} \right), \tag{1.5}$$

where W is the desired underlying image at iteration i, D is the image constructed from the data, A is the instrument PSF, and $*$ represents convolution.

An example of deconvolution on a NuSTAR image is shown in Figure 1.5 where we display different iteration numbers. The X-ray feature is smoothed with relatively few iterations and then maintains its general shape and orientation as the iteration number is increased while the spatial extent of the feature is reduced. Information on NuSTAR's PSF can be found in Section 2.1.1.

1.4.2 Flare EUV and X-ray Emission Lines

Spectroscopic measurements of emission lines can provide information on temperatures, densities, abundances, motions, and magnetic field strengths in the emitting source. Therefore, emission lines from atoms or ions that reside in the EUV and X-ray wavelength ranges are useful diagnostics when investigating the solar atmosphere.

In order to utilise the spectral diagnostic power of emission lines with regards to solar spectra, we use `CHIANTI` (Dere et al. 1997). `CHIANTI` provides a database of line information on atom and ion populations under optically thin equilibrium conditions which are common in cosmic plasmas and is actively updated for the research community (Del Zanna et al. 2021).

In this thesis, we study EUV images, which are sensitive to a number of emission lines, with the Solar Dynamic Observatory's Atmospheric Imaging Assembly (SDO/AIA; Lemen et al. 2012, see Section 2.5.1) in Chapter 3, 4, and 8. Additionally, we use magnetic field information derived from Helioseismic and Magnetic Imager measurements of the Fe I 6173 Å emission line (SDO/HMI; Schou et al. 2012, see Section 2.5.2) in Section 3.9.

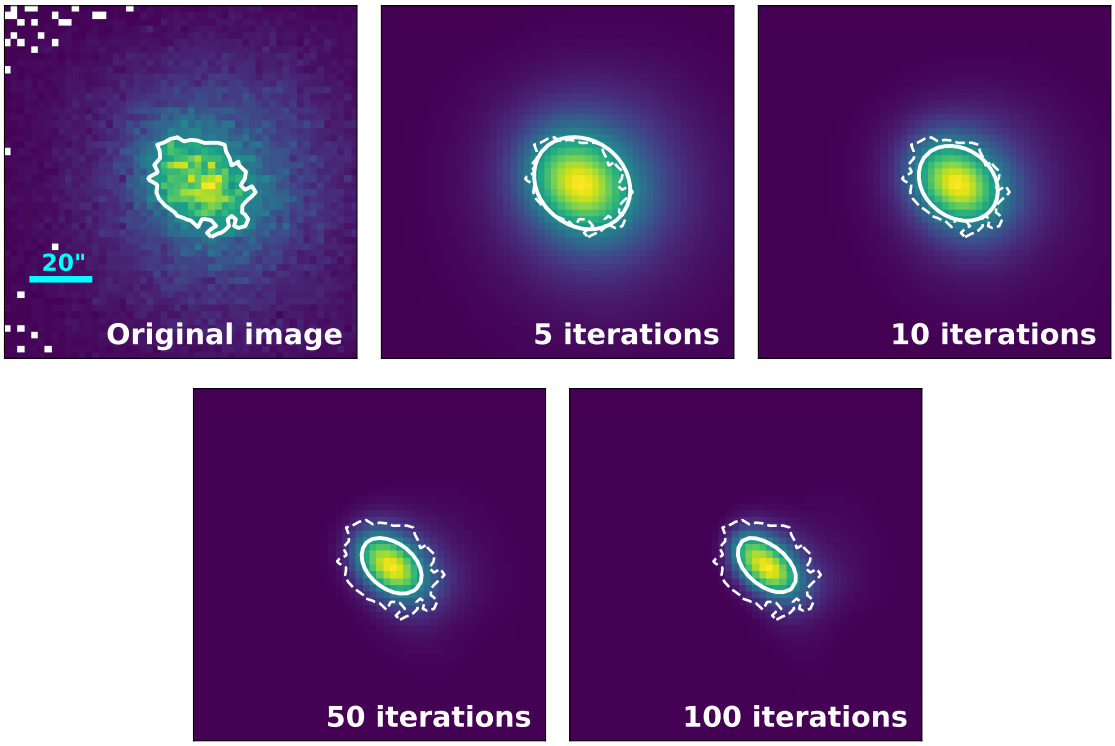

Figure 1.5: The top left panel shows the original image made from NuSTAR FPMB emission during micro-flare 10 from Section 3.7. The accompanying images depict this original image deconvolved with the stated number of iterations. White contours show the 50% level of each maximum normalised image. The contours from the original image are shown in all deconvolved panels as a dotted line. A reference scale of 20″ is shown in cyan.

We investigate the X-ray wavelengths (mainly with NuSTAR) also using images; however, we are able to study the emission spectroscopically (Chapter 3, 4, 6, 7, and 8). There are a number of spectral lines within the X-ray wavelength range but we make specific note of the Fe-line feature at 6.7 keV which is composed of numerous emission lines from Fe XXV and Fe XXVI (Phillips et al. 2012, pp. 136) and visible in many spectra presented in Chapter 3 and 8. The line feature is also visible in Figure 1.8 (orange line). The location of this line complex has helped discover the need for a livetime dependent response correction in NuSTAR solar spectral analysis (Section 2.1.5).

In addition to both EUV and X-ray wavelength ranges including *bound-bound* processes producing emission lines, they also contain *free-bound* (recombination processes) and *free-free* (bremsstrahlung, see Section 1.4.3) processes. In EUV wavelengths ($<$3 keV) the *bound-bound* transitions are dominant in the spectral emission producing very clear line profiles and a low surrounding continuum. However, in X-rays ($>$3 keV) the continuum emission processes (*free-bound* and *free-free*) dominate (Phillips et al.

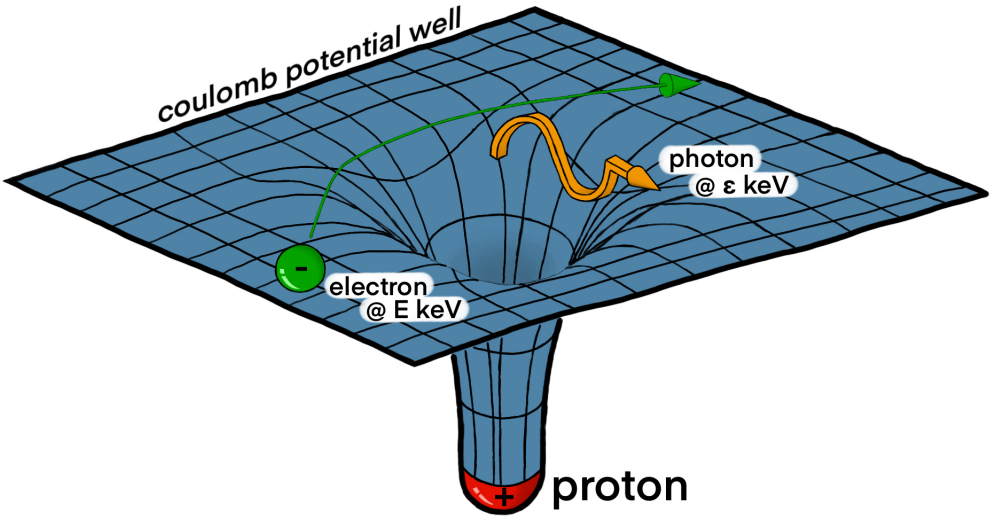

Figure 1.6: An illustration of the bremsstrahlung radiation process of an electron (green) with initial energy E being decelerated by the oppositely charged proton (red) where the path of the electron is shown with a green arrow. A photon (orange) is produced from the interaction with energy ε such that $\varepsilon \leqslant E$.

2012, pp. 8). Therefore, because of its significance in X-ray spectra, we discuss the bremsstrahlung emission process in Section 1.4.3.

1.4.3 Bremsstrahlung

Bremsstrahlung, meaning braking radiation, is the process by which an electron is decelerated in the presence of another electric charge, emitting a photon (Carroll & Ostlie 2014, pp. 278) and is the most common source of flare emission (Section 5.2.1 of Tandberg-Hanssen & Emslie 2009).

Figure 1.6 shows the process of bremsstrahlung emission where we note that the energy of the emitted photon must be less than or equal to the initial electron energy; i.e., $\varepsilon \leqslant E$. Figure 1.6 only shows the *free-free* radiation for a single interaction; however, in a plasma, flaring or otherwise, this process is continuously happening with a vast number of particles.

Figure 1.7 shows how the single interaction depicted in Figure 1.6 can occur on large scales with an electron flux interacting with numerous protons in a fully ionised plasma. An expression for the photon flux $I(\varepsilon)$ from the interaction shown in Figure 1.7 detected at a distance D from an isotropically emitting source is given by

$$I(\varepsilon) = \frac{1}{4\pi D^2} \int_V \int_\varepsilon^\infty n_p F(E) \sigma_B(\varepsilon, E) dE dV \quad [\text{photons s}^{-1}\ \text{keV}^{-1}\ \text{cm}^{-2}], \tag{1.6}$$

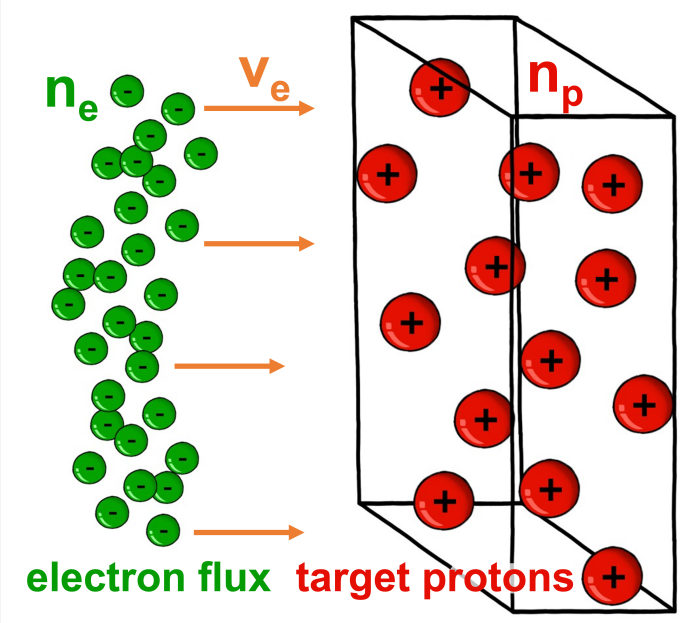

Figure 1.7: A simplistic illustration of how bremsstrahlung radiation can result from a plasma where multiple electrons (green) with number density n_e are propagating with velocity v_e into a slab of oppositely charged protons (red) with number density n_p. The scenario depicted in Figure 1.6 occurs for every electron-proton interaction.

where n_p is the number density of the protons, $F(E)$ is the differential electron flux in electron energy E (given by $d(n_e v_e)/dE$, with n_e being the electron number density and v_e the electron velocity as shown in Figure 1.7), and $\sigma_B(\varepsilon, E)$ is the differential bremsstrahlung cross-section in photon energy ε (Tandberg-Hanssen & Emslie 2009, pp. 105–116, 181). In Equation 1.6, we are integrating over the emitting source volume V. We also integrate over all electron energies greater than the given photon energy ($> \varepsilon$) as an electron with $E < \varepsilon$ cannot contribute to the emission of photons at ε (see Figure 1.6).

The bremsstrahlung cross-section represents the space where an electron will interact with an ion; however, the interaction is not constant across this area (Section 5.2 of Longair 2018). From the simple cartoon shown in Figure 1.6, we observed that an electron further from, or closer to, the proton will interact more weakly, or strongly. In addition, the temporal length of the interaction will influence the energy of the emitted photon which is dependent on the electron energy; for example, a higher energy electron will not interact with the proton for as long compared to a lower energy electron approaching the proton under the same conditions. Therefore, the cross-section is dependent on the incoming electron energy and the emitted photon energy (Tandberg-Hanssen & Emslie 2009, pp. 106–107).

From Equation 1.6, we can see that the photon flux I from bremsstrahlung is completely dependent

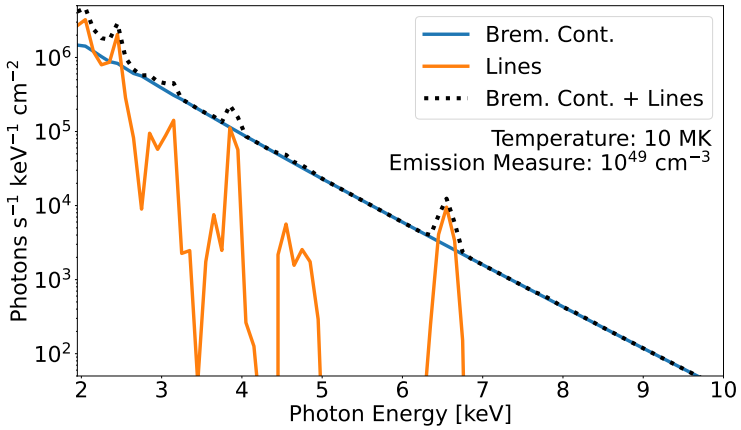

Figure 1.8: The thermal bremsstrahlung photon spectrum (blue) and line emission (orange) produced from an isothermal, completely ionised source with coronal abundances (Feldman et al. 1992) radiating isotropically and measured at Earth. The black dotted line shows the total thermal model. This plot has been produced using the Python package Sunxspex (Section 5.4).

on the differential electron flux spectrum $F(E)$. As such, we can define different electron spectra to include in Equation 1.6 and obtain the corresponding photon spectra that can then be compared to spectral observations through a process such as forward-fitting (see Chapter 5, 6, and 7). For example, we can investigate whether a spectrum is composed of thermal emission (Section 1.4.4), non-thermal emission (Section 1.4.5), or a combination of both. This then allows the electron distribution of the emitting source to be determined.

1.4.4 Models with Thermal Bremsstrahlung

As discussed in Section 1.2.2 and 1.3, and shown in Figure 1.4, flares can heat plasma to a wide range of temperatures. Therefore, it becomes useful to create a mathematical model of the photon spectrum that is expected from an isothermal, fully ionised emitting source. Figure 1.8 shows the photon thermal model being discussed in this section. First we discuss the steps to obtain the bremsstrahlung continuum (blue line in Figure 1.8) then the spectral line emission (orange line). The two thermal contributions are then combined to produce the total model (black dotted line).

In order to derive the thermal bremsstrahlung continuum emission, we define the *Maxwell-Boltzmann energy distribution* which describes the fraction of electrons in an energy range between E and $E + dE$ of a system in thermal equilibrium. The energy distribution, $f(E)$, is given by

$$f(E)dE = \frac{2}{\sqrt{\pi}} \frac{E^{1/2}}{(k_B T)^{3/2}} e^{-\frac{E}{k_B T}} dE, \tag{1.7}$$

where E is the electron energy, k_B is Boltzmann's constant, and T is the isothermal temperature of the plasma (Carroll & Ostlie 2014, pp. 234). Therefore, we can write

$$F(E) = n_e v_e f(E) \quad [\text{electrons s}^{-1}\ \text{keV}^{-1}\ \text{cm}^{-2}], \tag{1.8}$$

obtaining an expression for the thermal differential electron flux, $F(E)$ in Equation 1.6.

The next step is to obtain a suitable expression for the differential bremsstrahlung cross-section $\sigma_B(\varepsilon, E)$; however, $\sigma_B(\varepsilon, E)$ can quickly take a complicated form (Haug 1997). To progress with this example we make use of the Bethe-Heitler cross-section formula which takes the form

$$\sigma_B(\varepsilon, E) = \frac{7.9 \times 10^{-25}}{\varepsilon E} ln\left(\frac{1 + (1 - \varepsilon/E)^{1/2}}{1 - (1 - \varepsilon/E)^{1/2}}\right) \quad [\text{cm}^2\ \text{keV}^{-1}], \tag{1.9}$$

where this is only valid for non-relativistic cases and is integrated over all directions (Tandberg-Hanssen & Emslie 2009, pp. 107).

We can simplify Equation 1.9 further by utilising Kramers approximation (Holman et al. 2011) which removes the logarithmic term. Therefore, the bremsstrahlung photon spectrum produced from a fully ionised, isothermal plasma can be expressed as

$$I_{thermal\ brem.}(\varepsilon) = \sqrt{\frac{2}{\pi m_e}} \frac{1.6 \times 10^{-24}}{4\pi D^2 \varepsilon (k_B T)^{1/2}} \cdot n_e^2 V \cdot e^{-\frac{\varepsilon}{k_B T}} \quad [\text{photons s}^{-1}\ \text{keV}^{-1}\ \text{cm}^{-2}], \tag{1.10}$$

where m_e is the electron mass, V is the volume of the emitting source, and we have used $E = \frac{1}{2} m v_e^2$ and that in a fully ionised plasma we have $n_p = n_e$.

Since we have obtained an expression for the bremsstrahlung thermal continuum, we now need to calculate the spectral contribution from the emission lines. The photon flux for each line is calculated using

$$I_{line}(\varepsilon) = \frac{1}{4\pi D^2} \int_V G(\varepsilon, T, n_e) n_e^2 dV \quad [\text{photons s}^{-1}\ \text{keV}^{-1}\ \text{cm}^{-2}], \tag{1.11}$$

where $G(\varepsilon, T, n_e)$ is called the *contribution function* which is dependent on relative level populations of the emitting source abundances (Section 4.6 of Phillips et al. 2012). The information for the contribution function calculations can be obtained using `CHIANTI` which was discussed in Section 1.4.2.

In order to obtain the total thermal emission from an isothermal, fully ionised, homogeneous, isotropically emitting source we can combine the bremsstrahlung continuum (Equation 1.10, blue line in Figure 1.8) and line (Equation 1.11, orange line) emission. This is accomplished by simply summing

both components as

$$I_{thermal}(\varepsilon) = I_{thermal\ brem.}(\varepsilon) + I_{line}(\varepsilon) \quad [\text{photons s}^{-1}\ \text{keV}^{-1}\ \text{cm}^{-2}], \tag{1.12}$$

where $I_{thermal}(\varepsilon)$ is the total thermal model (black dotted line in Figure 1.8) and can be used in spectral analysis (see Section 5.2). Equation 1.12 should also have a term representing *free-bound* processes; however, this mechanism contributes significantly less than bremsstrahlung to the X-ray continuum (Section 4.7.2 of Phillips et al. 2012). Therefore, a *free-bound* term is not included in Equation 1.12 for simplicity.

It is clear that Equation 1.12 will model a photon spectrum produced from a plasma in thermal equilibrium at a given temperature (T), electron density (n_e), and source volume (V) allowing this model to be easily forward-fitted to an observed spectrum. However, we note that the combined quantity of the electron number density and the source volume has a significant physical meaning and is termed the *emission measure* (EM). The emission measure will be discussed next.

Emission Measure

The squared electron density n_e^2 and the source volume V are present in all processes that contribute to the total thermal model (Equation 1.12). The emission measure (EM) is the product of these two quantities being expressed as

$$EM = \int_V n_e^2 dV = n_e^2 V \quad [\text{photons cm}^{-3}], \tag{1.13}$$

which gives the number of photons emitted per unit volume (Phillips et al. 2012, pp. 94). The rightmost expression implies an isothermal and homogeneous emitting source.

If instead we consider a non-homogenous, multi-thermal plasma, the integral term cannot be evaluated as in Equation 1.13. We then consider the spatial distribution of the different temperature components (Craig & Brown 1976) allowing the emission measure to be re-expressed as

$$EM = \int_V n_e^2 dV = \int_T \xi(T) dT \quad [\text{photons cm}^{-3}], \tag{1.14}$$

where $\xi(T)$ is the differential emission measure (DEM) which is integrated over all temperatures to obtain the EM (Section 2.6 of Tandberg-Hanssen & Emslie 2009). The differential emission measure, $\xi(T)$, then gives the emitting power of the source between temperatures T and $T + dT$.

Calculating the DEM is not a straightforward process and is an ill-posed[4] inverse problem (Jefferies et al. 1972; Craig & Brown 1977). Several methods have been created in order to recover the DEM from a set of observations, one of which is described in Hannah & Kontar (2012) and used in Section 4.5 to recover the DEM of a microflare from EUV and X-ray observations.

The emission measure is a fundamental property of the emitting source and can be used in combination with temperature to calculate the thermal energetics of the source. We now discuss how this is performed.

Thermal Energy Estimates from Thermal Models

From the temperature, T, and emission measure, EM, of a thermally modelled spectrum (Equation 1.12), accompanied by the emitting structure volume (V) obtainable from high resolution images (EUV images, for example), we can obtain an instantaneous thermal energy estimate, E_{th}, for the emitting plasma. This can be calculated using

$$E_{th} = 3N_e k_B T = 3(fV \times EM)^{\frac{1}{2}} k_B T \quad [\text{erg}], \tag{1.15}$$

where N_e is the total number of thermal electrons and k_B is Boltzmann's constant (Hannah et al. 2008). The filling factor, f, is generally assumed to have a value of 1 unless stated otherwise. If, instead, we have a multi-thermal plasma then Equation 1.15 becomes

$$E_{th} = 3(fV)^{\frac{1}{2}} k_B \sum_{i=1}^{N} T_i EM_i^{\frac{1}{2}} \quad [\text{erg}], \tag{1.16}$$

where we sum over N different thermal model components with temperature and emission measure T_i and EM_i, respectively (Aschwanden et al. 2015).

The temperature and emission measure can also be used to compare emission observed by different instruments which is discussed in the following section.

Synthetic Flux Estimates

We can use the parameters derived from a fitted thermal model (Equation 1.12) to one instrument's data to help compare with emission seen by other instruments. This is done using

$$I_i = \frac{1}{A} \int_0^\infty R_i(T) \xi(T) dT \quad [\text{DN s}^{-1} \text{ pix}^{-1}], \tag{1.17}$$

[4]No unique solution.

where I_i is the emission rate, A is the apparent area of emission, $R_i(T)$ is the instrument temperature response (see Section 2.5.1 for a specific example), and $\xi(T)$ is the differential emission measure (see Boerner et al. 2012). The subscript i denotes the instrument channel in question. In Equation 1.17, we are integrating over all temperatures present in the observed plasma. However, if we assume an isothermal plasma is being observed between observatories (e.g., when investigating microflares with the Fe XVIII proxy discussed in Section 2.5.1 and NuSTAR from Section 2.1), Equation 1.17 can be re-written as

$$I_i \approx \frac{1}{A} R_i(T^*) EM^* \quad [\text{DN s}^{-1}\ \text{pix}^{-1}], \tag{1.18}$$

where T^* and EM^* is the characteristic temperature and emission measure from X-ray spectral fitting results. We can then calculate the emission expected to be seen from a plasma with that same temperature and emission measure for any instrument given its temperature response.

This is valid under a thermal assumption; however, as discussed in Section 1.4.3, the source electron distribution for the photon model can take any form and does not have to be a thermal one. Therefore, we investigate non-thermal electron distributions with regards to Equation 1.6 in Section 1.4.5.

1.4.5 Models with Non-thermal Bremsstrahlung

In addition to thermal electron distributions, we can also calculate the photon spectrum produced when a beam of superathermal electrons[5] are injected into an proton source. To do this we assume a source function ($F(E)$ in Equation 1.6) of the form

$$F(E > E_C) = AE^{-\delta} \quad [\text{electrons s}^{-1}\ \text{keV}^{-1}\ \text{cm}^{-2}], \tag{1.19}$$

which describes a single power-law distributions of electrons with normalisation constant A and spectral index δ (Tandberg-Hanssen & Emslie 2009, pp. 111). In Equation 1.19, we define a low-energy cut-off, E_C, for the injected non-thermal electron distribution.

Therefore, the electron distribution is only defined for $E > E_C$ such that the energy distribution[6] satisfies

$$\int_{E_C}^{\infty} f(E)dE = 1, \tag{1.20}$$

where, in some cases, an upper limit E_{high} may be defined for the distribution as well as the lower limit E_C[7].

[5] Electrons with motion much larger than the local characteristic thermal velocity (https://link.springer.com/referenceworkentry/10.1007/978-3-642-11274-4_1549).

[6] The distribution that describes the fraction of electrons at each energy, equivalent to the thermal bremsstrahlung case in Equation 1.7

[7] https://hesperia.gsfc.nasa.gov/ssw/packages/xray/doc/brm_thin_doc.pdf

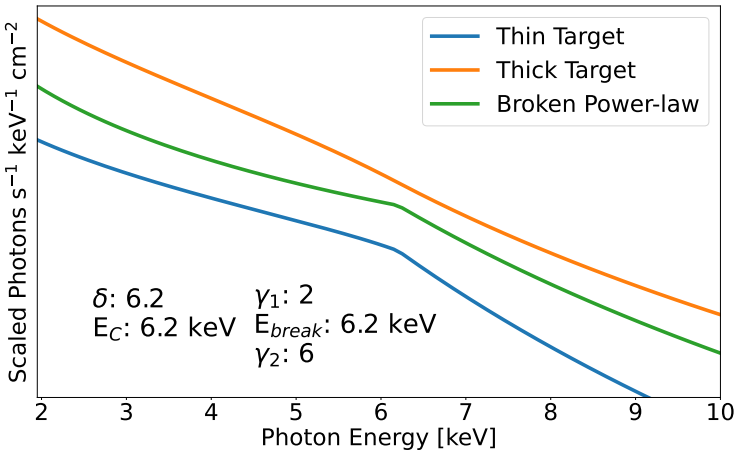

Figure 1.9: Scaled non-thermal bremsstrahlung photon spectra produced from injecting a single power-law of electrons, with electron spectral index δ and low-energy cut-off E_C, into a thin target (blue) and a thick target (orange). The final model shown (green) is an approximation of the thick target which models the photon spectrum directly with a broken power-law, calculated with a spectral index below and above the break energy (E_{break}) of γ_1 and γ_2, respectively. See text for more details on specific models. This plot has been produced using the Python package Sunxspex (Section 5.4).

Note that the total electron flux F_{tot} is then given by

$$F_{tot} = \int_{E_C}^{\infty} F(E)dE \quad [\text{electrons s}^{-1}\text{ cm}^{-2}], \tag{1.21}$$

where we have integrated over all electron energies above the defined lower limit (Tandberg-Hanssen & Emslie 2009, pp. 113).

Figure 1.9 shows all non-thermal models discussed in this section. In this section, we describe the scenario of the electron spectrum being injected into a thin target (blue line) and also a thick target (orange line) bremsstrahlung source. We also discuss how the energy in the electron spectrum can be calculated. We then describe an approximation to the thick target model which is much simpler and faster computationally (Figure 1.9, green line) along with how to relate this back to the electron spectrum and obtain the power estimate.

The Thin Target

In the thin target case, the target proton density is low with the electrons weakly interacting in the medium, like electrons propagating through the solar corona (Tandberg-Hanssen & Emslie 2009, pp. 108). The thin target case would also be applicable to an observation of a thick target scenario for a duration much shorter than the electron energy loss time. This process is generally straightforward

to model starting from Equation 1.6 and is shown in Figure 1.9 (blue line). We are able to follow a similar approach taken in the thermal bremsstrahlung case (Section 1.4.4) to illustrate its simplicity.

Taking the Bethe-Heitler bremsstrahlung differential cross-section from Equation 1.9, we derive a form for the thin target bremsstrahlung photon spectrum. This is given by

$$I_{thin}(\varepsilon) = \frac{7.9 \times 10^{-25} n_p V}{4\pi D^2 \varepsilon} \int_{\varepsilon}^{\infty} \frac{F(E)}{E} ln \left(\frac{1 + (1 - \varepsilon/E)^{1/2}}{1 - (1 - \varepsilon/E)^{1/2}} \right) dE \quad [\text{photons s}^{-1}\ \text{keV}^{-1}\ \text{cm}^{-2}], \tag{1.22}$$

where all notation has been defined previously. Inserting the single power-law electron distribution shown in Equation 1.19, we obtain

$$I_{thin}(\varepsilon) = \frac{7.9 \times 10^{-25} n_p V A}{4\pi D^2 \varepsilon} \int_{\varepsilon}^{\infty} E^{-(\delta+1)} ln \left(\frac{1 + (1 - \varepsilon/E)^{1/2}}{1 - (1 - \varepsilon/E)^{1/2}} \right) dE \quad [\text{photons s}^{-1}\ \text{keV}^{-1}\ \text{cm}^{-2}], \tag{1.23}$$

giving a simple form for thin target bremsstrahlung emission where we note that the electron distribution is only defined for $E > E_C$ as described for Equation 1.20 (Tandberg-Hanssen & Emslie 2009, pp. 111).

The Thick Target

If the electron spectrum is injected into a thick target, the target proton density is high and the electrons completely thermalise within the bremsstrahlung source (Tandberg-Hanssen & Emslie 2009, pp. 108). An example of a thick target photon model spectrum (Brown 1971) is shown in Figure 1.9 (orange); however, this is more difficult to model than the thin target case as the thermalisation of electrons complicates the integral.

We can approach this by thinking of an appropriate depth l into the target particle source that would be sufficient to thermalise all of the injected electrons (i.e., a *stopping depth*). This is equivalent to the depth reached by the electrons into the 'slab' of protons in Figure 1.7. This is where the derivation for the thick target deviates from the thin target as, instead of Equation 1.22, we write

$$I_{thick}(\varepsilon) = \frac{S}{4\pi D^2} \int_{\varepsilon}^{\infty} F(E_0) \int_{l_0(E=E_0)}^{l_1(E=\varepsilon)} n_p(l) \sigma_B(\varepsilon, E) dl dE_0 \quad [\text{photons s}^{-1}\ \text{keV}^{-1}\ \text{cm}^{-2}], \tag{1.24}$$

where S is the area of the proton target perpendicular to the depth and we have kept the Bethe-Heitler bremsstrahlung differential cross-section as $\sigma_B(\varepsilon, E)$ for simplicity. Therefore, in Equation 1.24, we are integrating over all the stopping depths of the electrons in the distribution that contribute to the photon energy ε, from those that penetrate the deepest to l_0 from $F(E_0)$ ($E = E_0$) to those that

penetrate the least to l_1 ($E = \varepsilon$). We note again that the electron distribution is only defined for $E > E_C$ (Equation 1.20).

From Equation 1.24, we can then perform a change of variables from l to E. This can be done by first using the relation between the travel depth l, the electron speed (v_e), and the time taken (t) with $l = v_e t$. Then we can use the relation

$$\frac{dE}{dt} = -\left|\frac{dE}{dt}\right| \quad [\text{keV s}^{-1}], \tag{1.25}$$

since we know the electrons are losing energy (Tandberg-Hanssen & Emslie 2009, pp. 109).

Using Equation 1.25, Equation 1.24 becomes

$$I_{thick}(\varepsilon) = \frac{n_p S A}{4\pi D^2} \int_{\varepsilon}^{\infty} E_0^{-\delta} \int_{\varepsilon}^{E_0} \sigma_B(\varepsilon, E) \frac{v_e}{|dE/dt|} dE dE_0 \quad [\text{photons s}^{-1}\ \text{keV}^{-1}\ \text{cm}^{-2}], \tag{1.26}$$

where we have also included the single power-law electron distribution from Equation 1.19. The final step is to obtain an analytical form for Equation 1.25 which describes the electron energy loss. We use the following expression

$$\frac{dE}{dt} = -\frac{4\pi r_e^2 m_e c^4 n_p ln(\Lambda)}{v_e} \quad [\text{keV s}^{-1}], \tag{1.27}$$

where r_e is the classical electron radius (Longair 2018, pp. 167)[8], m_e is the electron mass, c is the speed of light in a vacuum, and Λ is the ratio of the largest interaction distance to the smallest between the electron and proton (Longair 2018, pp. 133).

Using Equation 1.27, we can now establish an equivalent equation for the thick target as we have for the thin target in Equation 1.23. Equation 1.26 becomes

$$I_{thick}(\varepsilon) = \frac{SA}{4\pi D^2 \varepsilon} C \int_{\varepsilon}^{\infty} E_0^{-\delta} \int_{\varepsilon}^{E_0} ln\left(\frac{1 + (1 - \varepsilon/E)^{1/2}}{1 - (1 - \varepsilon/E)^{1/2}}\right) dE dE_0 \quad [\text{photons s}^{-1}\ \text{keV}^{-1}\ \text{cm}^{-2}], \tag{1.28}$$

where $C = 7.9 \times 10^{-25}/(2\pi r_e^2 m_e^2 c^4 ln(\Lambda))$. Therefore, Equation 1.28 is the form for thick target bremsstrahlung emission as described in Brown (1971).

In addition to having forms for the photon models produced from a single power-law of electrons being injected into a thin and thick target, we can use the electron distribution parameters to obtain the total power of the spectrum $> E_C$.

[8] $r_e = e^2/(4\pi\varepsilon_0 m_e c^2)$, where e is the electron charge and ε_0 is the permittivity of free space.

Non-thermal Energy from a Single Power-law Electron Distribution

In order to obtain the total power of the electron spectrum, we first calculate an expression for normalisation constant A in Equation 1.19. This can be done by substituting Equation 1.19 into Equation 1.21 giving

$$F_{tot} = \int_{E_C}^{\infty} AE^{-\delta} dE = A\frac{E_C^{1-\delta}}{\delta - 1} \quad [\text{electrons s}^{-1}\ \text{cm}^{-2}], \tag{1.29}$$

which is rearranged to find the expression for A. Therefore, to find the total power flux of the single power-law injected electron spectrum $> E_C$ ($P(\geqslant E_C)$) we calculate

$$P(\geqslant E_C) = \int_{E_C}^{\infty} AE^{-\delta} E dE = A\frac{E_C^{2-\delta}}{\delta - 2} = F_{tot}\frac{\delta - 1}{\delta - 2}E_C \quad [\text{keV s}^{-1}\ \text{cm}^{-2}]. \tag{1.30}$$

This gives a simple expression to calculate the energetics of the electron distribution. In practice, the thick target model used in OSPEX (Section 5.3.1) utilises the total electron rate ($e^- s^{-1}$) which is the product of F_{tot} and the area of radiation S[9]. Therefore, we define an $S \times F_{tot}$ term as the rate N_N.

The non-thermal power from a single power-law of electrons is calculated with

$$P(\geqslant E_C) = 1.6 \times 10^{-9}\frac{\delta - 1}{\delta - 2}N_N E_C \quad [\text{erg s}^{-1}], \tag{1.31}$$

which can be used directly with the thick target model parameters as described in Section 3.4.1 and 4.4 (Wright et al. 2017).

Broken Photon Power-law

Historically, calculating the thick target model was too computationally cumbersome and impractical for direct use in large scale fitting analysis. However, it has been shown that the power-law photon spectrum reflects the underlying power-law electron spectrum if the observed emission is assumed to be produced from a thick target source (Brown 1971).

A broken power-law photon model is required to represent a spectrum produced by a single power-law electron distribution due to the low-energy cut-off. The broken power-law spectrum ($I_{phot.}(\varepsilon)$) for

[9] https://hesperia.gsfc.nasa.gov/ssw/packages/xray/doc/brm_thick_doc.pdf

photon energies ε takes the form

$$I_{phot.}(\varepsilon) = \begin{cases} N\varepsilon^{-\gamma_1}, & \varepsilon \leqslant \varepsilon_{break} \\ N\varepsilon_{break}^{\gamma_2-\gamma_1}\varepsilon^{-\gamma_2}, & \varepsilon > \varepsilon_{break} \end{cases} \quad [\text{photons s}^{-1}\ \text{keV}^{-1}\ \text{cm}^{-2}], \tag{1.32}$$

where N is the normalisation constant and γ_1 & γ_2 are the photon spectral indices below and above the photon break energy ε_{break} (shown in Figure 1.9, green line). Note ε_{break} is different to the previously discussed electron low-energy cut-off, E_C. This method has been used in large scale microflare studies (Hannah et al. 2008), allowing the recovery of information on underlying single power-law electron distributions which is discussed in the following section.

Non-thermal Energy from Broken Power-law Photon Models

We can use the broken power-law parameters from a fitted photon model to calculate the power output from an assumed single power-law of accelerated electrons with a thick target approximation (Section 1.4.5). We calculate the power, $P(\geqslant E_C)$, in the accelerated electron distribution with

$$P(\geqslant E_C) = 9.5 \times 10^{24}\gamma^2(\gamma-1)\beta\left(\gamma - \frac{1}{2}, \frac{3}{2}\right)KE_C^{(1-\gamma)} \quad [\text{erg s}^{-1}], \tag{1.33}$$

where γ is the photon power-law index above ε_{break}, β is the beta function (Brown 1971; Lin 1974), K is the normalisation constant of the photon power-law spectrum above the break, and E_C is the low-energy cut-off of the electron distribution (Equation 4 in Hannah et al. 2008).

First we must calculate the normalisation constant of the photon power-law above the break if not already available. This is done using

$$K = N\varepsilon_{break}^{\gamma_2-\gamma_1} \quad [\text{ph s}^{-1}\ \text{cm}^{-2}\ \text{keV}^{-1}], \tag{1.34}$$

where N & γ_1 represent the normalisation constant and the index below the break energy (ε_{break}) in the photon spectrum and K & γ_2 represent the normalisation and the index above ε_{break}.

One challenge is using the broken power-law photon spectrum to obtain the low-energy cut-off from the thick target equation. This difficulty arises as the photon spectrum break and low-energy cut-off are not directly related with $\varepsilon_{break} < E_C$ (Hannah et al. 2011). However, we can use the result from statistical studies to estimate the low-energy cut-off (E_C) in order to find the non-thermal power by using Equation 6 found in Hannah et al. (2008),

$$E_C \approx 0.15\gamma + (1.86 - 0.04\gamma)\varepsilon_{break} - 3.39 \quad [\text{keV}], \tag{1.35}$$

where γ is equivalent to γ_2 in Equation 1.34. Again, like Equation 1.33, this works under the assumption of thick target emission. This method is used in Chapter 3 and 8 in the absence of the thick target model.

2

Instrumentation

In this chapter, we will cover the instruments used throughout this thesis with particular emphasis on the Nuclear Spectroscopic Telescope ARray which will be discussed in Section 2.1. We will then briefly discuss the Geostationary Operational Environmental Satellite in Section 2.2, the Reuven Ramaty High-Energy Solar Spectroscopic Imager in Section 2.3, Solar Orbiter's Spectrometer/Telescope for Imaging X-rays in Section 2.4, and the Solar Dynamic Observatory's Atmospheric Imaging Assembly and Helioseismic & Magnetic Imager in Section 2.5.1 and 2.5.2, respectively.

2.1 The Nuclear Spectroscopic Telescope ARray

The Nuclear Spectroscopic Telescope ARray (NuSTAR; Harrison et al. 2013) is an astrophysical telescope with capabilities of observing the Sun above 2.5 keV. NuSTAR directly focuses X-rays onto two focal plane modules (FPMA & B) using Wolter-I type optics mounted on a 10 m mast and is able to achieve greater sensitivity and lower background than instruments like RHESSI (Hannah et al. 2016) and others discussed in this chapter. NuSTAR first observed the Sun in 2014 (Grefenstette et al. 2016; Hannah et al. 2016; Kuhar et al. 2017) with many following solar campaigns. See Section 1.3.1 for more detail on NuSTAR solar observations.

In Section 2.1.1, we discuss the physical components of the NuSTAR observatory as well as how the telescope pointing is controlled. The point spread function of NuSTAR's optics are also presented in Section 2.1.1. Since NuSTAR is an astrophysical telescope it has a limited throughput for observing the Sun; therefore, we discuss the difficulties this presents in Section 2.1.2. We discuss the effect sources outside NuSTAR's field of view (FOV) can have on a given observation in Section 2.1.3. In Section 2.1.4, we describe the NuSTAR data pipeline (`NuSTARDAS`[1] within HEASoft[2]) and how it can

[1] https://heasarc.gsfc.nasa.gov/docs/nustar/analysis/
[2] https://heasarc.gsfc.nasa.gov/docs/software/lheasoft/

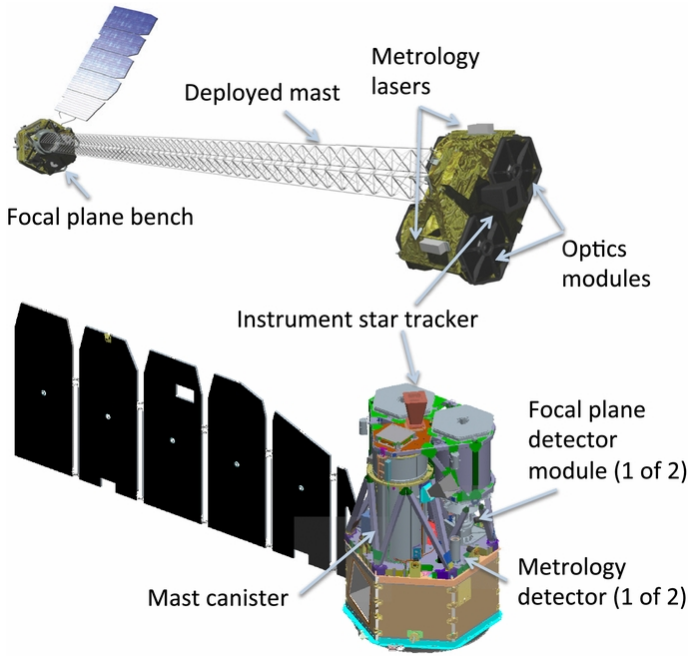

Figure 2.1: The NuSTAR observatory components (labelled) in their deployed configuration (top) and their stowed configuration (bottom). This figure is taken from Harrison et al. (2013).

be used to produce the data products required for analysis. These data products can then be employed to create higher order data products such as the instrument spectral response matrix (Section 2.1.5) and temperature response (Section 2.1.6).

2.1.1 Detectors and Optics

NuSTAR consists of two telescopes, each with an optics module on the 10 m mast which focusses X-rays onto a focal plane module (FPM A or B). Figure 2.1 shows the deployed version of NuSTAR (top, the stowed configuration is shown on the bottom) with the FPMs being located on the focal plane bench (Harrison et al. 2013). Each FPM has an FOV of 12′ by 12′ and is made up of four CdZnTe detectors (DET-ID 0–4), each separated by a chip-gap (Bhalerao 2012).

The chip-gap is indicated in a NuSTAR FPMA image by four navy arrows in Figure 2.2 (left panel) where it cuts through the emission observed from an active region (AR12721, see Chapter 3 and 4 for more on this solar campaign). To further highlight the splitting of AR12721, the emission per detector images are shown in the right panels. Note that the image created in Figure 2.2 (left panel) is integrated over a 5 min period with the bright AR not being the initial target and the pointing being chosen for the emission feature more centrally located in the FOV (see Chapter 3).

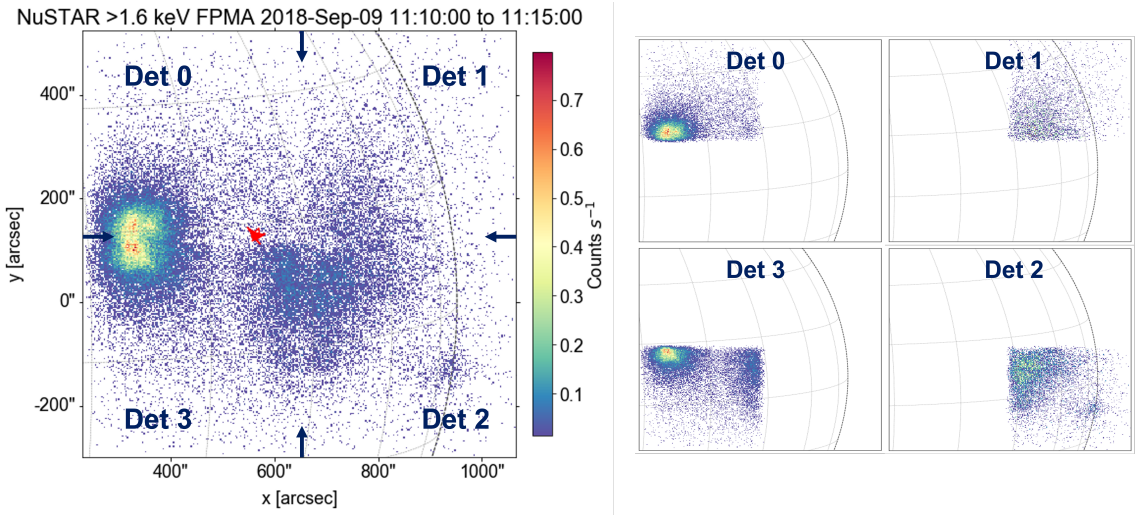

Figure 2.2: A NuSTAR FPMA >1.6 keV count rate image of AR12721 on 2018 September 9 (see Chapter 3 and 4) integrated between 11:10–11:15 UTC (left panel). The detectors are labelled in navy and the chip-gap is indicated with arrows of the same colour. The position of the optical axis over the integration time is plotted in red. To make the positions of the detectors and the chip-gap clearer, the same emission and field of view is also plotted per detector (right panels).

NuSTAR's pointing is controlled with three Camera Head Units (CHUs), or star tracker heads, aimed in orthogonal directions with respect to the spacecraft that track distant stars (Harrison et al. 2013). Throughout one solar observation the CHU combination used for the pointing changes, when one gets blinded by the Earth for example, potentially causing drastic changes in the observed emission if the source is close to the chip-gap or detector edge. Consequently, careful consideration must be taken when comparing data even from the same orbit. The solar pointing accuracy of NuSTAR is limited by the thermal motions of the observatory and can have a pointing uncertainty of ~1.5′ (Grefenstette et al. 2016). There is a fourth CHU that faces in the direction of the observed object that is mounted on the instrument optics bench (labelled in Figure 2.1); however, this is blinded by the Sun during solar observations. Therefore, NuSTAR often has to be co-aligned to data that has better pointing accuracy, e.g. SDO/AIA.

Another optical aspect to take note of is the position of the optical axis (OA), shown in Figure 2.2 (left panel, red), which is nominally located approximately 60″ into DET-ID 0 from the FOV centre (NuSTAR Observatory Guide[3]). However, the OA will move due to thermal motions of the mast during observations, which unsurprisingly becomes worse when observing the Sun. The thermal conditions cause the alignment between the optics and the focal bench to change and potentially moving the OA location by up to 60″ (Harrison et al. 2013). The metrology instruments, along with the focal plane mounted CHU (labelled in Figure 2.1) are used to help correct for the OA translational

[3] https://heasarc.gsfc.nasa.gov/docs/nustar/NuSTAR_observatory_guide-v1.0.pdf

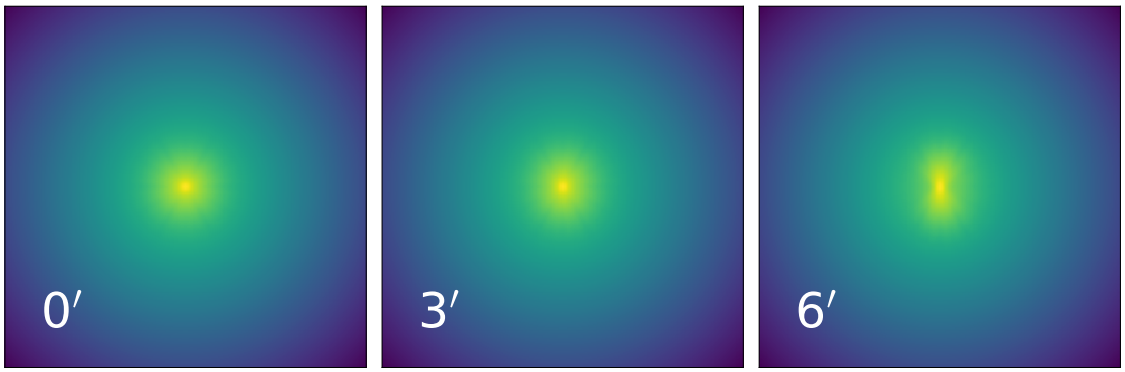

Figure 2.3: NuSTAR FPMA point spread functions at 0′ (left panel), 3′ (middle panel), and 6′ (right panel) from the optical axis.

motion; however, as stated, when observing the Sun these corrections are more uncertain as this CHU is unreliable. The position of the OA provides us with valuable context when discussing NuSTAR's point spread function.

Point Spread Function

Figure 2.3 shows NuSTAR's point spread function (PSF[4]) as it changes shape with radial distance from the OA (Madsen et al. 2015). The orientation of the PSF also changes with azimuthal angle such that the vertical feature in the rightmost panel (Figure 2.3) remains orthogonal to the line drawn between its centre and the OA. Therefore, the position of the OA is required to utilise the correct NuSTAR PSF as emission detected further away will be stretched more at right angles to the optical axis.

NuSTAR has an angular resolution of 18″ due to its PSF (Harrison et al. 2013); therefore, we use deconvolution as described in Section 1.4.1, and performed in Chapter 3 and 4, in order to recover the image with the PSF effects removed. This can help reveal detail in the X-ray image and allows for an easier spatial comparison between the NuSTAR X-ray images and those obtained from higher resolution instruments.

2.1.2 Instrument Livetime

As X-ray photons are detected by the FPM detectors their time, location, and energy are recorded in a triggered readout format, where a detected photon is called an event (Harrison et al. 2013). For every event, NuSTAR has a deadtime period of 2.5 ms during which another detection cannot be

[4]The function which describes how a point-like source is altered in a final image due to the instrument's optics.

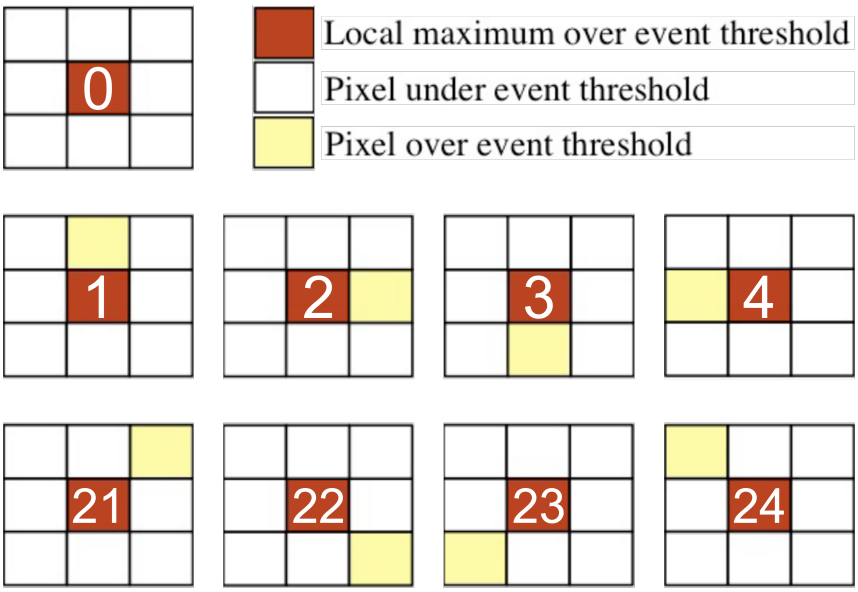

Figure 2.4: The definition of a select number of NuSTAR grades. This describes the pixels that record a photon detection for nine scenarios. The top left indicates when one pixel is found to record the event (grade 0), the middle row shows when the photon is split over two contiguous pixels (grades 1–4), and the bottom row shows when two pixels diagonal to each other detect an event (grades 21–24) which are deemed to be non-physical (Grefenstette et al. 2016). These panels have been taken from The NuSTAR Data Analysis Software Guide (Perri et al. 2017)[6].

made. This means that NuSTAR's throughput is limited to 400 counts s^{-1} FPM^{-1}, which is suitable for astrophysical sources but can pose challenges when observing the Sun as it becomes easier for faint sources to be missed if another region is much brighter.

From the concept of deadtime, we define the FPM livetime. The livetime is the length of time within a given period that is not deadtime due to a photon detection[5]. Therefore, the event rate is given by

$$Rate = \frac{events}{livetime} \quad [\text{events s}^{-1}], \tag{2.1}$$

since the FPM is only open to photon detection during livetime periods (Bhalerao 2012).

As photons reach the FPMs, there are a number of ways they can be detected by the pixels. Figure 2.4 shows a number of these detection scenarios where a single pixel event is labelled as grade 0. If two contiguous pixels are above the detection threshold for the event then these are labelled grade 1–4 and if two diagonal pixels are involved then these are grade 21–24. Note that grades 21–24 are non-physical (Grefenstette et al. 2016) since if a photon were to interact with these two pixels it should have also been detected by a contiguous pixel.

[5]Livetime is also discussed in its fractional form where it is the fraction of a given time period not spent as deadtime.
[6]https://heasarc.gsfc.nasa.gov/docs/nustar/analysis/nustar_swguide.pdf

When working with NuSTAR solar data, we can easily reach livetimes $\ll$1% which means that pile-up can become an issue. In NuSTAR, pile-up can occur either by two photons being detected as one in the same pixel by the on-board electronics or by two photons being classified as a single 'split-pixel' event by the post-processing software (Grefenstette et al. 2016). The former scenario can be investigated by comparing the pile-up viewed from the brightest astrophysical source, Sco X-1, which is effectively a point source with emission from the Sun where Grefenstette et al. (2016) found that this type of pile-up can be neglected. The latter, simpler style of pile-up is solved by removing any event grades $>$0 from the analysis. As pile-up affects the higher energy events more than the lower ones[7], correcting for pile-up will steepen the slope of the spectrum at these higher energies.

2.1.3 Ghost-rays and Stray Light

In addition to bright sources within the NuSTAR FOV taking a significant fraction of the throughput, sources outside the FOV can also have a measurable effect. One way for this to occur is via *ghost rays*. Ghost rays are photons which only reflect once off the optics, where a properly focussed photon will reflect twice (Madsen et al. 2015; Grefenstette et al. 2016). These photons can increase the background emission drastically and, if the source is bright enough, cover NuSTAR's full FOV with its emission.

Another source of photons from outside the field of view comes in the form of *stray light*. Stray light originates from the areas of the sky that are not obscured by the physical body of NuSTAR (Bhalerao 2012; Harrison et al. 2013). However, stray light does not greatly impact solar observations due to the brightness of the Sun in comparison to other sources. An event may be detected from a cosmic ray strike; however, these detections are vetoed by the readout electronics onboard NuSTAR (Grefenstette et al. 2016).

2.1.4 Data Pipeline

Once NuSTAR observation files have been downloaded from the archive[8] they must be pre-processed with the NuSTAR data pipeline, `NuSTARDAS`, within HEASoft which is described in The NuSTAR Data Analysis Software Guide (Perri et al. 2017). For this we use the `nupipeline` routine which takes Level 1 files, those containing the original telemetry information, and processes them to Level 3 which contain all the same data but now with added calibration information (Perri et al. 2017).

[7] The lower the photon energy, the lower the chance to cause two or more pixels to be above the detection threshold.
[8] The NuSTAR Master Catalog.

Calibration of the Data

The calibration data is found using the calibration database `CALDB`[9]. During this initial stage of processing the information from the laser metrology system can be used to correct for the thermal motions of the mast (Section 2.1.1). Next, the attitude information of NuSTAR is processed from the CHU data followed by the cataloging of bad and hot pixels. The event grades are assigned (Figure 2.4) which involves correcting detected photon energies if the photon was split over different pixels. Finally, gain corrections are made which converts the charge associated with a photon's detection (electronic units) to more physical energy units and the sky position of each count is calculated (Perri et al. 2017). The final product is a Level 1a file with a list of all detected counts, each with their energy, position, grade, and time of detection.

Filtering the Data

The pipeline can then be used to screen the event files with respect to certain observational and instrumental parameters, taking Level 1a products and outputting Level 2 files (Perri et al. 2017). This stage can be used to remove unwanted pixels and filter the data with respect to attitude and orbital conditions such as for South Atlantic Anomaly (SAA) passage or occultation by the Earth and its atmosphere. The data can also be filtered to a specific time range or spatial region with Good Time Interval (GTI) and Region (REG) files, respectively.

High Level Data Product Extraction

The final processing step involving the NuSTAR data pipeline discussed here is for the creation of Level 3 data products. These are high level data products which include time profiles, spectral files, and image extraction (Perri et al. 2017). From this step we can obtain the Pulse Height Amplitude (PHA) file, Redistribution Matrix File (RMF), and the Ancillary Response File (ARF) that are all needed for spectral analysis (see Section 6.1). The PHA file contains the count information from the chosen files obtained from the filtering process. The RMF and ARF files dictate how the incoming photon flux was converted to a count spectrum with the former detailing the photon-to-count conversion and the latter containing the effective area information.

Figure 2.5 shows NuSTAR's effective area (left panel), redistribution matrix (middle panel), and resulting spectral response matrix (SRM, right panel) for a given observation (microflare 3's pre-flare time in Section 3.4.1). From the left panel we see at low energies, <10 keV, NuSTAR's response

[9] https://heasarc.gsfc.nasa.gov/docs/heasarc/caldb/caldb_intro.html

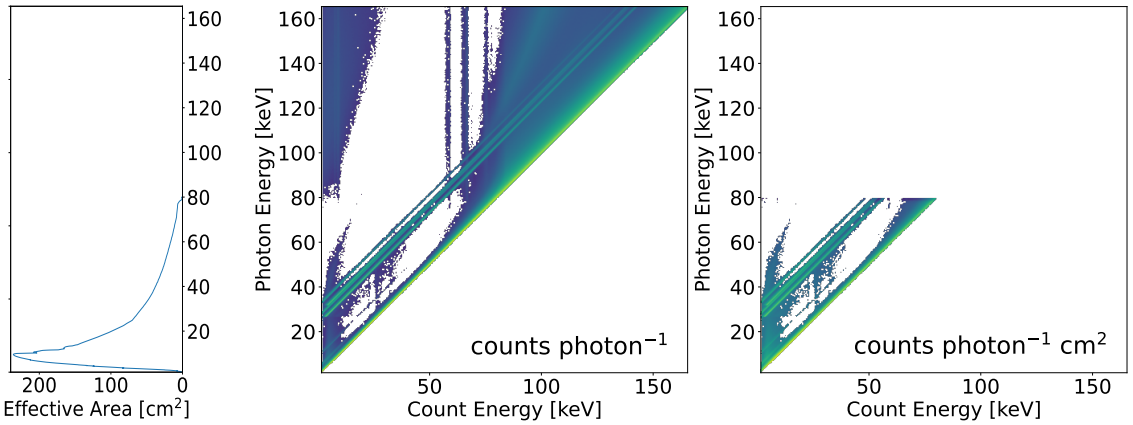

Figure 2.5: NuSTAR effective area (left panel), redistribution matrix (middle panel), and spectral response matrix (SRM, right panel) for Section 3.4.1's microflare 3 pre-flare time. The SRM is created by multiplying the redistribution matrix rows by the effective area at the corresponding photon energy. Note that the y-axis for all panels are shared.

decreases due to the thickness of the Platinum (Pt) on the detector surface (Madsen et al. 2015) with a hard cut-off at 80 keV due to a Pt K-shell absorption edge (Bhalerao 2012). The effective areas are then multiplied along the corresponding photon energy rows in the redistribution matrix (middle panel) in order to create the SRM which will be discussed further in Section 2.1.5. The redistribution matrix shows the strongest conversion between photons and counts along the diagonal, the photopeaks in each photon energy spectrum, with the off-diagonal terms originating from backscatter peaks and compton edges from the X-ray photons interacting with other parts of the observatory and not just the detectors.

Solar Specific Processing

An additional processing step is required for solar observations where we convert the R.A./Dec. coordinates in the event count-list to solar helioprojective coordinates. This is performed with the NuSTAR solar Python repository[10]. Images can then be produced by spatially and temporally binning the new event list in solar coordinates, allowing for an easier comparison with images from other solar observations (Grefenstette et al. 2016). The NuSTAR images can then be co-aligned using other solar observatories, helping with NuSTAR's pointing uncertainties (Chapter 3 and 4).

[10] https://github.com/NuSTAR/nustar_pysolar

2.1.5 Instrument Spectral Response

The spectral response matrix (shown in Figure 2.5, right panel) dictates how an incoming photon flux is converted to a count spectrum. Therefore, this is crucial information when analysing NuSTAR's observed X-ray spectrum which is discussed further in Chapter 5 and more specifically in Section 5.1. We use the SRM to forward-fit (Chapter 5 and 6) parametric models (Section 1.4.4 and 1.4.5) to the NuSTAR count spectrum in order to determine the electron distribution properties that best explain the observation. The conversion from a photon spectrum to a count spectrum is shown in Figure 5.1, Section 5.1.

As well as optimising over model parameters in the forward-fitting process, it may be necessary to also investigate spectral response parameters, such as the gain discussed in Section 2.1.4. For example, Duncan et al. (2021) found evidence that the gain in the initial calibration steps is not characterised correctly for very high count, low livetime ($\ll$1%) regimes when observing the Sun. This is discussed briefly in the following section.

Spectral Response Gain Correction

As discussed in Section 2.1.4, the gain describes how the measured photon energies are mapped to count energies; therefore, when applying a change to the gain, we simply change the energies that correspond to the count spectrum bin channels. The new energies (E_{new}) for a model count spectrum then becomes

$$E_{new} = \frac{E_{original}}{G_{slope}} - G_{offset} \quad \text{[keV]}, \tag{2.2}$$

where $E_{original}$ is the array of original count channel energies for the model spectrum and G_{slope} & G_{offset} are the gain parameters[11]. The default G_{slope} & G_{offset} values are 1 and 0 keV when calculating a count model with the SRM (Section 2.1.5 and 5.1) which causes no change to the original gain calibration from Section 2.1.4 with $E_{new} = E_{original}$.

Figure 2.6 shows an example of how the 6.7 keV Fe complex (Section 1.4.2) is moved with relation to the gain parameters. The default count model spectrum ($G_{slope} = 1$ and $G_{offset} = 0$ keV) is shown in blue. The orange line shows the count spectrum with $G_{offset} = 0.1$ keV, the green with $G_{slope} = 0.9$, and the red count spectrum has both gain parameters changed with $G_{offset} = 0.1$ keV and $G_{slope} = 0.9$.

Duncan et al. (2021) discusses a variation in the NuSTAR gain that is specific to solar data analysis for observations with low livetime (<1%). This is most obvious in the inconsistent location of the 6.7 keV

[11] https://heasarc.gsfc.nasa.gov/xanadu/xspec/xspec11/manual/node34.html

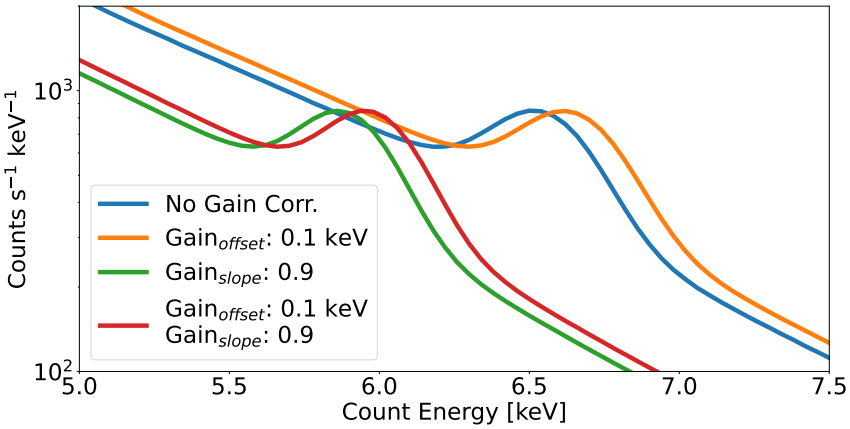

Figure 2.6: NuSTAR count model spectrum calculated with different gain parameter values. The default spectrum (blue, $G_{slope} = 1$ and $G_{offset} = 0$ keV) is shown, with the orange spectrum showing the result with values $G_{slope} = 1$ and $G_{offset} = 0.1$ keV, the green showing $G_{slope} = 0.9$ and $G_{offset} = 0$ keV, and the red showing values of $G_{slope} = 0.9$ and $G_{offset} = 0.1$ keV (see text and Equation 2.2 for more detail). This plot has been produced using the thermal model and gain correction methods in the Python spectral fitting package Sunxspex (Section 5.4).

line feature (Section 1.4.2) in the observed spectrum compared to the computed count model. The observed 6.7 keV Fe line complex is found to be centred at ~6.4 keV, a difference much greater than the quoted gain uncertainty of 0.04 keV (Madsen et al. 2015).

Alternatives that could explain the ~6.4 keV line feature were explored, for example Doppler shifts or Fe-Kα emission (Emslie et al. 1986). However, the Doppler shift needed for a difference of ~0.3 keV is ~10^4 km s^{-1}, far larger and in the opposite direction compared to plasma velocities observed during chromospheric evaporation (Section 1.2.1) that have upflow velocities of ~400 km s^{-1} (Antonucci et al. 1984), and the presence of a Fe-Kα 6.4 keV feature does not explain the absence of the 6.7 keV Fe complex. Therefore, a gain discrepancy for solar spectra was concluded. Including a gain correction in the spectral analysis also better fitted the Ca XIX ~3.9 keV line feature observed with NuSTAR (Duncan et al. 2021). We, therefore, utilise a gain correction in the spectral analysis presented in Section 6.1.3, 6.1.4, and 8.1.2.

In addition to the spectral response of NuSTAR which dictates the photon-to-count conversion, we can also calculate NuSTAR's temperature response from its SRM. We discuss NuSTAR's temperature response in Section 2.1.6.

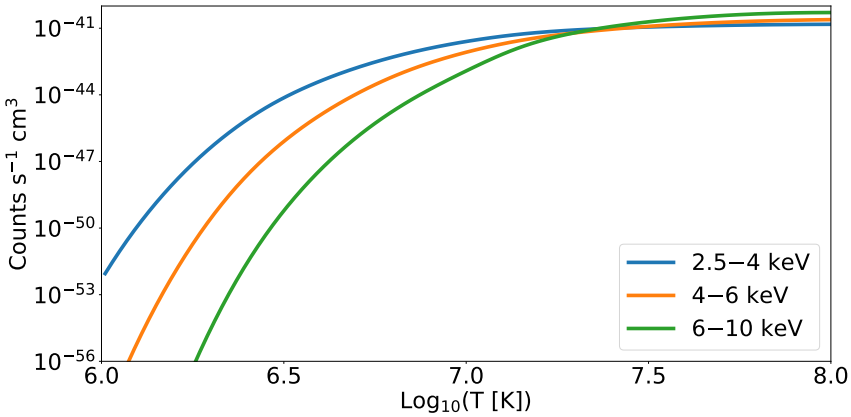

Figure 2.7: NuSTAR temperature responses for three energy ranges, 2.5–4 keV (blue), 4–6 keV (orange), and 6–10 keV (green). This plot has been produced using the thermal model from the Python spectral fitting package Sunxspex (Section 5.4).

2.1.6 Instrument Temperature Response

The temperature response of an instrument describes the expected emission for a range of isothermal sources integrated over a unit volume. For example, Figure 2.7 shows NuSTAR's temperature response for three different energy range bands, 2.5–4 keV, 4–6 keV, and 6–10 keV. NuSTAR's temperature response is monotonic and has more sensitivity at higher temperatures for all energy ranges, with the higher energy range response (6–10 keV, green) being less sensitive to cooler temperatures than the lower energy range response (2.5–4 keV, blue).

The temperature responses in Figure 2.7 are calculated from the SRM (Section 2.1.5) shown in the right panel of Figure 2.5 by obtaining emission measure normalised thermal model count spectra. Each temperature response ($R(T)$) for a given source temperature (T) is found using

$$R(T) = \int_{E_{low}}^{E_{high}} \frac{I(T, EM)SRM(E)}{EM} dE \quad [\text{counts s}^{-1}\ \text{cm}^3], \tag{2.3}$$

where $I(T, EM)$ is the thermal photon model spectrum, EM is the emission measure (Section 1.4.4), and E_{low} & E_{high} represent the count energy range for the response. Equation 2.3 is then performed for a range of temperatures where an energy range is chosen, this produces the NuSTAR temperature responses shown in Figure 2.7.

With the temperature responses calculated, we can then sum over the same energies in the observed spectrum to obtain the total count rates in each range. Dividing the total count rates by the temperature responses then produces *loci-curves* which can be used to estimate the emission measure and temperature of the source, if isothermal (Section 5.5.1 of Phillips et al. 2012), or indicate the

multi-thermal nature of the emitter where the loci-curves can then be used to obtain the differential emission measure (DEM introduced in Section 1.4.4 while detailed and calculated in Section 4.5). Wright et al. (2017) and Cooper et al. (2020, see Chapter 4) use loci-curves to calculate the DEM of microflare emission finding NuSTAR's sensitivity to the higher temperature range invaluable.

2.2 The Geostationary Operational Environmental Satellite

The Geostationary Operational Environmental Satellite (GOES) is a series of satellites beginning in 1974 with a new satellite in operation approximately every three years (Simões et al. 2015). GOES-16 and GOES-17 are the latest satellites and are used to monitor both the Earth and the Sun continuously with GOES-18 being launched on 2022 March 1[12]. In this thesis we focus on the X-ray Sensor (GOES/XRS) which detects solar soft X-rays in the 0.5–3 Å and 1–8 Å wavelength ranges[13]. The longer wavelength channel (1–8 Å) is used for the GOES flare classification scale (Table 1.1).

The NuSTAR microflares discussed throughout Chapter 3 are too weak to be reliably observed by GOES; therefore, we utilise the SolarSoft IDL routine `goes_flux49.pro`[14] which assumes coronal abundances by default and uses `CHIANTI V7.1` (Dere et al. 1997; Landi et al. 2013). This routine takes a temperature and emission and returns the expected GOES flux, or synthetic flux (see Section 1.4.4), from such an isothermal source. From this we compare to Table 1.1 and obtain the microflare's GOES *equivalent* class. The `CHIANTI` version being used by the routine is not the most recent one; however, the GOES equivalent class calculated in this way is only used for relative comparison purposes between microflares and not for any quantitative analysis.

Note that recent calibration on the GOES satellite series has revealed the need for a scaling factor for GOES satellites 1–15[15]. Therefore, all short (0.5–3Å) and long (1–8Å) channel X-ray fluxes are to be divided by 0.85 and 0.7, respectively. To maintain consistency, we primarily quote the same GOES or GOES equivalent class for each microflare as stated in the relevant published works; however, the new scaled class will also be presented if the value changes in the same number of significant figures. For example, a class of A0.1 remains the same but A0.2 would be scaled to A0.3 which would be indicated.

[12] https://spacenews.com/atlas-5-launches-goes-t-weather-satellite/
[13] https://www.swpc.noaa.gov/products/goes-x-ray-flux
[14] https://hesperia.gsfc.nasa.gov/ssw/gen/idl/synoptic/goes/goes_flux49.pro
[15] https://hesperia.gsfc.nasa.gov/rhessidatacenter/complementary_data/goes.html

2.3 The Reuven Ramaty High-Energy Solar Spectroscopic Imager

The Reuven Ramaty High-Energy Solar Spectroscopic Imager (RHESSI; Lin et al. 2002) launched in 2002 and uses rotating modulation collimator techniques to image solar flares in X-rays and gamma-rays (Hurford & Curtis 2002; Hurford et al. 2002). RHESSI is sensitive to an energy range of 3 keV–17 MeV and used nine germanium detectors and nine pairs of differently pitched grids for collimation.

We investigate a RHESSI observation in Section 6.2.1 where we perform similar analysis for a six-detector summed spectrum of a solar flare described in Fletcher et al. (2007) but with the new Python spectral fitting software Sunxspex (Section 5.4 and Chapter 6). Furthermore, using Sunxspex, we are able to perform simultaneous fitting on all six separate spectra for the first time (Section 6.2.2).

2.4 Solar Orbiter

Solar Orbiter (Müller et al. 2020) is a mission lead by the European Space Agency and was launched on 2020 February 10. The observatory contains ten instruments that will take in-situ measurements of the solar wind in the heliosphere and high resolution observations of the Sun. Solar Orbiter aims to investigate the driving mechanism for the solar wind, the origin of the coronal magnetic field, the role that solar transient sources have in heliospheric variability, and other solar-related phenomena while observing the Sun at a close distance, up to 0.28 au (Zouganelis et al. 2020).

In Section 2.4.1, we will briefly introduce the Spectrometer/Telescope for Imaging X-rays (Krucker et al. 2020), one of the remote observing instruments on Solar Orbiter. We investigate its spectral data in Section 6.2.3 and 6.3.1.

2.4.1 The Spectrometer/Telescope for Imaging X-rays

The Spectrometer/Telescope for Imaging X-rays (STIX; Krucker et al. 2020) is a similar instrument to RHESSI in that it uses an indirect imaging technique with 32 coarsely-pixelised (CdTe) detectors, each measuring hard X-rays behind a pair of collimating grids. STIX is sensitive to an energy range of 4–150 keV and has a variable time resolution for active and quite solar activity. Since 2020, STIX has observed numerous microflares at the Sun from GOES class A to B (see Figure 1.4).

In Section 6.2.3, we investigate an A class microflare which is at the lower limits of STIX's sensitivity using the aforementioned Sunxspex spectral fitting package, the STIX microflare was also observed

at the upper limits of NuSTAR's sensitivity. In Section 6.3.1, we simultaneously fit both spectra from STIX and NuSTAR for the first time.

2.5 The Solar Dynamic Observatory

The Solar Dynamic Observatory (SDO; Pesnell et al. 2012) was launched on 2010 February 11 and is composed of three instruments, the Atmospheric Imaging Assembly (Lemen et al. 2012), the Helioseismic and Magnetic Imager (Schou et al. 2012), and the Extreme Ultraviolet Variability Experiment (Woods et al. 2012). The SDO instruments have observed the Sun almost continuously since 2010 providing a crucial UV/EUV and magnetic field context to many solar phenomena and observations made by other observatories.

In Section 2.5.1, the Atmospheric Imaging Assembly is briefly discussed with reference to the instrumental temperature responses. The Helioseismic and Magnetic Imager is then introduced in Section 2.5.2, describing the data products used in later sections (Section 3.9).

2.5.1 The Atmospheric Imaging Assembly

SDO's Atmospheric Imaging Assembly (SDO/AIA; Lemen et al. 2012) provides two UV and seven EUV full-disk images of the Sun with a spatial resolution of 1.5″ and cadence down to 12 s. Six of the EUV telescopes have wavelength band passes centred on various Fe ion lines (from Fe VIII to Fe XXIV) and are called the 94 Å, 131 Å, 171 Å, 193 Å, 211 Å, and 335 Å channels. The other EUV channel, 304 Å, is centred on a He II line (Lemen et al. 2012). The two UV channels are stated to observe the continuum around 1600 Å and 1700 Å with the former also observing C IV; however, Simões et al. (2019) shows that many lines contribute to these channels.

SDO/AIA observations are used throughout Chapter 3, 4, and 8 to provide a higher degree of spatial context when analysing NuSTAR microflares. In addition, we also utilise the temperature responses of the SDO/AIA channels (Figure 2.8), along with NuSTAR temperature responses (see Section 2.1.6), when performing DEM analysis (see Section 4.5). The temperature response ($R(T)$) for a given channel is calculated using

$$R(T) = \int_0^\infty G(T,\lambda)W(\lambda)d\lambda \quad [\mathrm{DN\ s^{-1}\ pix^{-1}\ cm^5}], \tag{2.4}$$

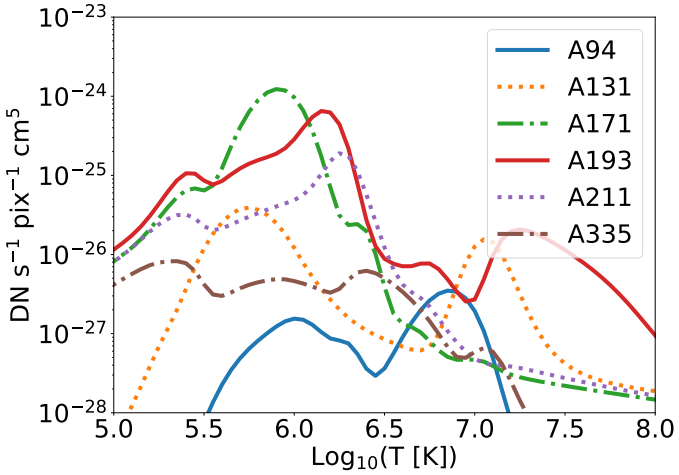

Figure 2.8: SDO/AIA temperature responses for the 94 Å, 131 Å, 171 Å, 193 Å, 211 Å, and 335 Å channels.

where $G(T, \lambda)$ is the contribution function (see Equation 1.11) and $W(\lambda)$ is the wavelength response of the chosen channel (Boerner et al. 2012). The SDO/AIA responses in Figure 2.8 are obtained using the SolarSoft IDL routine `aia_get_response.pro`[16].

Note that the units for the SDO/AIA temperature responses in Equation 2.4, DN s^{-1} pix^{-1} cm^5, are different to NuSTAR's in Equation 2.3, counts s^{-1} cm^3. Taking into account the observable units for each instrument, DN s^{-1} pix^{-1} for SDO/AIA and counts s^{-1} for NuSTAR, the temperature responses differ by an area factor (cm^2) where the area should be the emitting region. Therefore, the area of emission deduced from observations is important when using both SDO/AIA and NuSTAR temperature responses in the same analysis.

Since the SDO/AIA channels have a response to multiple lines (O'Dwyer et al. 2010), even though they are centred on a dominant one, *proxy* channels can be created through a linear combination of several responses in order to further isolate a desired line feature. This is discussed in the next section where we use a combination of the 171 Å and 211 Å channel emission to isolate the higher temperature emission in the 94 Å channel, removing the lower temperature response peak visible in Figure 2.8 between $\log_{10}$(T)~5.5–6.6 (blue line).

The Fe XVIII Proxy Channel

From the SDO/AIA temperature response, the 94 Å channel has a significant response to material at 4–10 MK ($\log_{10}$(T)~6.6–7, Figure 2.8 and 2.9, blue line). Sensitivity in this range is crucial with regards to microflares. This is corroborated by the spectral analysis of many observed NuSTAR

[16] https://hesperia.gsfc.nasa.gov/ssw/sdo/aia/idl/response/aia_get_response.pro

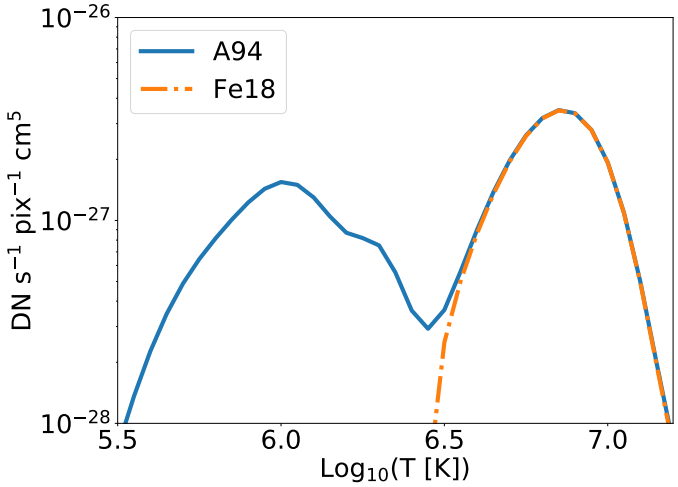

Figure 2.9: SDO/AIA temperature responses for the 94 Å and Fe XVIII proxy channels.

microflares (Hannah et al. 2016; Wright et al. 2017; Hannah et al. 2019; Glesener et al. 2020; Duncan et al. 2021), and also with respect to the temperatures typically observed in hot AR cores of ~4 MK and flaring coronal plasma (O'Dwyer et al. 2010; Lemen et al. 2012; Warren et al. 2012).

The 94 Å wavelength channel has a significant response from Fe XVIII emission; however, it also has a significant sensitivity to other emission lines formed at lower temperatures (Lemen et al. 2012). An estimate of the non-Fe XVIII line contribution (continuum + Fe VIII–Fe XIV) from an example AR is determined from differential emission measure analysis using Hinode's EUV Imaging Spectrometer (Hinode/EIS; Kosugi et al. 2007; Culhane et al. 2007). A linear combination of the 171 Å and 211 Å channels is then used as a proxy for the "non-Fe XVIII line contribution" and subtracted from the 94 Å channel detected emission to isolate the hotter Fe XVIII component (Del Zanna 2013).

To help compare the X-ray data to the available SDO/AIA EUV observations and provide more spatial detail, we utilise an Fe XVIII proxy. The proxy takes the form

$$I_{\text{Fe XVIII}} \approx \frac{I_{94\,\text{Å}}}{\tau_{94\,\text{Å}}} - \frac{I_{171\,\text{Å}}}{450\tau_{171\,\text{Å}}} - \frac{I_{211\,\text{Å}}}{120\tau_{211\,\text{Å}}} \quad [\text{DN s}^{-1}\ \text{pix}^{-1}], \tag{2.5}$$

where $I_{\text{Fe XVIII}}$ is the emission from the Fe XVIII proxy channel with $I_{94\,\text{Å}}$, $I_{171\,\text{Å}}$, and $I_{211\,\text{Å}}$ being the measured 94 Å, 171 Å, and 211 Å channel emission, respectively (Del Zanna 2013). Note that τ_i is the corresponding correction factor to account for channel i's degradation from launch. This isolates the 94 Å response between 4–10 MK ($\log_{10}$(T)~6.6–7) as shown in Figure 2.9 (orange dash-dotted line).

2.5.2 The Helioseismic and Magnetic Imager

SDO's Helioseismic and Magnetic Imager (SDO/HMI; Schou et al. 2012) studies both the solar interior and surface by measuring Doppler shifts, intensities, and vector magnetic field of the Fe I 6173 Å absorption line (Norton et al. 2006). SDO/HMI is able to create full-disk maps with a cadence of 45 s with a spatial resolution of 0.91″ (Schou et al. 2012). Using SDO/HMI, in combination with SDO/AIA and other instruments like NuSTAR, we are able to map solar events through the atmosphere and investigate any corresponding behaviour that might be present in the photosphere or photospheric magnetic field.

In Section 3.9, we investigate the presence of mixed polarity photospheric magnetic flux at, or close to, the apparent footpoints of eight microflares. Using SDO/HMI 45 s cadence magnetograms, we are able to obtain flux cancellation rates at the footpoint locations of the two brightest and most complex microflares of Chapter 3 where consistency is found with other microflare studies. In addition, the investigation of SDO/HMI, and SDO/AIA, data products remains the subject of future work with regards to those microflares discussed in Chapter 8.

3

NuSTAR Observations of a Repeatedly Microflaring Active Region

The work in this chapter is based on the scientific research presented in Cooper et al. (2021).

NuSTAR has been enriching a novel area of solar X-ray observations at incredibly weak scales since 2014 (Grefenstette et al. 2016; Hannah et al. 2016). In this chapter, we describe analysis on a NuSTAR solar observation campaign, performed in 2018 September, and present events that are not easily observed even with solar dedicated observatories. Therefore, we can search for high temperature or non-thermal emission in very small microflares, probing the physics involved at these weak flare scales.

First we discuss a number of microflares that were captured by NuSTAR in September 2018 in Section 3.1 and subsequent sections, the subject of (Cooper et al. 2021). We then delve deeper into the weakest microflare of these data in Chapter 4, the focus of Cooper et al. (2020).

3.1 NuSTAR Solar Observation Campaign: 2018 September 9–10

NuSTAR observed the Sun on 2018 September 9–10 where 6 hour-long dwells were performed, co-ordinated to study a region with the FOXSI-3 sounding rocket on September 7 (Musset et al. 2019). However, an active region, AR12721, emerged on September 8 dominating NuSTAR's field of view (FOV). The original FOXSI-3 target still produced visible X-ray emission but was too weak relative to AR12721 to be studied. NuSTAR observed numerous X-ray microflares across both days from AR12721, with only the brightest two events being detected by GOES/XRS. All other NuSTAR microflares were far too weak to appear in this whole-Sun integrated light curve. When analysing this NuSTAR data-set, we are only able to use additional observations from SDO as other common

solar dedicated observatories were either focussed on the FOXSI-3 region or not observing during that time.

Due to AR12721 not being the initial target, the pointing of NuSTAR was not optimal. Therefore, some microflares became obscured by the detector chip-gaps, only allowing analysis using data from one FPM. Additional difficulty arose from the various CHU state changes throughout each hour-long observation, sometimes further obscuring the detection of X-ray emission from potential events, moving the most prominent emission across the chip-gap. Therefore, we took careful consideration when choosing whole microflare times if there was a CHU change and any FPM that recorded an event but displayed behaviour indicative of the above issues was discounted. In addition to chip-gap proximity issues, some events were truncated when occurring close to NuSTAR's eclipse or South Atlantic Anomaly (SAA) passage.

To aid in locating the position and finer scale features of the microflares observed, we obtain 12 s cadence EUV image data from SDO/AIA (Section 2.5.1) approximately over a 4.5 h period on each day centred on the middle time of dwell 2 and 5. Figure 3.1 shows the EUV light curves for the 94 Å, 131 Å, 171 Å, and 211 Å channels in panel a and f for September 9 and 10, respectively. Comparing these EUV time evolutions of AR12721 to their co-temporal $>$2.5 keV NuSTAR X-ray light curves from both FPMA & B, panels c and h, we only see similar behaviour in the 94 Å for the two brightest microflares, the others show no corroboratory features. The SDO/AIA and NuSTAR light curves are produced from the boxed regions indicated in panel d and e for September 9 and panel i and j for the following day. The NuSTAR FPMA FOV images show how much AR12721 dominates over the initial FOXSI-3 target.

Figure 3.1, panel b and g, shows the Fe XVIII channel from Equation 2.5 (Section 2.5.1) integrated over the same region as the other native SDO/AIA channels. When comparing these to the X-ray time profiles (panel c and h) it is clear that the hotter contribution to the 94 Å shows incredible agreement with NuSTAR, not only providing information inferred from the emission observed itself (e.g., temperature range, synthetic flux estimations) but also a reliable means to quickly and easily locate the microflaring structure such that analysis for all SDO/AIA channels can be focussed on these sites. In addition, as discussed in Section 2.1, NuSTAR has a pointing uncertainty $<1.5'$ during solar observations (Grefenstette et al. 2016). Since we clearly find a significant similarity in the emission observed by NuSTAR and the Fe XVIII proxy, we also use this synthetic EUV channel to co-align the X-ray image data. We then apply a single shift to the NuSTAR X-ray map per event where any spatial shift in X-ray source, either in time or in energy, can be investigated. Due to NuSTAR's relatively low spatial resolution, we still assume a conservative co-alignment uncertainty of $\sim 10''$.

The NuSTAR (grade 0, FPMA+B, $>$2.5 keV counts) time profiles in Figure 3.1 show the presence of 10 different microflaring events (summarised in Table 3.1). All of these weak X-ray microflares were

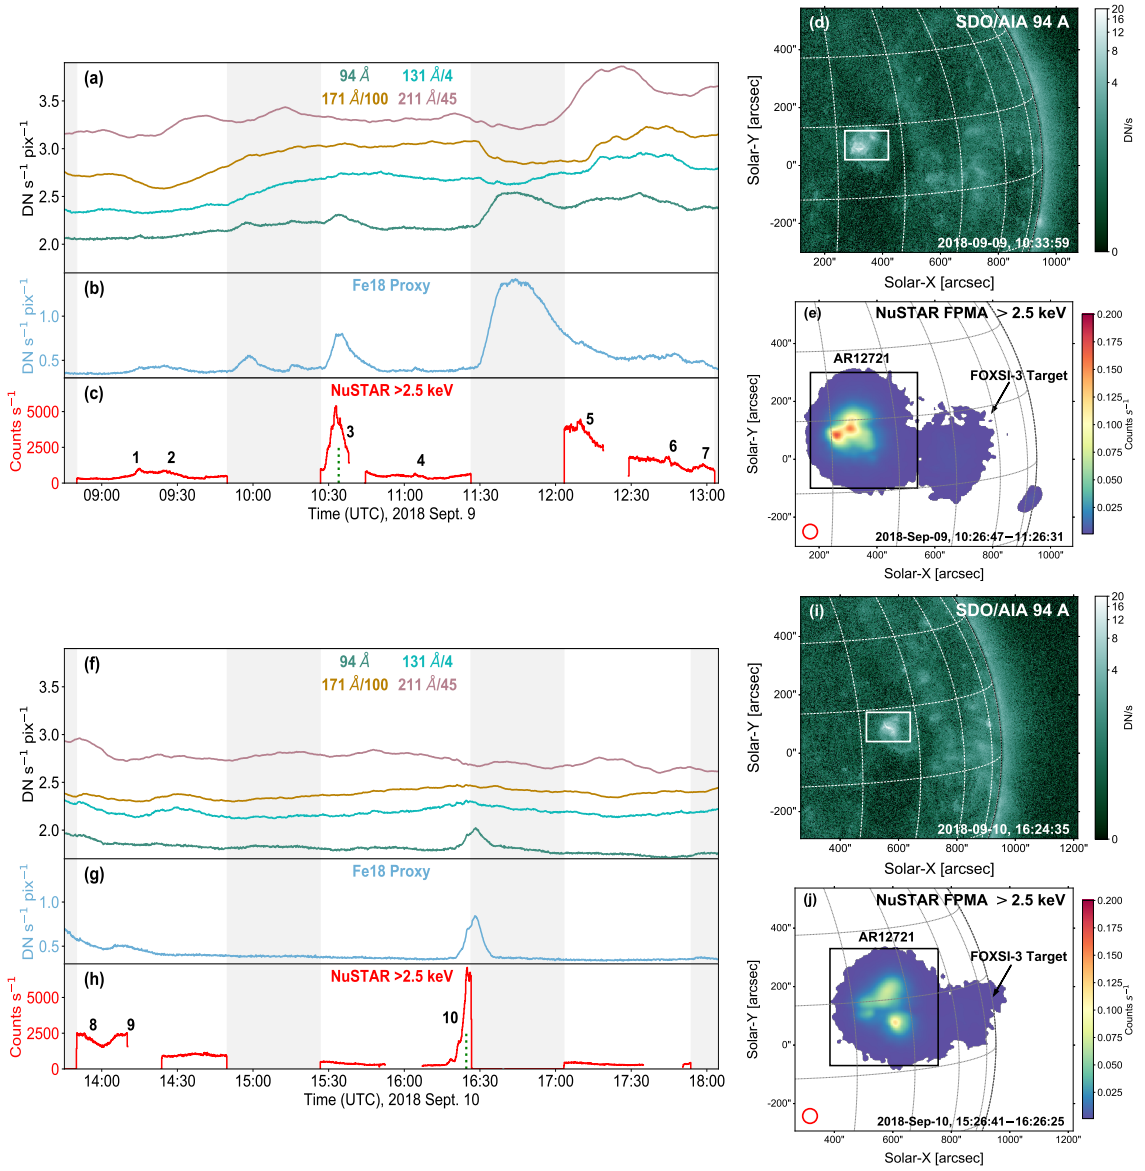

Figure 3.1: NuSTAR and SDO/AIA, including Fe XVIII proxy, time profiles from AR12721 on 2018 September 9 (panels a–c) and 10 (panels f–h) produced from the relevant boxed areas in panels d and e, and i and j, respectively. NuSTAR's eclipse is shown as shaded regions with other gaps in the data being caused by SAA passage. Constant y-axis limits are used across both day's time profile panels. Light curve channel names are displayed in each panel with any scaling factor shown. The vertical dotted green lines in panel c and h indicate the times used to produce the SDO/AIA 94 Å images in panel d and i, respectively. Each NuSTAR FOV image in panel e and j are created by integrating all FPMA grade 0 emission >2.5 keV over orbit 2 and 5. The 50″ diameter red circles in the bottom left of panel e and j indicate the minimum region size used for each microflare's area in spectral analysis. The numbers in panel c and h indicate identified microflares and the NuSTAR light curves and images are livetime corrected. This figure is taken from Cooper et al. (2021).

Table 3.1: NuSTAR event summary of all ten microflares observed on 2018 September 9–10. The 'NT' column indicates whether a particular microflare is observed to have direct evidence for non-thermal emission.

Obs. Date	Orbit No.	Microflare No.	Solar ID	Livetime (%)	Thermal Energy (Decade, [erg])	Direct NT Evidence
Sept. 9	1	1	SOL2018-09-09T09:15	14.3	10^{26}	No
	1	2	SOL2018-09-09T09:25	13.5	10^{26}	No
	2	3	SOL2018-09-09T10:33	6.0	10^{27}	Yes
	2	4	SOL2018-09-09T11:04	14.7	10^{26}	No
	3	5	SOL2018-09-09T12:10	4.4	10^{26}	No
	3	6	SOL2018-09-09T12:44	7.8	10^{26}	No
	3	7	SOL2018-09-09T12:58	9.0	10^{27}	No
Sept. 10	4	8	SOL2018-09-10T12:53[a]	4.4	10^{28}	No
	4	9	SOL2018-09-10T14:06	8.0	10^{28}	No
	5	10	SOL2018-09-10T16:28[b]	7.0	10^{28}	No

[a] Peak time was pre-NuSTAR, time from peak Fe XVIII

[b] Peak time was post-NuSTAR, time from peak Fe XVIII

observed to be present, if less obvious, in the Fe XVIII proxy channel upon comparison to the NuSTAR light curve. There does not seem to be any indication of significant volumes of plasma being heated to temperatures higher than what the Fe XVIII proxy is sensitive to in any of the ten microflares. This is observed from the SDO/AIA 131 Å channel (Figure 3.1, panel a and f) which is sensitive to the lower temperatures seen by the 171 Å and 211 Å channels in addition to plasma emission of ~10 MK (Lemen et al. 2012). Therefore, since the 131 Å light curve closely follows these lower temperature channels in evolution, it can be concluded that most emitting material is likely constrained to $<$10 MK. This suggests Fe XVIII is representative of the bulk of the hotter material observed by NuSTAR.

From Figure 3.1, it is clear that the Fe XVIII proxy qualitatively matches the observed X-ray emission, allowing the ten microflares to be corroborated in EUV emission. However, the emission evolution measured from two separate events in the proxy channel is not necessarily expected to be identical to NuSTAR. For example, microflare 3 and 10 have similar Fe XVIII levels but have noticeably different peak values in the X-ray regime. This could be explained with the different responses of instrument. The Fe XVIII channel has a peak response to material at log(T)~6.9 (7.1 MK) (O'Dwyer et al. 2010), whereas NuSTAR's response increases monotonically through this temperature region. Therefore, similar rates of increase in emission seen by Fe XVIII for two different events may not mean constant proportional changes are observed in the X-ray flux for those same events.

In this section, we present findings from the ten microflares observed by NuSTAR on 2018 September 9–10. We first discuss the analytical similarities of each microflare's investigation in Section 3.2. We then describe the analysis and problems more specific to each individual microflare/orbit while investigating their EUV and X-ray temporal, spatial, and spectral evolution in Section 3.3 to 3.8. In Section 3.9, we use 45 s cadence data from SDO/HMI to investigate the presence of mixed polarity photospheric magnetic flux in the proximity to microflare footpoints before discussing the final results in Section 3.10.

3.2 AR12721 X-ray Microflares

The ten microflares, shown in Figure 3.1, all vary in temporal and spatial complexity. Some microflares appear by themselves as one loop feature, as individual loops close to other bright structures, or some as a complex network of loops. This makes the analysis of each event vary in difficulty, particularly due to NuSTAR's spatial resolution. Dynamic or hot features that are spatially disparate in SDO/AIA can be indistinguishable with NuSTAR. Therefore, every event has to be investigated individually in order to ensure the emissions seen by both observatories are from the same comparable plasma, and not from disparate sources.

We create NuSTAR images by spatially binning the detected counts from the event list over chosen set of filters to the data using the NuSTAR solar Python repository[1]. These filters include choosing a time range, spatial area, energy range, and/or grade range to integrate over. We deconvolve the resulting X-ray map with NuSTAR's PSF using the Richardson-Lucy method (Richardson 1972). Richardson-Lucy deconvolution is a Bayesian-based iterative method of removing the distortion NuSTAR's point spread function (PSF) adds to the resulting X-ray emission map (see Section 1.4.1). The number of iterations is chosen to be high enough as to reveal more about the X-ray source shape but not too high as to start resolving noise. This helps with co-alignment between NuSTAR and other observatories since NuSTAR's pointing accuracy is not optimal when observing the Sun.

Before being able to deconvolve the NuSTAR images using Equation 1.5, we first need to identify the correct PSF to use (see Section 2.1.1). NuSTAR's PSF profile changes shape with radial distance from the optical axis (OA) and orientation with azimuthal angle (Madsen et al. 2015). Accounting for these parameters does not reveal any additional or different structures; however, the orientation of the contours appear to match slightly better with the features seen in Fe XVIII. The angle between AR12721 and NuSTAR's OA was estimated to be $1.0'$–$4.4'$ with an azimuthal angle between 0°–40° over both days.

[1] https://github.com/NuSTAR/nustar_pysolar

We investigate the 2.5–4 keV and 4–10 keV NuSTAR energy ranges as the emission from the former tends to closely follow that emission seen by the synthetic Fe XVIII channel and a clear increase occurs >4 keV for each microflare. We find no significant emission >10 keV for any microflare discussed in this section.

The number of iterations used in the deconvolution of the lower and higher energy range was 100 and 50, respectively, for the two brightest microflares in these data (microflare 3 and 10). For the weaker microflares we chose 80 and 40 iterations for the same energy bands. We investigate iterations from 25 to 200 in increments of 25 and observed that the general microflare shape does not change for any event, they only become more pronounced and compact when increasing the number of iterations.

By investigating the orientation and general shape of the NuSTAR contours, we can gain greater confidence on where the dominant X-ray emission is localised. In the past, footpoints of microflares have been observed to be the location of HXR non-thermal emission with the loop itself being mainly comprised of thermal emission (Krucker et al. 2002; Hannah et al. 2011). Therefore, contour plots help indicate approximately where X-ray emission is being produced in the corresponding EUV structure, although we do not find a significant spatial separation between different energy ranges in the microflares studied here. The X-ray and EUV structures are co-aligned across both observatories using cross-correlation; however, a manual shift was also applied to some microflares (approximately 2″–5″). All total shifts were consistent with NuSTAR's pointing uncertainty.

The microflare X-ray data was fitted with the XSPEC spectral fitting tool (Arnaud 1996, see Section 5.3.2), utilising a specific statistic to help with the low count regimes in the spectra (Cash 1979). This is particularly important as the higher energy emission from these weak microflares is commonly found to have few counts.

The thermal and non-thermal nature of the microflaring spectra are investigated using thermal models and broken power-law models, respectively. During the fitting process, we assume coronal abundances (Feldman et al. 1992) and all spectra were created by integrating over a >50″ diameter circular region centred on the brightest emission in the time of interest (red circle in panel e and j in Figure 3.1).

The region for the spectral integration is chosen to be large enough to account for NuSTAR's PSF and include all microflare emission. Any surrounding AR emission included in this area would only add to the hot AR core temperature component (3–4 MK) and not greatly impact on the microflare portion. Note, we do not apply a NuSTAR response gain correction for any microflare, studied in Duncan et al. (2021) and discussed in Section 2.1.5, as the livetime was not sufficiently high for this correction to be needed.

3.3 Orbit 1: Microflare 1 and 2

Within the first orbit, we easily identify two X-ray microflares (labelled 1 and 2 in Figure 3.1 and 3.2) with NuSTAR at approximately 09:15 UTC and 09:25 UTC. These microflares are more difficult to observe in the native SDO/AIA EUV channels; however, the Fe XVIII proxy helps to direct the search for the specific microflare structures. The synthetic Fe XVIII appears to show two spatially resolved loops. The EUV light curves produced in Figure 3.2 (top right panel), which are created from the area enclosed by the corresponding coloured contours in the top left panel, appears to indicate that the most southern loop structure (navy) is the structure responsible for both microflare 1 and 2 as seen in the X-ray time profile (purple and red). The X-rays from microflare 2 also appear to coincide with the peak of the slower evolving EUV emission (cyan) shown in the most northern loop structure.

We overlay the integrated NuSTAR X-ray emission contours on the average Fe XVIII emission over the defined time ranges and show that NuSTAR does not indicate the presence of multiple resolved structures (Figure 3.2, middle row) despite the presence of two distinct loop structures approximately NuSTAR's PSF FWHM apart. It is likely, however, that both loop structures are contributing to the overall emission seen by NuSTAR even if only one is producing the impulsive profiles. Due to AR12721's location in NuSTAR's FOV we are only able to utilise FPMB for spatial and spectral analysis. We co-align NuSTAR FPMB images with the Fe XVIII synthetic emission using a single shift per CHU combination. Between times 09:13:36–09:31:56 UTC there is a noticeable shift in the X-ray location relative to the Fe XVIII emission. This drift could be real change in the source location. However, it could also be due to a slight change in pointing since the shift does not correspond to any Fe XVIII behaviour or due to the deconvolution process since the drift is less pronounced in the raw NuSTAR images. The movement of the source meant that a single shift would not align the X-ray and EUV source for the full time being investigated.

We find that the structure responsible for microflare 1 is clearly located within the navy contours in Figure 3.2 (top left panel) despite the brighter Fe XVIII source being present in the cyan contours. In the top right panel, microflare 1 becomes more obvious in the X-ray higher energy range (4–10 keV, red) compared to the lower energy emission (2.5–4 keV and Fe XVIII, purple and navy, respectively).

Due to the drift and resolution of NuSTAR, and the complex temporal and spatial nature of these events, there is some ambiguity as to which Fe XVIII loop source microflare 2 originates. The emission from microflare 2 could be from either or both loops since both show a peak at the same time in Figure 3.2. This may suggest that there is a physical connection between both loop structures meaning that processes in one could be affecting the material in the other, or that a third structure is interacting with both loops.

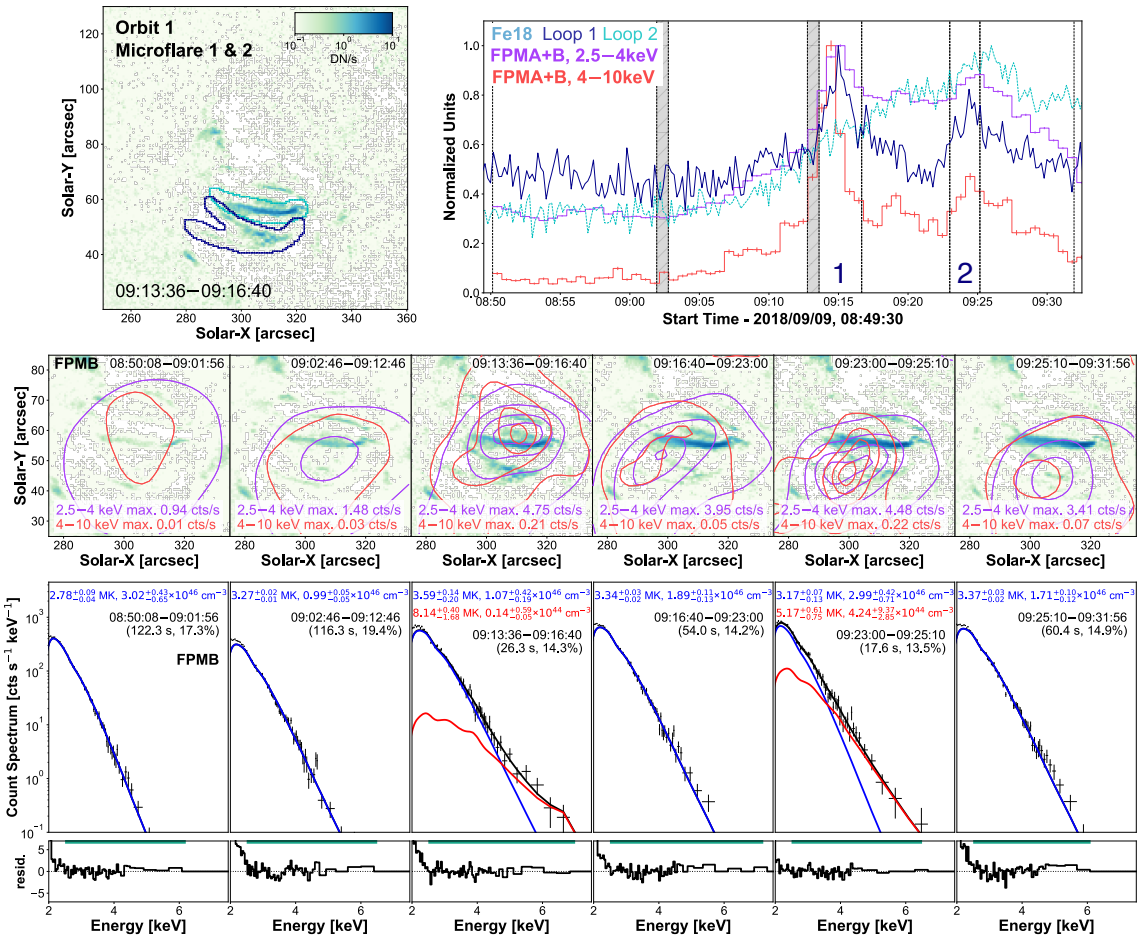

Figure 3.2: The SDO/AIA Fe XVIII synthetic channel emission integrated over the time associated with microflare 1 (top left panel). Two likely sources of microflare 1 and 2's NuSTAR emission are outlined in navy and cyan contours, these coloured contours are used to produce the two Fe XVIII maximum normalised time profiles of the same colour (top right panel). Also in the top right panel, the time profiles of two NuSTAR FPMA+B energy ranges (2.5–4 keV: purple, 4–10 keV: red) with 40 s binning are plotted and were integrated over the full AR. Vertical dashed lines are plotted to indicate a pre-flare time (08:50:08–09:01:56 UTC), a rise time (09:02:46–09:12:46 UTC), microflare 1's time (09:13:36–09:16:40 UTC), and microflare 2's time (09:23:00–09:25:10 UTC). The dark grey regions show times that are not used because of a noticeable shift in source location from CHU state changes. NuSTAR images are created by integrating over the time shown (middle row) and displayed as contours on top of the average Fe XVIII emission over that same time. The contour levels indicate the top 2%, 20%, 50%, and 80% of the maximum emission across all time ranges for each energy range (2.5–4 keV (purple): 4.75 counts s^{-1}, 4–10 keV (red): 0.22 counts s^{-1}). A single shift was applied to the NuSTAR images (middle row) per CHU combination, derived through cross-correlation. The X-ray spectrum for each time range is shown directly below each corresponding image (bottom row). The temperature, emission measure, time range, and livetime of each spectrum is shown while the fitting range is displayed with a horizontal green line. Uncertainties for parameters were obtained from MCMC analysis and are 1-σ equivalent with the higher/lower uncertainty on temperature corresponding to the lower/higher emission measure uncertainty. Note that the same y-range displayed in the top left panel is used for all similar panels on the 2018 September 9–10 microflares for relative spatial context. This figure is taken from Cooper et al. (2021).

We perform spectral fitting at each of the defined times, shown in Figure 3.2 (bottom row), where the times are also visible on each respective panel. Initially, the spectra are fitted with a single APEC[2] thermal model using the XSPEC spectral fitting software (see Section 5.3.2) where we find appropriate fits in all times except those corresponding to the peak microflare emission; the peak times need an additional thermal model to represent the emission from the microflare heated material, as expected. We find the pre-flare plasma is modelled well by a common hot AR temperature of ~2.8 MK with an emission measure of 3.0×10^{46} cm^{-3}. The rise time then appears to show that the pre-flare material evolves to 3.3 MK and 1.0×10^{46} cm^{-3} before heating even more material to 3.6 MK during the peak of microflare 1.

Microflare 1 also produces a counts excess above the enhanced pre-flare emission model. Fitting this excess with another APEC thermal model, we find this microflare reaches temperatures >8 MK with an emission measure of 1.4×10^{43} cm^{-3}. Realistically, a continuum of temperatures will be produced in the microflare (09:13:36–09:16:40 UTC); therefore, it is expected that the model representing the pre-flare emission evolves and that an additional thermal model is required to represent the microflare emission spectrum (as seen in Figure 3.2, bottom row). After microflare 1, the excess vanishes while more material is heated to >3 MK.

Microflare 2 (09:23:00–09:25:10 UTC) is fitted with cooler a temperature (5 MK) and higher emission measure (4×10^{44} cm^{-3}) compared to microflare 1. After microflare 2 the light curves show a gradual decrease and the spectra indicate a relaxation back to a similar state as seen outside of the microflare times. All times investigated of AR12721 from orbit 1 indicate the presence of ~3–4 MK plasma which is in agreement with previous hot AR temperatures found by NuSTAR (Wright et al. 2017; Glesener et al. 2017; Hannah et al. 2019; Cooper et al. 2020). The X-ray emission then shows a return to pre-flare temperatures, but with more material, after microflare 2 before migrating into the chip gap after 09:31:56 UTC.

Despite the difference in evolution of these events both indicate the higher energy range (4–10 keV) peaks before the lower energy emission of Fe XVIII and 2.5–4 keV (Figure 3.2, top right panel). This could elude to the presence of hotter material early on in the event or the presence of accelerated electrons. We also observe that despite the most northern source generally being the brightest (Figure 3.2, top left panel, cyan), it is the structure that hosts the more impulsive emission (navy) that appears to be the origin of the hot microflare material.

Since the microflare loops can be located with the Fe XVIII proxy, we can estimate the emitting volumes to obtain instantaneous thermal energy values for both microflares. The upper limit volumes are obtained by modelling the structure in the cyan contours as one loop and the one in navy contours as

[2] https://heasarc.gsfc.nasa.gov/xanadu/xspec/manual/XSmodelApec.html

two loops, where one loop volume is assumed to be half-torus in shape. We can, therefore, use the SDO/AIA images (Figure 3.2, top left panel) to obtain constraints on the microflaring loop volume with

$$V_{half\text{-}torus} = \pi^2 R^2 r \quad [\mathrm{cm}^3], \tag{3.1}$$

where R is taken as half the distance between footpoints and r is the loop tube radius. Thus, we obtain 2.2×10^{26} cm^3 and 8.4×10^{25} cm^3 for the navy and cyan contour structures, respectively.

When a volume has been obtained for a microflaring loop we can rearranging Equation 1.13 to obtain

$$n_e = \sqrt{\frac{EM}{V}} \quad [\mathrm{cm}^{-3}], \tag{3.2}$$

which can be used to calculate the electron densities for each flaring loop. The corresponding volumes and emission measures are used to find densities of approximately 3×10^8 cm^{-3} and 1×10^9 cm^{-3} for microflare 1 and 2, respectively. These densities are reasonable for flaring loops which are slightly larger than those found under typical coronal conditions (Gallagher et al. 1999; Hannah et al. 2008).

Using Equation 1.15 with the count excess microflare model parameters, and assuming the structure in the navy contours is solely responsible for both microflares, we find instantaneous thermal energies of $1.9^{+1.5}_{-0.3}\times10^{26}$ erg and $6.5^{+3.5}_{-2.4}\times10^{26}$ erg for microflare 1 and 2, respectively. These energy values are upper limits as we assume a volume filling factor of 1, this also makes the density estimates lower limits. If we consider both loops contributing to the emission seen for microflare 2 then the estimated thermal energy increases by ~18% ($7.7^{+4.1}_{-2.8}\times10^{26}$ erg). From the temperature and emission measure from the excess spectral models we can also calculate the GOES equivalent class for each event. This is performed via the SolarSoft IDL routine `goes_flux49.pro`[3] using default coronal abundances with `CHIANTI V7.1` (Dere et al. 1997; Landi et al. 2013). Equivalent GOES classes of A0.001 and A0.01 were obtained for microflare 1 and 2, respectively. Spectral fit model parameters and derived values are shown in Table 3.2.

3.4 Orbit 2: Microflare 3 and 4

During NuSTAR's second orbit we detected two microflares, the first being one of the brightest microflares in these data and the second being the weakest, microflare 3 and 4, respectively.

[3] `https://hesperia.gsfc.nasa.gov/ssw/gen/idl/synoptic/goes/goes_flux49.pro`, see Section 2.2

3.4.1 Microflare 3

Microflare 3 causes a noticeable impulsive microflare profile in Figure 3.1 (panel b and c), peaking at ~10:30 UTC, and appears to have a simple light curve (Figure 3.3, top right panel); however, the images show (top left panel) an evolving complex loop network, resulting in a squashed 'x' shape, being heated. The microflare can easily be split into three 3 min time ranges to investigate the rise, peak, and decay phase (10:28:30–10:37:30 UTC). In addition, a pre-flare time can also be defined since microflare 3 appears to occur in isolation, in space and time.

The X-ray emission in both FPMs were sufficiently clear of the chip-gap and so data from both independent detectors are able to be included when creating contours and spectral fitting (Figure 3.3, middle row, three right panels). FPMA & B are combined by first finding the relative solar-x and -y shift between their observed emission through cross-correlation, deconvolving each with their respective PSFs then adding corresponding energy channels. This does not significantly change the shape of the contours from only one FPM being used. A single shift, determined from the peak time (10:31:30–10:34:30 UTC), is applied to co-align NuSTAR and SDO/AIA across the four times and both energy ranges. There is not a significant shift in the origin of the emission progressing from the rise time range to the decay time range; however, at the peak time, the 4–10 keV contours extend in the direction of the lower leg feature of the squashed 'x' shape. This could indicate some continuing impulsive process even when the overall network loop structure is finishing its rise phase. There does not appear to be a significant difference in location and shape between the two NuSTAR energy ranges.

The pre-flare time is determined to be from the first available time of orbit 2 up to the start of the rising emission (10:26:50–10:28:30 UTC). The X-rays observed appears to be dominated by the general AR emission since the image of this time does not show any concentrated emission in the site of where microflare 3 takes place (Figure 3.3, middle row, left panel). The simultaneous FPMA & B spectral fit supports this (Figure 3.3, bottom row, left panel) as it is consistent with typical hot AR temperatures of ~4 MK (Warren et al. 2012).

We then use microflare 3's pre-flare model as a fixed component, with a temperature 4.1 MK and emission measure of 6.3×10^{46} cm^{-3}, in all subsequent spectra to represent the surround AR emission. Spectral fitting indicates that the rise phase of microflare 3 reaches the hottest temperatures (~7.5 MK) with an emission measure of 4.0×10^{44} cm^{-3} (Figure 3.3, bottom row). It then cools to 6.7 MK at peak emission then to 5.8 MK during the decay. Although the spectral fitted temperature decreases the emission measure drastically increases when progressing to the peak time (16.0×10^{44} cm^{-3}) then again to the decay phase (23.7×10^{44} cm^{-3}). The densities of the microflaring loop, calculated from the volume estimates discussed later in this section, also support this increase in material throughout

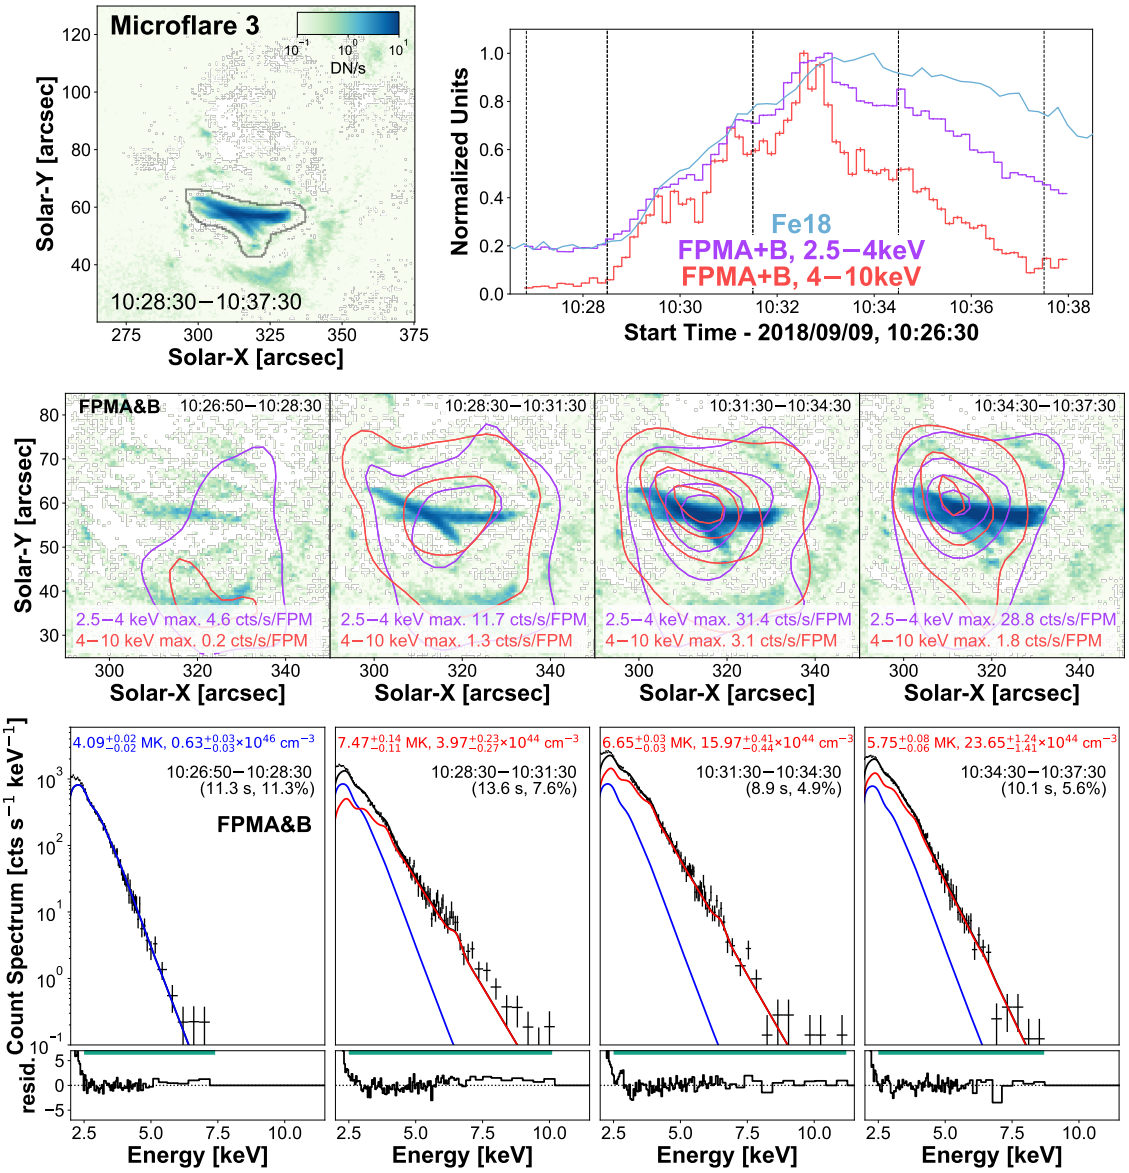

Figure 3.3: Microflare 3's Fe XVIII integrated emission (top left panel) where the black contours show the region used to produce the Fe XVIII time profile (blue, top right panel). NuSTAR FPMA+B 2.5–4 keV (purple) and 4–10 keV (red) time profiles are also shown with 10 s binning which has been integrated over the full AR. The middle row shows the average Fe XVIII emission with NuSTAR contours produced from each of the four times indicated by the vertical dashed lines in the time profile panel. The contour levels are 2%, 20%, 50%, and 80% of the maximum emission found from all time ranges for each energy range (2.5–4 keV (purple): 31.4 counts s^{-1} FPM^{-1}, 4–10 keV (red): 3.1 counts s^{-1} FPM^{-1}). Spectral fits of the four times are displayed in the bottom row, with temperature, emission measure, time range, and livetime shown. The pre-flare thermal parameters (bottom left panel, blue) were fixed components in all three microflare times. This figure is taken from Cooper et al. (2021).

the microflare's evolution, increasing from $\sim 4\times10^8$ cm^{-3} during the rise time (Equation 3.2) to $\sim 7\times10^8$ cm^{-3} in the decay. Therefore, this is consistent with more and more chromospheric material *evaporating* into the coronal loops as the microflare evolves, similar to a standard flare scenario (Fletcher et al. 2011).

We use the NuSTAR temperature and emission measure fitted values to compute the expected Fe XVIII emission and then compare this to the Fe XVIII observation. Using Equation 1.18, we find good agreement between the pre-flare subtracted Fe XVIII proxy flux and the flux calculated from folding the excess X-ray model parameters through the Fe XVIII temperature response. The excess models fitted to the X-ray spectra reveal that NuSTAR sees ~42%, ~65%, and ~75% during the rise, peak, and decay phase, respectively. The minor disagreement could be due to the multi-thermal nature of the microflare in addition to NuSTAR and the Fe XVIII channel having slightly different sensitivities to different temperatures as discussed in Section 2.1.6 and 2.5.1.

Inspecting the rise phase spectral fit in Figure 3.3 (bottom row), we find a significant counts excess still present >7 keV after the two thermal models have been fitted. This suggests an additional model is required to sufficiently represent the X-ray spectrum. Initially fitting another thermal model, we find that an unphysically high temperature of ~95 MK is required to account for the excess (Figure 3.4). Therefore, we investigate the possibility of the counts excess being produced by non-thermal emission and introduce a broken power-law photon model[4]. Ideally we would utilise the thick target model (Brown 1971, and see Section 1.4.5); however, at the time of analysis the thick target model was only available in OSPEX (Section 5.3.1) which was not used for this fitting because of the low count regime. Therefore, we use a double power-law photon model to replicate the photon distribution produced from a single power-law of non-thermal electrons.

We show the results of this fitting in Figure 3.5 where it is found that, in addition to the fixed per-flare thermal model, the power-law model represents the spectrum very well. We find that the microflare thermal model's temperature is reduced to 6.8 MK and emission measure 5.5×10^{44} cm^{-3}, still making it the hottest phase of the flare and all previous arguments remain unchanged. The broken power-law model fit produces a break energy of 6.2 keV, a photon index of 8.3, and a normalisation constant of 0.8 ph s^{-1} cm^{-2} keV^{-1} (at 1 keV). The photon index below the break was fixed at a value of 2.

The unreliable thermal fit shown in Figure 3.4 is further emphasised when inspecting the sampled posterior distribution from the MCMC analysis (Figure 3.6). We find that not only is the higher temperature component unphysical, but also terribly constrained in the sampled posterior that has significant sample density across the whole temperature prior (third diagonal panel). This is reflected in the incredibly large error range presented in Figure 3.4. However, when considering the sampled

[4] https://heasarc.gsfc.nasa.gov/xanadu/xspec/manual/node142.html

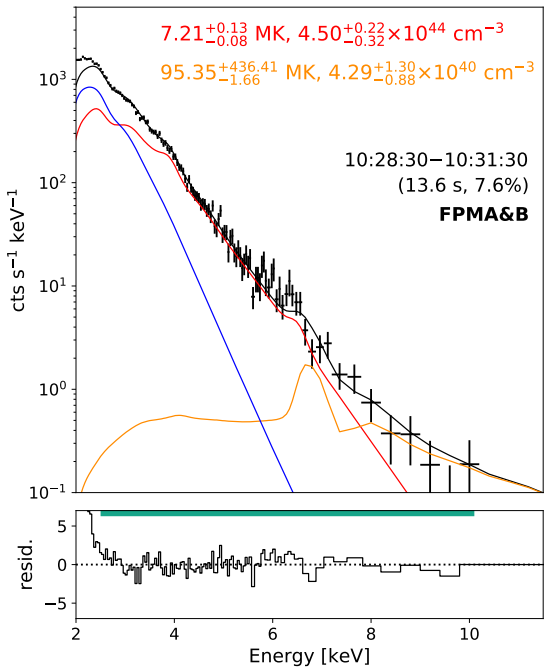

Figure 3.4: The impulsive phase spectrum of microflare 3, shown in Figure 3.3 (10:28:30–10:31:30 UTC), is fitted with the pre-flare model (blue) being fixed and a thermal model (red), as before. An additional thermal model (orange) is introduced in an attempt to represent the count excess; however, this produces an unphysical temperature. The spectra from FPMA and B were fitted simultaneously over the energy range indicated by the horizontal green line with the effective exposure and livetime indicated in brackets.

posterior from the case presented in Figure 3.5 we find that tighter, well behaved distributions that produce Gaussian-like marginalised sample densities are found along the diagonal in Figure 3.7. See Section 5.2 for a detailed explanation of MCMC analysis.

This is a scenario similar to that discussed in Glesener et al. (2020) where the spectrum of an A5.7 microflare was fitted with two thermal models where one of them produced an unreasonably high temperature and was only well-fitted once a non-thermal component was included.

Therefore, we find a power of $7.0^{+3.7}_{-2.3}\times10^{24}$ erg s^{-1} using Equation 1.33, 1.34, and 1.35 with the fitted broken power-law parameters. Over the 3 minute rise phase, this results in $1.3^{+0.7}_{-0.4}\times10^{27}$ erg. In addition to the non-thermal model fitting the higher energy spectra better, we also find better agreement during the rise time between the Fe XVIII proxy and the 6.8 MK temperature model which now predicts ~57% of the EUV emission instead of ~42%.

Due to the complex loop structure of microflare 3, the volume (V) is calculated using

$$V_{area} = A^{3/2} \quad [\text{cm}^3], \tag{3.3}$$

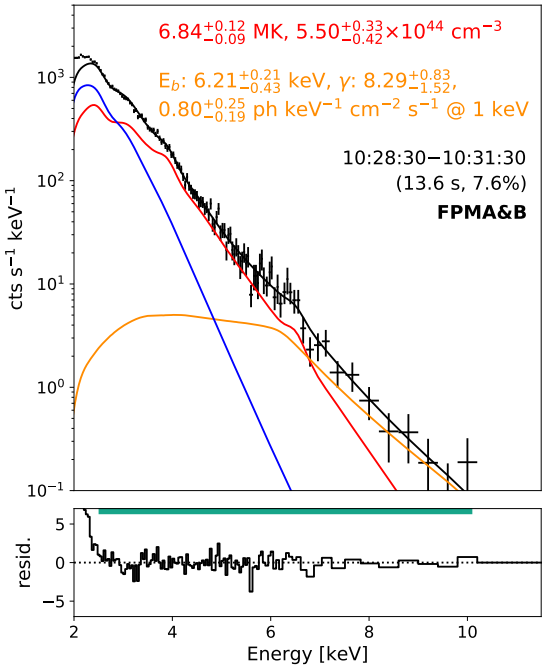

Figure 3.5: The impulsive phase spectrum of microflare 3, shown in Figure 3.3 and 3.4 (10:28:30–10:31:30 UTC), is fitted again with the pre-flare model (blue) being fixed and a thermal model (red), as before. However, an additional photon broken power-law model (orange) is introduced to represent non-thermal electron emission thought to be causing the count excess. The spectra from FPMA and B were fitted simultaneously over the energy range indicated by the horizontal green line with the effective exposure and livetime indicated in brackets. This figure is taken from Cooper et al. (2021).

with the area (A) estimated from the averaged Fe XVIII emission (Figure 3.3, top left panel). We find an upper limit volume of 3.5×10^{27} cm^3 (density of $\sim4\times10^8$ cm^{-3}, Equation 3.2) during the rise which then evolves to a volume of 4.9×10^{27} cm^3 (density of $\sim7\times10^8$ cm^{-3}) in the decay time. Therefore, we find thermal energy releases of $3.9^{+0.1}_{-0.1}\times10^{27}$ erg from the 6.8 MK rise component, $6.5^{+0.1}_{-0.1}\times10^{27}$ erg during the peak time, and $8.1^{+0.1}_{-0.1}\times10^{27}$ erg for the decay phase. Although the non-thermal energy release is not consistent with what appears to be required from the thermal component, this could be due to the non-thermal component not being solely responsible for all heating or that we do not observe all of the non-thermal emission. In addition, the volume used to calculate the thermal energy is no doubt an upper limit because of the nature of its calculation.

We have been able to investigate microflare 3 in several time ranges; however, this is not the case with the majority of microflares in these data. In order to help compare microflare 3 to the others, we perform spectral fitting over the full duration (10:28:30–10:37:30 UTC), shown in Figure 3.8. The photon power-law model from the rise time and pre-flare thermal model is fixed during the fit. The power-law normalisation factor was scaled to 0.3 ph s^{-1} cm^{-2} keV^{-1} as it is assumed that

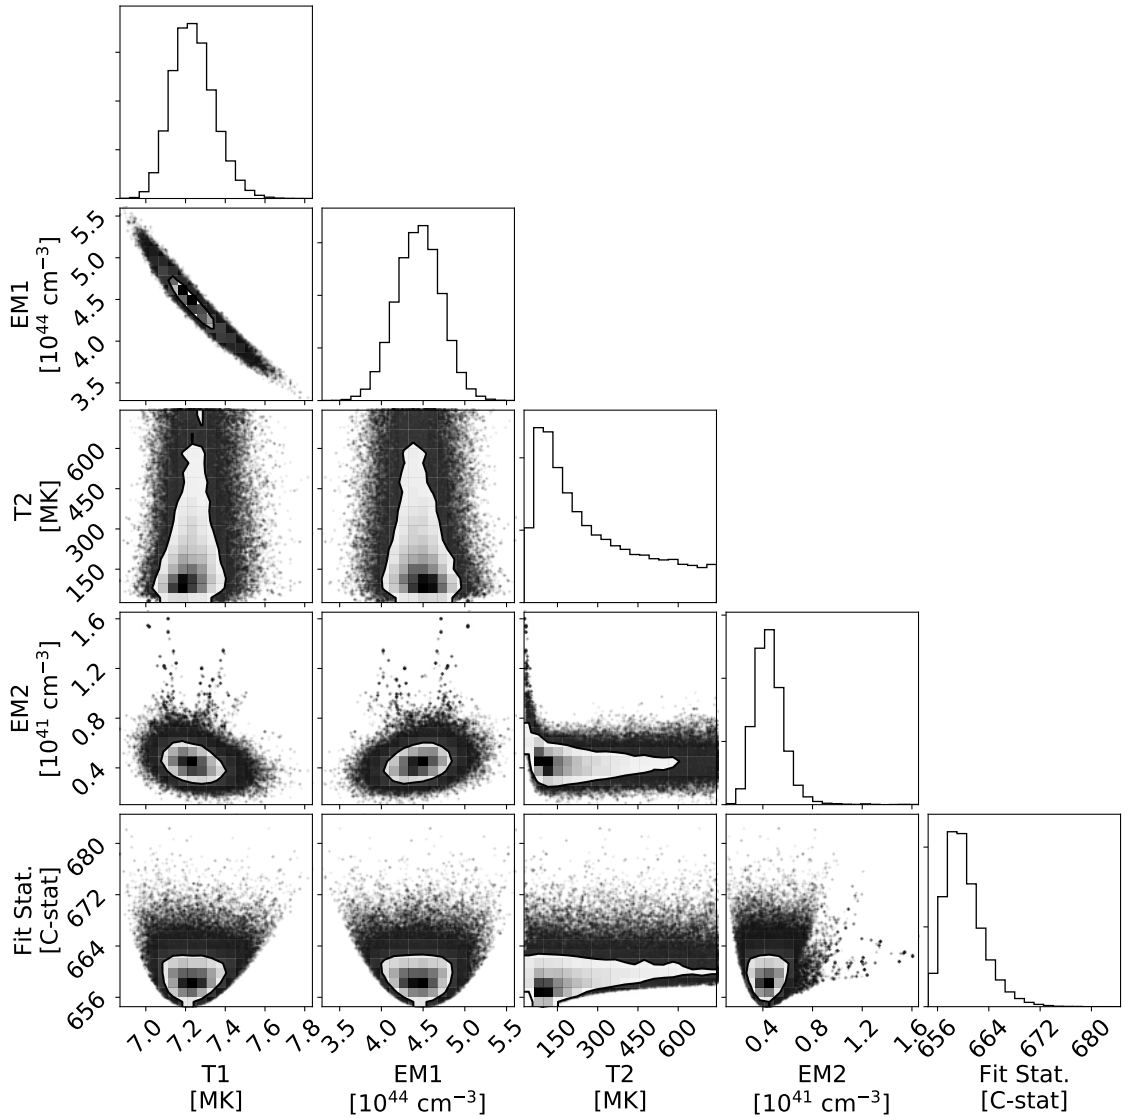

Figure 3.6: Corner plot of the MCMC analysis used to produce Figure 3.4 performed in XSPEC. The contours in the off-diagonal panels show 68% (1-σ equivalent) of the MCMC walker density and how each parameter relates to every other one, while the diagonal panels show the marginalised density distributions of each parameter.

the non-thermal particle acceleration is only significant in the rise time. Fitting the full time of microflare 3 with an additional thermal model we find a temperature and emission measure of 6.5 MK and 11.6×10^{44} cm^{-3}, respectively. We obtain the average thermal energy to be $6.49^{+0.04}_{-0.03}\times10^{27}$ erg with an equivalent GOES class of A0.1 at the defined peak phase. Therefore, this would suggest that microflare 3 is one of the weakest X-ray non-thermal microflares currently in literature. Microflare 3's spectral fit parameters for all investigated times are displayed in Table 3.3.

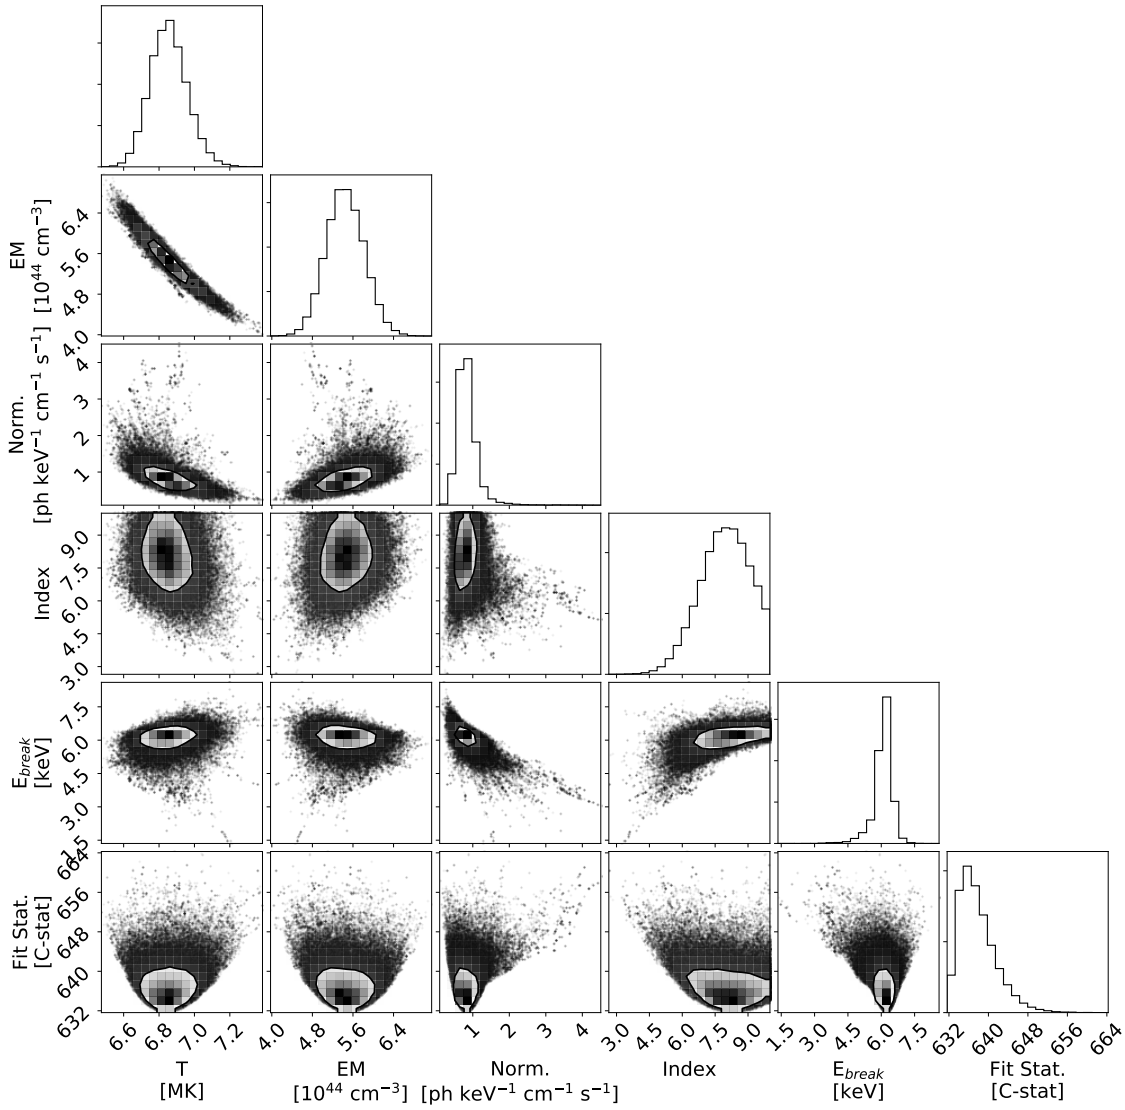

Figure 3.7: Similar format to Figure 3.6 but for the MCMC analysis used to produce Figure 3.5.

We note that the spectral fit from the whole time (10:28:30–10:31:30 UTC) still recovers a similar result to the fit during the rise time but with a slightly lower temperature and higher emission measure. This would be expected due to the evolution of the flare within the 9 min window. However, in the full time spectral fit, which is analogous to how many other microflares must be analysed due to their short duration, the non-thermal model is almost completely hidden and dominated by the thermal emission. Therefore, this illustrates the importance of careful time selection by the investigator, or the ability to select a time range from the detection, even at these very weak microflare scales to avoid

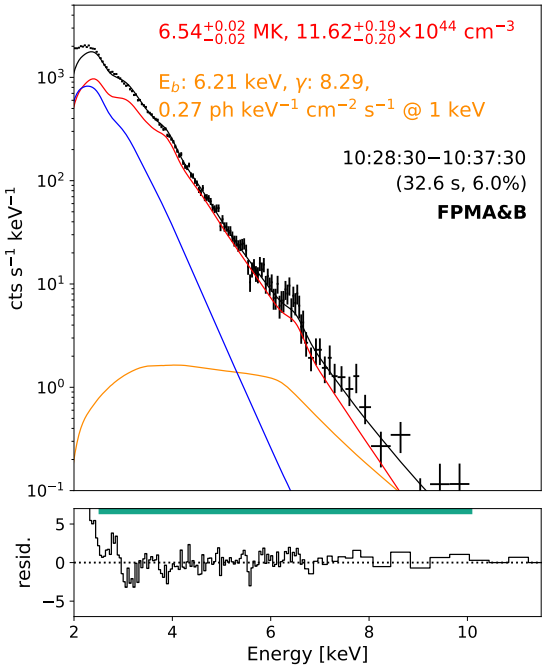

Figure 3.8: The spectrum from microflare 3's full time (10:28:30–10:31:30 UTC) with the photon power-law from Figure 3.5 fixed and scaled. The spectra from FPMA and B were fitted simultaneously over the energy range indicated by the horizontal green line with the effective exposure and livetime indicated in brackets.

any potential non-thermal emission being completely hidden or being indistinguishable from thermal emission.

We return to microflare 3's rise time spectrum in Section 6.1.5 where we utilise a newly developed spectral fitting software in Python. This software allows us to fit FPMA & B simultaneously using Poisson statistics, much like XSPEC used throughout this section; however, we can fit the count excess over the thermal model with the thick target model allowing us to directly model a power-law of accelerated electrons instead of trying to link a photon spectrum back to the electron distribution properties.

3.4.2 Microflare 4

From one of the brightest microflares on 2018 September 9–10 to the weakest; occurring during the same orbit as microflare 3 we observe microflare 4, SOL2018-09-09T11:04, which is discussed in detail in Chapter 4 and is the subject of Cooper et al. (2020). This is the weakest X-ray microflare from an AR in literature (Figure 3.9). Microflare 4 is only apparent when investigating the higher NuSTAR energy ranges (>4 keV) and is then able to be found visually in the lower energy range, 2.5–4 keV,

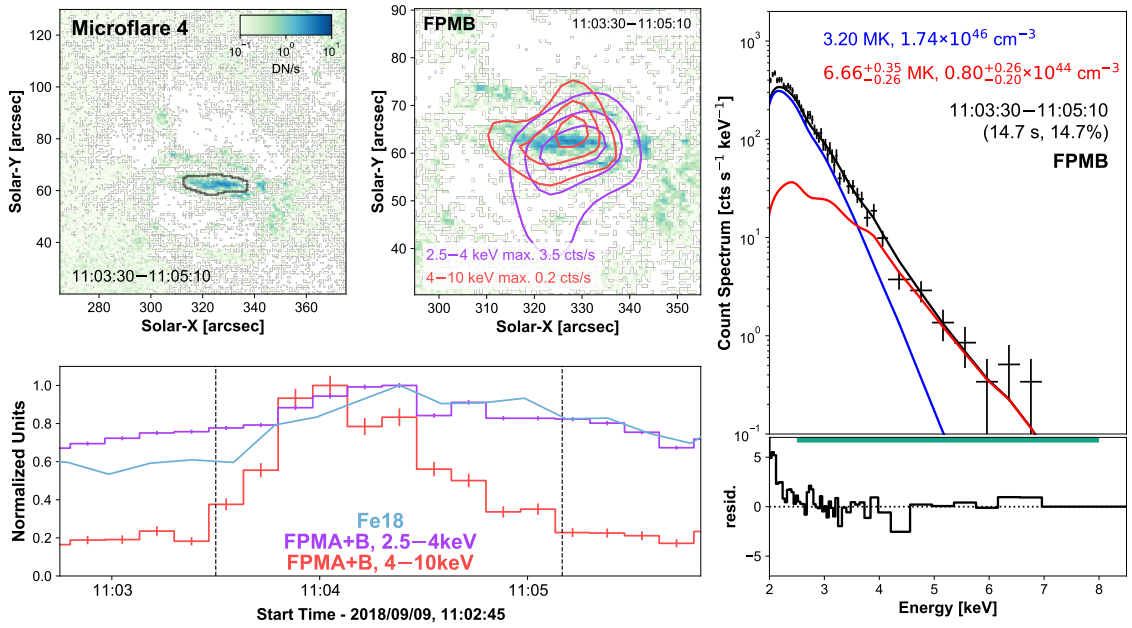

Figure 3.9: Integrated SDO/AIA Fe XVIII emission for microflare 4 over the time indicated (top left panel), the black contour shows the region used to produced the Fe XVIII time profile (bottom left panel). NuSTAR X-ray FPMB 2.5–4 keV (purple) and 4–10 keV (red) contours are shown on top of the average Fe XVIII emission (top middle panel) with 20%, 50%, and 80% levels of the maximum X-ray emission for each X-ray energy shown. Microflare 4's time is also indicated by the two vertical dashed lines in the bottom left panel where the X-ray light curves are integrated over the full AR with 10 s binning (for a more detailed analysis of microflare 4 see Cooper et al. 2020, and Chapter 4). This figure is taken from Cooper et al. (2021).

and the EUV via the Fe XVIII proxy. This event also benefits from being isolated in NuSTAR's FOV helping maintain a high livetime throughout its evolution as this microflare could have easily been completely masked if another, brighter source was present. We are also able to determine a suitable pre-flare time. Performing a spectral fit (Figure 3.9, right panel), we find a temperature and emission measure of 6.7 MK and 8.0×10^{43} cm^{-3}, respectively. Microflare 4's spectral fit values are shown in Table 3.2.

We obtain a volume of 1.9×10^{25} cm^3 by modelling microflare 4's structure as a half-torus using Equation 3.1 from the average Fe XVIII emission (Figure 3.9, top left panel) estimating a density of $\sim2\times10^9$ cm^{-3} from Equation 3.2. This provides an instantaneous thermal energy of $1.1^{+0.1}_{-0.1}\times10^{26}$ erg, when using Equation 1.15 with the microflare excess thermal model parameters, and an equivalent GOES class of ~A0.005 (scaled: A0.007, see Section 2.2).

3.5 Orbit 3: Microflare 5, 6, and 7

During orbit 3, the last dwell observation on September 9, we clearly observe three microflares (microflare 5, 6, and 7) located in the proximity of AR12721. They can all be observed above the decaying emission from a relatively large microflare which occurred at ~11:45 UTC (Figure 3.1, top panel), between orbit 2 and 3. The loops responsible for this missed microflare appears to have a footpoint anchored in the top right of the AR. The post-microflare decaying loops elevating the 'background' emission can be seen most clearly in microflare 5's integrated Fe XVIII images and NuSTAR contours (Figure 3.10, top left and middle panel), and can also be seen to a lesser extent in the same panels for microflare 6 and 7 (Figure 3.11 and 3.12). Although the decaying loop is very bright compared to the microflares, there is a large spatial separation; large enough to isolate and remove it from the light curves (bottom left panels) and spectral fitting (right panels) in Figure 3.10, 3.11, and 3.12. This means that the microflare is more prominent in these time profiles than that shown in Figure 3.1 and helps constrain the microflare thermal model fits. We note that the exclusion of the decaying loop had little effect on the fitted temperatures and emission measures; however, the microflare emission was more visible and easier fitted when the decaying loop, which contributed greatly to the lower temperature model, was removed. Both FPMA & B were able to be utilised for microflare 5; however, only FPMB could be used for microflare 6 and 7 for quantitative analysis.

There appears to be two loop structures that could be responsible for microflare 5 once the large decaying loop is removed, these are shown in the black and grey contours in Figure 3.10 (top left panel). By investigating the Fe XVIII light curves from both loops we find that the brighter bottom EUV source monotonically increases, whereas the emission from the fainter top loop (black contours in Figure 3.10) indicates similar behaviour as seen in X-rays, with a maximum 2 s after the lower energy X-ray peak. This is similar reasoning to that discussed in Section 3.3 for the location of microflare 2.

The Fe XVIII time profiles in Figure 3.11 (bottom left panel), produced from their respective contour regions in the integrated Fe XVIII images (top left panel), show that microflare 6 is produced from a similar location as microflare 5 (indicated by black contours in Figure 3.11). However, microflare 7 comes from the bright Fe XVIII structure above microflare 5 and 6, shown by the grey contours in Figure 3.11, which are then black in Figure 3.12.

To corroborate the location of microflare 6, we look at the orientation of the higher NuSTAR energy range; both the loop in SDO/AIA Fe XVIII and the 4–10 keV X-ray range are angled similarly (Figure 3.11, top middle panel), although the position might suggest the right footpoint is heated more than its left one. A CHU state change occurring immediately after microflare 6 makes it especially difficult to analyse, possibly missing some significant X-ray emission during the decay when the source moves into the chip gap in both FPMs for several minutes (12:45:50–12:53:21). Microflare 6 appears

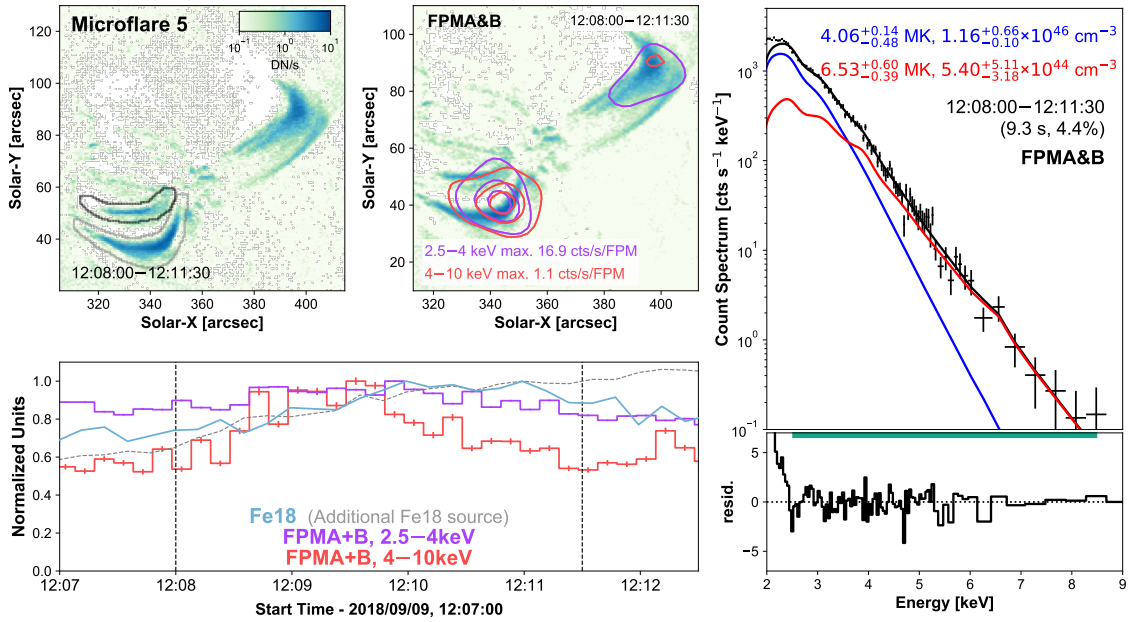

Figure 3.10: The same format as was used for microflare 4 in Figure 3.9 but for microflare 5. Both FPMA & B are able to be used in the spectral fitting and contour construction while the time profiles (bottom left panel) do not include the large decaying loop emission to the right of the black and grey contour regions in the top left panel. The black and grey contours are used to produce the blue and grey time profiles. This figure is taken from Cooper et al. (2021).

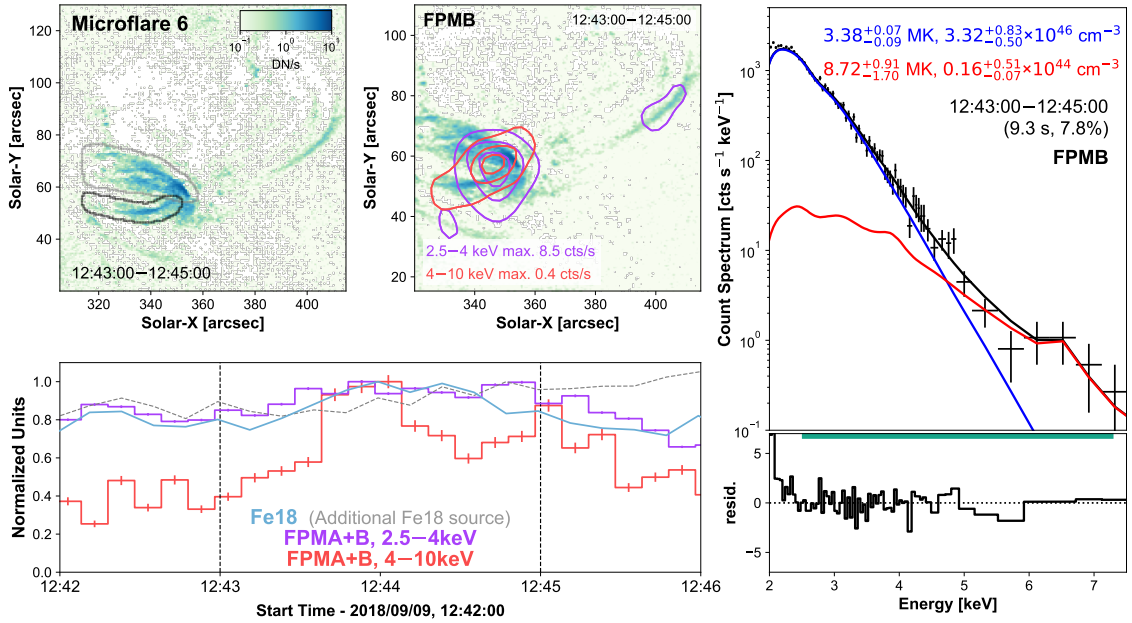

Figure 3.11: The same format as was used for microflare 4 in Figure 3.9 but for microflare 6. The black and grey contours are used to produce the blue and grey time profiles. This figure is taken from Cooper et al. (2021).

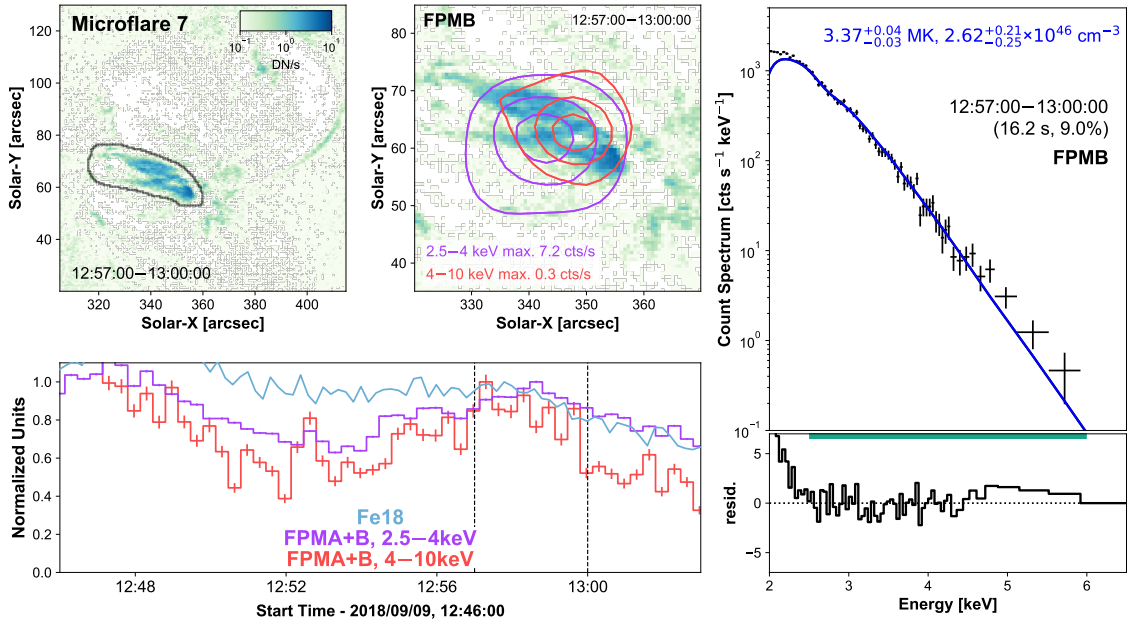

Figure 3.12: The same format as was used for microflare 4 in Figure 3.9 but for microflare 7. The NuSTAR light curves in the bottom left panel (purple and red) are binned in 20 s segments and the emission from a microflare missed due to an abrupt CHU state change is visible at ~12:47 UTC. This figure is taken from Cooper et al. (2021).

to be the last event to occur in the general structure responsible for microflare 5 when considering the Fe XVIII emission.

Inspecting the light curves from microflare 7's arc structure between microflare 6 and 7's time (from the grey contours in Figure 3.11 and black in Figure 3.12), we find qualitative evidence of another microflare at ~12:47 UTC. However, this was removed from future analysis as the aforementioned CHU change meant the NuSTAR emission could not be used.

A weak, but noticeable rise in emission is observed to be caused by microflare 7. The microflare signature is particularly noticeable in X-rays with slightly similar time profiles to Fe XVIII (Figure 3.12, bottom left panel). The time used for analysis, 12:47–13:00, is determined from the first usable time after the disruptive CHU state change and a significant drop in the 4–10 keV X-ray range. The impulsive emission appears to originate from the base of the arc structure where the higher NuSTAR energy range also indicates localised emission at the western footpoint of the general loop structure seen in Fe XVIII and 2.5–4 keV.

All three microflares investigated during orbit 3 (microflare 5, 6, and 7) show co-temporal signatures in the SDO/AIA 131 Å and 171 Å channels to the time profiles seen in Figure 3.10, 3.11, and 3.12 (bottom left panels) when isolating the emission using the black contour regions (Figure 3.13). The

structures in these SDO/AIA channels that are located in the contours appear to show activity from smaller sub-structures to the larger one identified with the Fe XVIII proxy. This enhanced emission across multiple SDO/AIA channels, therefore, helps corroborate that the correct loop structure has been identified as the source of the impulsive X-ray emission.

Even though microflare 6 and 7 are both very weak events the X-ray spectra show a clear hot excess beyond a single thermal model fit (Figure 3.10 and 3.11, right panels). We find very hot temperatures in microflare 6 and 7 of 6.5 MK and 8.7 MK, respectively. Microflare 7 shows a slight excess >5 keV above a single thermal model that is consistent with hot AR temperatures (Figure 3.12, right panel). Unfortunately, the fitting parameters when an additional model is introduced are not well constrained due to the weak nature of the excess; this makes any interpretation of a two-model fit difficult.

We investigate this weak excess by comparing microflare 7 to another microflare, microflare 2, which displays similar relative increases in emission suggested by their time profiles. Microflare 7 (Figure 3.12) appears to double in 2.5–4 keV emission, similar to that seen from microflare 2 (Figure 3.2); however, microflare 7 does not have the same drastic rise in the 4–10 keV X-ray range. This could be due to microflare 7 having half the livetime fraction, due to the 11:45 UTC decaying microflare, when compared to microflare 2 which would reduce the recording of higher energy counts detected by NuSTAR or the elevated 'background' emission masking the excess making it difficult to observe and fit.

Again, we model the loops as half-tori (Equation 3.1) finding that microflare 5 and 6 have a volume of 7.9×10^{25} cm^3 (densities of approximately 3×10^9 cm^{-3} and 5×10^8 cm^{-3} from Equation 3.2) while two half-tori were used to estimate microflare 7's volume of 5.2×10^{26} cm^3 (density of $\sim7\times10^9$ cm^{-3}). We then use Equation 1.15 with the volumes and relevant spectral fit parameters finding instantaneous thermal energies of $5.6^{+1.8}_{-1.7}\times10^{26}$ erg, $1.3^{+0.8}_{-0.2}\times10^{26}$ erg, and $5.2^{+0.2}_{-0.2}\times10^{27}$ erg—with equivalent GOES classes of A0.03, A0.002, and A0.2 (scaled: A0.04, A0.003, and A0.3, see Section 2.2)—for microflare 5, 6 and 7, respectively. We note that the energy calculated for microflare 7 is an estimate of the overall energy release of the elevated AR emission model and not of the microflare excess like for microflare 5 and 6. Orbit 3's microflare spectral fit, and derived, values are shown in Table 3.2.

3.6 Orbit 4: Microflare 8 and 9

NuSTAR's fourth orbit, the first orbit on September 10, contains the signatures of two separate heated loop events. Both events are observed to be from the post-flare loop arcade of a microflare that occurred an hour before the NuSTAR observation commenced on September 10 and was only

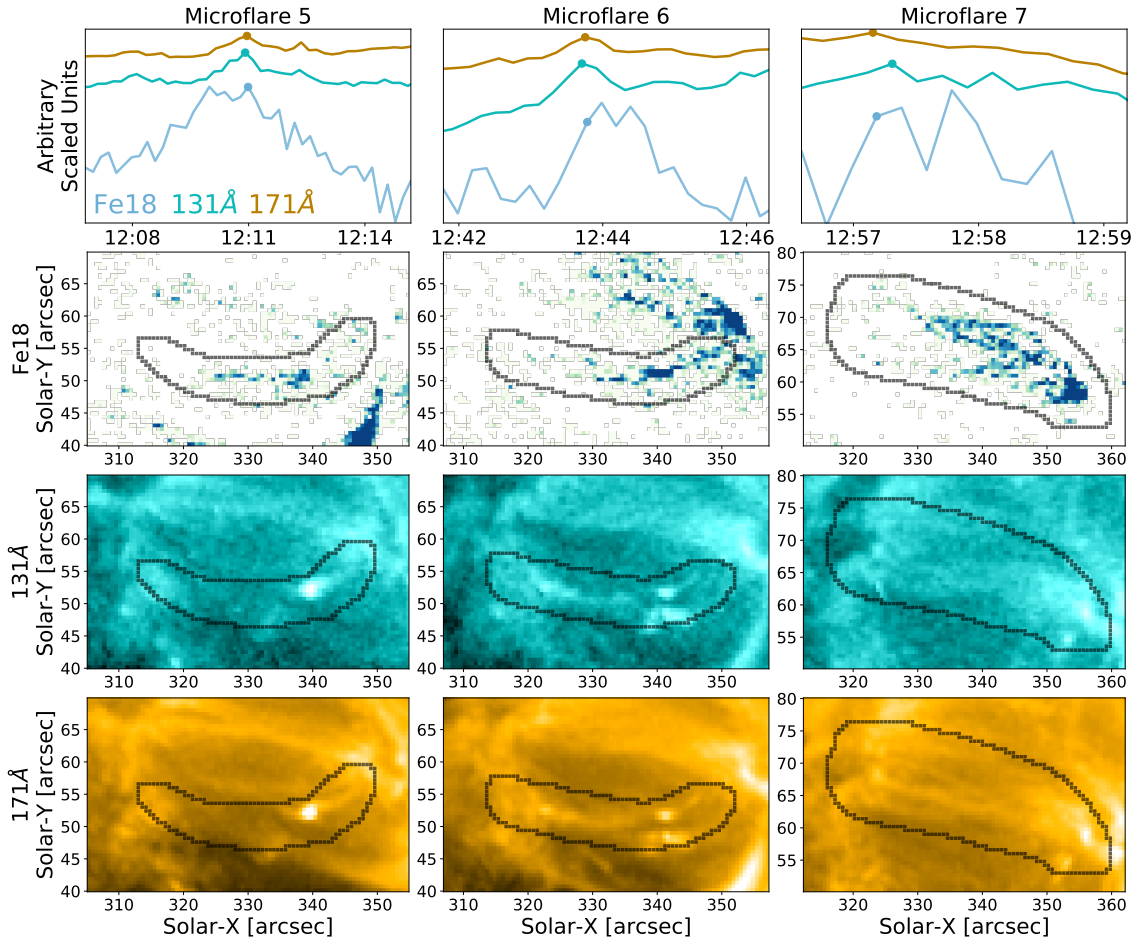

Figure 3.13: Microflare 5, 6, and 7's Fe XVIII (blue), 131 Å (blue), and 171 Å (yellow) time profiles from the loop structures identified in the Fe XVIII proxy emission (top row). The Fe XVIII, 131 Å and 171 Å images, depicted in the same colours as their respective time profiles and taken around their peak emission shown by a circle on each light curve, show a zoomed in region around each microflare's structure with the same black contours shown in Figure 3.10, 3.11, and 3.12 for spatial context and comparison to the pervious Fe XVIII maps. At these peak times, which are well within the Fe XVIII and X-ray microflaring time ranges, each channel shows a corroboratory bright and compact source within the black contour region.

observed by SDO/AIA. We find that microflare 8 shows the decay of this missed microflare and microflare 9 occurs in microflare 8's post-flare arcade loops (Figure 3.14).

NuSTAR X-ray emission for microflare 8 and 9 shows the top of the heated loop arcade (Figure 3.14, middle panels) while the Fe XVIII emission (top left and middle panels) shows the same source as well as the loop footpoints. The footpoints of these loops, only evident in EUV, are seen as two negatively sloped lines south-east of the heated loop tops. The loop-top source is indicated with black contours in the integrated Fe XVIII of Figure 3.14 (top left panel). There does not appear to be any obvious spatially impulsive evolution in Fe XVIII as the loops progress; the X-ray emission contours show a slight elongation of the source perpendicular to the loops which is possibly related to microflare 9. Only the rise and peak of microflare 9 can be analysed as an SAA passage meant the decay was missed.

Despite the non-quiet behaviour of microflare 8 and 9's time profiles, we find that the FPMA & B fitted spectra are well represented with a single thermal model even though microflare 9 especially has a light curve where an excess might have been expected (Figure 3.14, bottom row). We find temperatures that are in the upper range of hot AR core temperatures with microflare 8's decaying loops having a temperature of 4.1 MK and emission measure 8.9×10^{45} cm^{-3} which then evolves to 4.3 MK and 7.6×10^{45} cm^{-3} during microflare 9.

We estimate the whole volume of the overall structure, again using a half-torus shape with Equation 3.1, finding a volume of 1.7×10^{27} cm^3 and densities of approximately 1×10^9 cm^{-3} (Equation 3.2) for both microflares. However, due to the clear space between the loop footpoints and the loop-top source we employ a filling factor of 0.5. Calculating the equivalent GOES classes for these events, we find a class of A0.1 for both microflare 8 and 9 which is comparable to that of the at largest microflares in these data. We also find instantaneous thermal energy releases of $1.07^{+0.03}_{-0.03}\times10^{28}$ erg and $1.03^{+0.02}_{-0.03}\times10^{28}$ erg for microflare 8 and 9, respectively. The spectral fit values for microflare 8 and 9 can be found in Table 3.2.

3.7 Orbit 5: Microflare 10

The brightest X-ray microflare recorded by NuSTAR from AR12721 is microflare 10, despite going into nighttime before the peak of the event as comparison with the Fe XVIII time profile suggests (Figure 3.1, panel g and h, and Figure 3.15, top right panel). Like microflare 3 (Figure 3.3), microflare 10 has a spatially complex structure and is an amalgamation of various microflaring loops (Figure 3.15, top left panel).

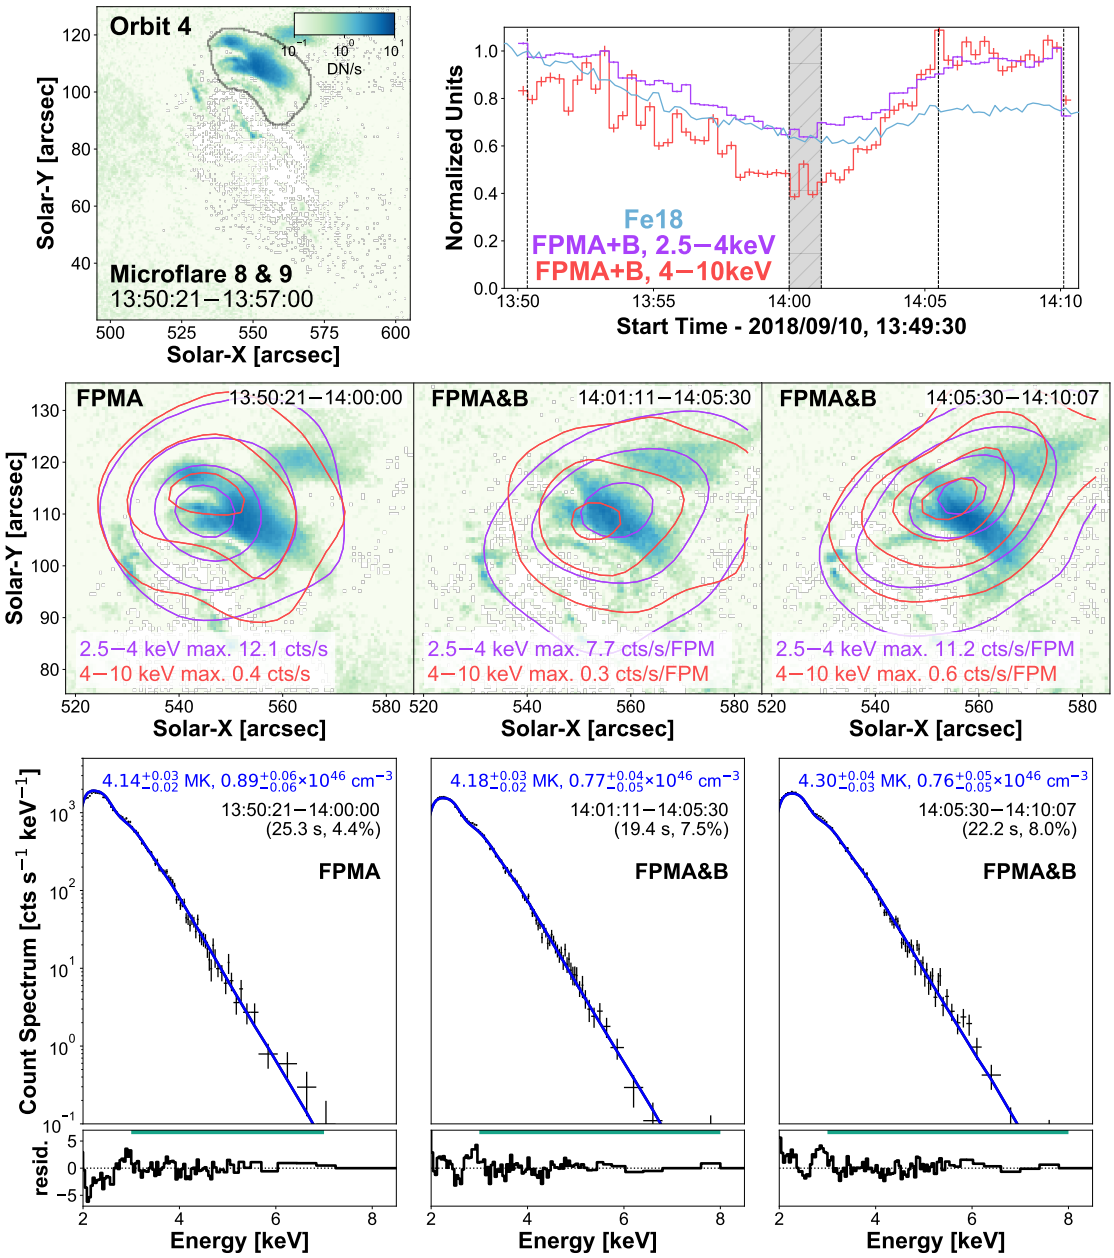

Figure 3.14: Integrated over microflare 8's time, an SDO/AIA Fe XVIII image (top left panel) shows the loops being investigate and shows the region used to produce the EUV time profiles in black contours (top right panel). Normalised Fe XVIII (blue) and NuSTAR (2.5–4 keV: purple, 4–10 keV: red, 20 s binning) light curves are plotted where the X-ray time profiles have been integrated over the full AR. A significant CHU change occurred between microflare 8 and 9, indicated by a dark grey region, which was excluded from analysis. NuSTAR X-ray contours (middle row) show the 5%, 20%, 50%, and 80% of the maximum emission across all time ranges for each energy range (2.5–4 keV (purple): 12.1 counts s^{-1} FPM^{-1}, 4–10 keV (red): 0.6 counts s^{-1} FPM^{-1}) and are plotted on the average Fe XVIII emission of the time ranges shown and by vertical dashed lines in the light curve panel. Results from the X-ray spectral fits for the same times are shown in the bottom row where the energy ranges used for fitting is indicated with a green horizontal line. This figure is taken from Cooper et al. (2021).

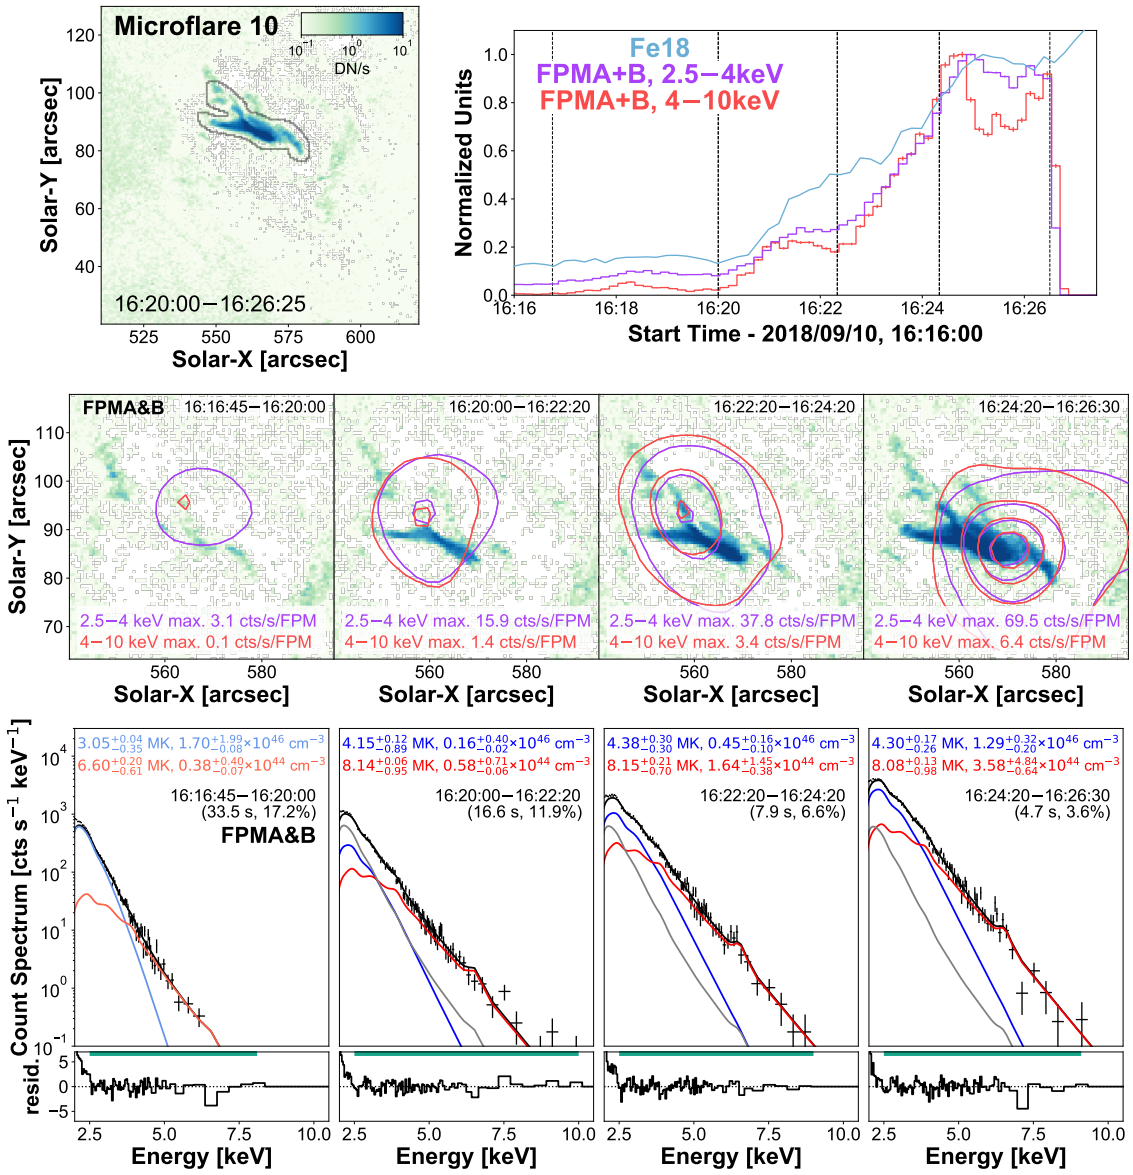

Figure 3.15: The same format as was used for microflare 3 in Figure 3.3 but for microflare 10. Both pre-flare thermal model fits (faded blue and red, bottom left panel) are kept as fixed components for all other microflare 10 spectral fits (grey). Contour levels are 2%, 20%, 50%, and 80% of the maximum emission across all time ranges for each energy range (2.5–4 keV (purple): 69.5 counts s^{-1} FPM^{-1}, 4–10 keV (red): 6.4 counts s^{-1} FPM^{-1}). This figure is taken from Cooper et al. (2021).

3.7 *Orbit* 5: *Microflare* 10

We investigate microflare 10 across four stages, a pre-flare (16:45–20:00 UTC), initial rise (20:00–22:20 UTC), continued rise (22:20–24:20 UTC), and plateau (24:20–26:30 UTC) which are indicated by vertical dashed lines in Figure 3.15's time profile (top right panel). NuSTAR contours produced from each phase are overlain on the Fe XVIII images averaged over the NuSTAR time range (Figure 3.15, middle row) where we find that the pre-flare time does not appear to show a localised source in both X-rays and EUV, similar to the scenario presented for microflare 3. However, NuSTAR's FPMA & B contours show a coincidence with the EUV source location where there is a small shift from the north-east of the structure to the south-west, appearing to agree with the evolution seen in Fe XVIII.

A stark difference between microflare 3 and 10 is the multi-thermal nature of their respective pre-flare phases. Microflare 3's pre-flare is fitted well with one thermal model; however, in order to adequately represent microflare 10's pre-flare we require two thermal models (Figure 3.15, bottom row, left panel). As with all other microflares in this analysis, the pre-flare components were kept fixed for all other successive phases (Figure 3.15, bottom row, right three panels, grey). However, we find that the hot, 6.6 MK component has a negligible effect on any resulting fit other than in the pre-flare stage, but is kept for consistency. After the pre-flare time we find that the microflaring plasma is kept heated at ~8.1 MK while an increasing amount of material at this temperature is continually increasing with an emission measure rising from 0.6×10^{44} cm^{-3} to 3.6×10^{44} cm^{-3}. The hot AR temperature component at ~3–4 MK also undergoes a similar enhancement.

Since we can fit a well determined pre-flare component to the microflare spectra it then becomes possible to compare the pre-flare subtracted Fe XVIII emission to the spectral fitting results. We use the microflare excess model parameters from microflare 10's spectral fits to investigate the consistency between the emission observed by NuSTAR and the SDO/AIA Fe XVIII proxy channel. We do not find good agreements at any time with NuSTAR seeing ~112% at the initial rise time, ~128% during the continued rise, and ~135% at the plateau. This disagreement could indicate that there is non-thermal emission present making the NuSTAR models inconsistent with the Fe XVIII proxy, could be due to each instruments different sensitivities to observed temperature distribution, or could be because the 4 MK plasma model component is at the edge of the Fe XVIII proxy channel's temperature response but contributes the majority of the total synthetic flux value making any comparison between NuSTAR and the composite SDO/AIA channel more problematic.

Little suggests that a non-thermal component is needed to fit microflare 10's spectra, other than the synthetic flux inconsistency, since the presented total model fits appear to describe the count spectra well (Figure 3.15, bottom row). However, we investigate its possible inclusion by replacing the higher temperature thermal model with a broken power-law which then reduces the peak temperature of microflare 10 to 5.2 MK. Therefore, we deem this not physical as this would mean microflare 10, the brightest and most spatially complex microflare in this study, would not reach temperatures above

those found during its own pre-flare phase (6.6 MK). Moreover, including this broken power-law model does not resolve the incompatibility found between the synthetic flux comparison made with the Fe XVIII proxy emission.

It has been observed that during the impulsive phase of microflares that non-thermal emission can be present (Glesener et al. 2020). We, therefore, may expect this should be observed in microflare 10 since it is as bright as microflare 3, which does have strong evidence for non-thermal emission, despite the X-ray peak being missed due to NuSTAR's SAA passage. However, this analysis does not prove one way or another that a non-thermal component is present in microflare 10's spectra. The lack of clear non-thermal signatures could be due to the more complicated spatial evolution of microflare 10 with loops of various sizes and orientations being heated at different times; therefore, this could produce many thermal components which may lead to any non-thermal emission being less obvious or completely hidden.

Using the same method described for microflare 3, we approximate microflare 10's volume for each indicated time (Equation 3.3) finding upper volume limits for the initial rise, continued rise, and plateau to be 9.9×10^{26}cm^3, 1.4×10^{27} cm^3, and 3.6×10^{27} cm^3, respectively, with densities increasing from 1.5×10^{9} cm^{-3} to 2.2×10^{9} cm^{-3} (Equation 3.2). We then use Equation 1.16 with both microflare excess thermal models for each time to obtain the multi-thermal energy release from each microflare stage. From the initial rise to the plateau we obtain energies of $3.0^{+1.3}_{-0.1}\times10^{27}$ erg, $6.1^{+0.7}_{-0.4}\times10^{27}$ erg, and $1.6^{+0.2}_{-0.1}\times10^{28}$ erg. Therefore, we calculate that microflare 10 has the largest energy release from these data, with its plateau stage having an estimated equivalent GOES of A0.3 (scaled: A0.4, see Section 2.2), while still only being in its impulsive phase. Microflare 10's spectral fit values are summarised in Table 3.4.

3.8 Orbit 6

We do not find any impulsive emission indicative of a microflare, or any other event, during the last 1 h dwell focussed on the western solar hemisphere (right hand side) as seen from panel f, g, and h in Figure 3.1 and, therefore, was not studied further. However, this orbit could be used in future quiescent AR studies precisely because of the lack of X-ray and EUV activity as viewed from NuSTAR and SDO/AIA, respectively.

3.9 A Magnetic Perspective

From visual inspection of SDO/HMI images (Section 2.5.2), the vast majority of the ten microflares appear to have photospheric line-of-sight magnetic field activity (Figure 3.16). In eight of the ten microflares, we find the loop footpoints anchored in large unipolar regions of magnetic field and at least one footpoint in close proximity to a smaller mixed polarity region or the loop overlying two opposite polarity patches. The SDO/AIA 1600 Å and 1700 Å channels show bright chromospheric emission at the footpoint locations derived from the SDO/AIA 94 Å and Fe XVIII proxy channels in Figure 3.16. We show the bright response seen by the 1600 Å channel in Figure 3.17.

All NuSTAR microflares on 2018 September 9–10 originate in differently located sets of loops; however, their footpoints are all similarly located. Every labelled microflare takes place in the AR core and is east-west orientated with footpoints rooted in two large regions of opposite line-of-sight magnetic polarity. In addition, the AR does not drastically change orientation or general structure in X-rays, EUV, or magnetic field from the first to the last observation.

3.9.1 Photospheric Magnetic Flux Cancellation and/or Emergence

We show potential flux cancellation or emerging regions at eight microflare start times with red arrows on SDO/HMI magnetograms (Figure 3.16). Emission from the SDO/AIA 94 Å and hotter Fe XVIII channels is shown with green and blue contours, respectively. The levels are chosen to best display the microflaring loops in question and the mixed polarity regions around the footpoints of microflare 1, 2, 3, 5, 6, 8, 9, and 10 are indicated with red arrows. The final microflare, microflare 10, has several footpoint locations with mixed polarity regions and are labelled a, b, and c.

From the NuSTAR temperature and emission measure spectral analysis results and strong Fe XVIII presence, we find microflares similar to those previously studied where evidence of photospheric magnetic flux cancellation has been observed at their footpoints (Chitta et al. 2017a,b, 2018, 2019). Therefore, this study corroborates previous research which indicates that Fe XVIII intensity is correlated to photospheric magnetic flux emergence or cancellation (Asgari-Targhi et al. 2019; Chitta et al. 2020).

The presence of the mixed magnetic polarity regions close to the footpoints of some of these events allows the possible presence of similar geometries as to those described in models that involve photospheric flux cancellation being the mechanism for solar heating events (Priest et al. 2018; Syntelis et al. 2019). Studying photospheric flux cancellation can help provide insight into the mechanism and energy release behind solar flares and is an important area to study as the properties of any impulsive

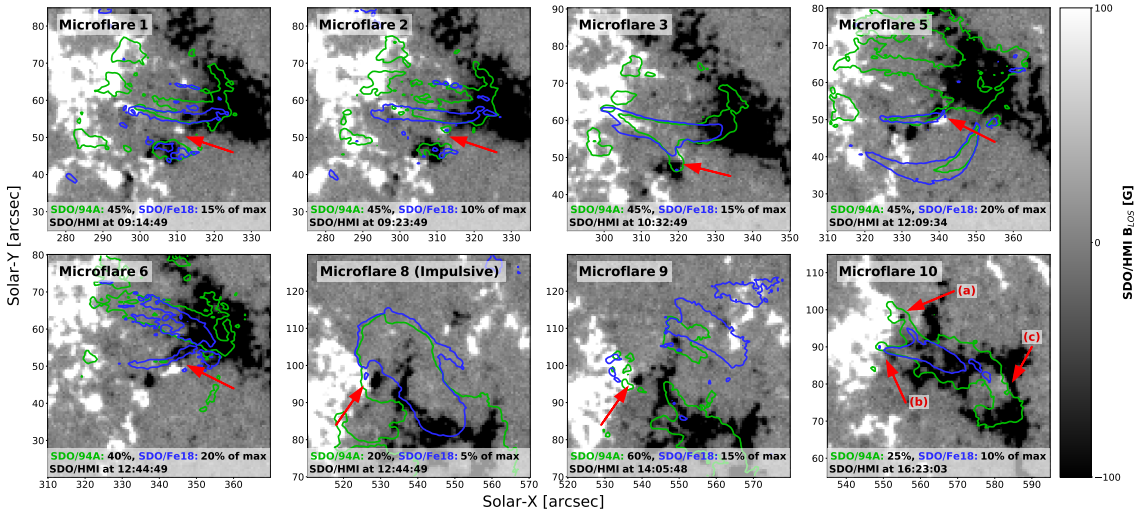

Figure 3.16: SDO/HMI magnetograms at the start of eight microflares (time shown) with SDO/AIA 94 Å (green) and Fe XVIII proxy (blue) contours showing the integrated coronal EUV microflare emission over the times previously stated, with an exception for microflare 8. Contour levels are chosen to appropriately show the loop brightening for each microflare and to help roughly locate footpoint locations. The red arrows indicate the mixed polarity regions in the proximity of each microflare's footpoints while the black and white colourmap represents negative and positive magnetic polarity, respectively. Microflare 8's impulsive phase was not observed by NuSTAR and so was determined to be at 12:40–12:50 UTC (September 10) from the Fe XVIII emission. This figure is taken from Cooper et al. (2021).

dynamic activity can provide estimates on the energy release from the phenomenon and whether or not this could be significant in the flaring process (Chitta et al. 2020).

Magnetic flux cancellation in the lower solar atmosphere may also be the mechanism by which several microflares can appear to occur in the same general loop structure (Chitta et al. 2020) which is observed multiple times throughout these data. Microflare 1 and 2 appear to originate from the same set of loops which overlies a newly emerging positive polarity. Similarly, we find microflare 5 and 6 also seem to be produced by the same general loop structure with a dynamic positive magnetic polarity area in the proximity of the western footpoint. In addition, we find a negative polarity region at the eastern footpoint of microflare 9 which could explain its presence in microflare 8's decaying loop structure which also has the negative magnetic region during its impulsive phase (Figure 3.16)

Although microflare 7 does not have a visible *parasitic* polarity[5] at either footpoint, the asymmetric loop heating observed (Figure 3.12, top left and middle panels) could be explained by magnetic flux emergence/cancellation at the western footpoint by an unresolved mixed polarity region (Chitta et al. 2017a). Chitta et al. (2017a) showed that unipolar regions as viewed by SDO/HMI can actually have oppositely polarised regions using instruments with greater spatial resolution and magnetic flux

[5]A small polarity region appearing to cause the flux cancellation of a larger magnetic flux region of the opposite polarity.

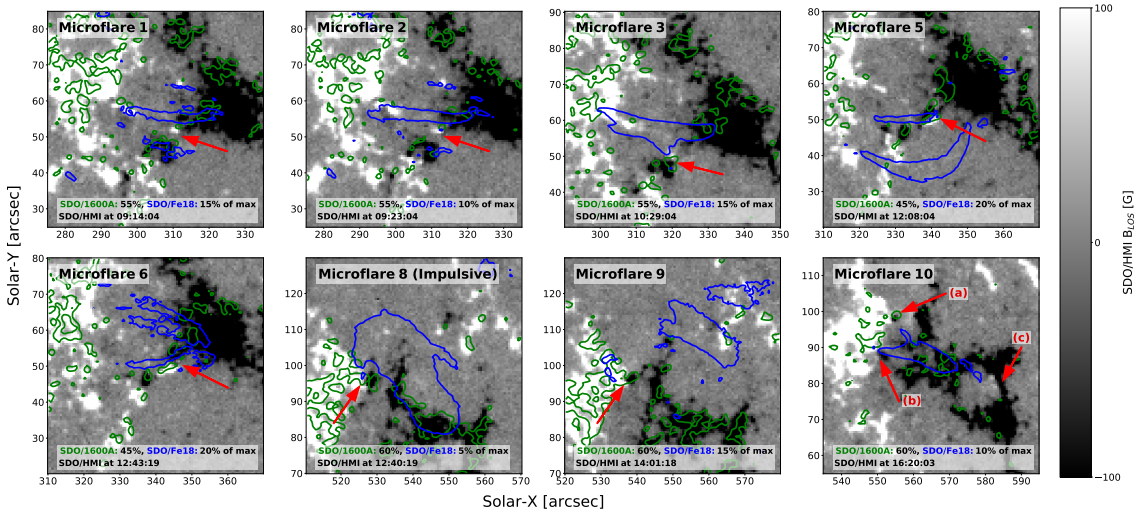

Figure 3.17: The same as Figure 3.16 but the green SDO/AIA 94 Å contours have been replaced by SDO/AIA 1600 Å emission. The contour levels for the UV channel are chosen to highlight the activity close to the loop footpoints.

cancellation at a single footpoint could be the reason a loop is heated asymmetrically (Testa & Reale 2020). A similar scenario may have occurred for microflare 6 and unresolved mixed polarity regions may also be present for microflare 4.

We only continue to investigate the footpoint magnetic configuration of the brightest two microflares (microflare 3 and 10) as these parasitic polarity areas are easily isolated from the surrounding regions of the same polarity and only appear to significantly interact with one area of the opposite polarity (Figure 3.16).

3.9.2 Microflare 3 and 10

The photospheric magnetic flux of the parasitic polarity of microflare 3 and each indicated for microflare 10 (a, b, and c) are shown in Figure 3.18. Both positive parasitic polarities (microflare 3's and microflare 10's c) indicate the presence of flux emergence and cancellation during the onset of their respective microflares (Figure 3.18, top left and bottom right). However, both of microflare 10's negative polarities (a and b) show a constant flux cancellation (Figure 3.18, top right and bottom left). We ensure no other magnetic flux with the same polarity as the parasitic region was included when producing Figure 3.18 and we only sum line-of-sight magnetic field strengths with magnitudes >17 G to exclude pixels below noise levels (Pesnell et al. 2012).

Flux cancellation rates are found to be on the order of 10^{14}–10^{15} Mx s^{-1} at the negative polarities

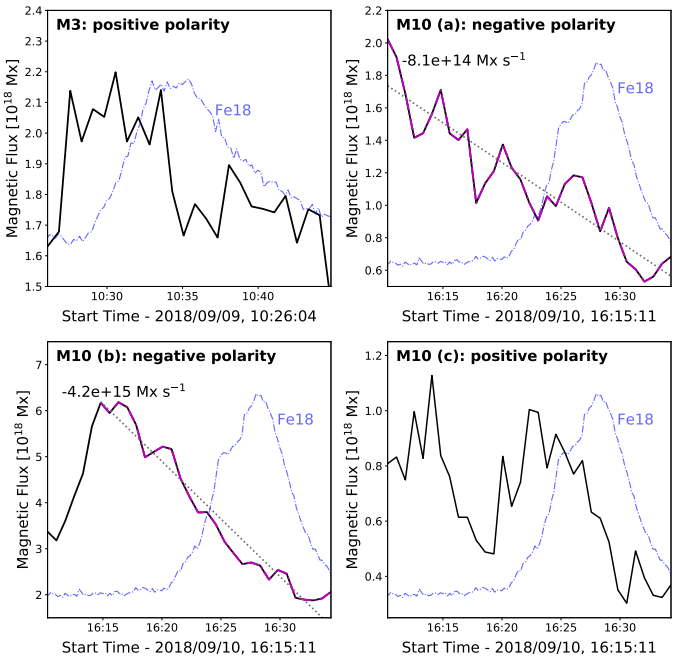

Figure 3.18: The evolution of microflare 3's positive (top left panel) and microflare 10's positive (bottom left panel) and negative (top right and bottom left panel) polarity, indicated in Figure 3.16. Microflare 10's steady negative magnetic flux cancellation is fitted (grey, dotted) with the resulting cancellation rate displayed in the top left of their respective panels. The magenta dash-dotted line represent the times considered for the linear fit and the AR Fe XVIII emission (blue, dash-dotted) for microflare 3 and 10 is also shown. This figure is taken from Cooper et al. (2021).

associated with microflare 10's footpoints. The magnetic flux cancellation from these locations is monotonic in nature and similar in magnitude to that correlated with brightenings and the production of A and B GOES class microflares (Chitta et al. 2017a, 2018). It has been estimated, depending on current sheet lengths and structure magnetic field strengths, that similar cancellation rates could release 10^{27}–10^{28} erg of magnetic energy over a 30 min period (Chitta et al. 2020). This would certainly suggest that the flux cancellation observed is sufficient to be responsible for microflare 10, a sub-A class flare with >15 min of flux cancellation at locations a and b shown in Figure 3.18.

3.10 Summary and Conclusions

In this chapter, we present the largest study of NuSTAR X-ray microflares that originate from a single AR. From this analysis, we significantly increase the number of analysed NuSTAR microflares in published works (discussed in Section 1.3.1) which aids in providing a more statistical perspective of flares and their physics at this scale, an example of which is shown in Figure 1.4.

The higher energy HXR NuSTAR emission of 4–10 keV shows more impulsive and earlier peaking time profiles when compared to the lower energy ranges of 2.5–4 keV and Fe XVIII for the majority of the microflares investigated here. This could indicate either the presence of hotter material being present in the early microflare phases and/or non-thermal emission in these sub-A class microflares, similar to the properties observed in far bigger and brighter flares.

We find each microflare presented unique challenges when attempting to isolate the enhanced AR emission or microflare excess; however, each flare is easily identified when using the HXR observations NuSTAR provides. A microflare could be investigated further once identified in X-rays as it could then be found and studied using the Fe XVIII proxy channel. The Fe XVIII emission could then be used to identify the specific loop structures producing each microflare which, in turn, allows us to direct our attention to corresponding structures in other SDO/AIA wavebands.

Good agreement can be found between time profiles produced from specific SDO/AIA Fe XVIII loop emission and those made from NuSTAR AR integrated 2.5–4 keV light curves with the 4–10 keV energy range giving an insight into the flare's more impulsive nature. We note that NuSTAR's HXRs are invaluable at identifying these incredibly weak events as they show more pronounced compared to surrounding material over a much larger area. Thus, emphasising the importance of being able to image, and observe spectra from, the Sun at these HXR ranges in order to observe events that are much weaker and fainter in EUV.

The inclusion of HXR data is imperative when considering microflare 1 and 2 or 5 and 6. Here we find that it is not the most obvious and brightest source in Fe XVIII (cyan contours in Figure 3.2 and grey contours in Figure 3.10 and 3.11) that is deemed to be the location for the microflaring emission, but a more transient and weaker Fe XVIII structure. This is only informed via the impulsive X-ray emission observed by NuSTAR and is important to consider when identifying the same, or dependent, structures between EUV and X-ray emission, particularly in any analysis which utilises image data.

Moreover, microflaring AR temperatures are confirmed to range from 3–4 MK plasma (Glesener et al. 2017; Wright et al. 2017; Hannah et al. 2019; Mitra-Kraev & Del Zanna 2019) up to ~10 MK (Duncan et al. 2021) when compared to other studies. Therefore, it is crucial to probe these plasma temperatures in the X-ray regime in order to be able to fully investigate the spatial, temporal, and spectral nature of the thermal and non-thermal microflare emission occurring. This emphasises the need for imaging X-ray spectroscopy for solar observations in order to probe both quiescent and flaring AR temperatures and, thus, solar atmospheric heating.

We find in eight out of ten microflares that loop footpoints appear to be in the proximity of, or rooted in, mixed photospheric magnetic polarity regions. The mixed polarity area provides the possibility of flux cancellation or emergence at the microflare footpoints which may explain the repeated heating

seen in microflares 1 to 2, 5 to 6, and 8 to 9. We also quantitatively show the presence of positive flux emergence during microflare 3 and 10's impulsive phase as well as two negative magnetic flux regions at two additional footpoint locations of microflare 10. Chitta et al. (2020) showed that magnetic flux emergence/cancellation in the photosphere that is interacting with overlaying coronal structures may be correlated to the onset of AR core microflares, performing a substantial role in the impulsive heating process. Flux cancellation rates of 10^{14}–10^{15} Mx s^{-1} are found for microflare 10's two negative parasitic polarity regions and are consistent with other studies (Chitta et al. 2017a, 2018).

We find that the more complicated loop structures produce the brightest events (microflare 3 and 10) with relatively simpler single loop structures producing weaker microflares. We are able to investigate the more spatially complex microflares temporally where microflare 3 demonstrates that the hottest temperatures are obtained in its impulsive phase. This is similar to what has been found in previous analysis and corroborates that this phase cannot go unnoticed when investigating the peak temperatures that flares obtain (Mitra-Kraev & Del Zanna 2019; Testa & Reale 2020). The same conclusion is harder to make with microflare 10 as NuSTAR entered night before the estimated peak; however, some of the highest temperatures were recorded during its impulsive stage, reaching $>$8 MK (Table 3.4).

Non-thermal emission is found to be present, injecting 1.3×10^{27} erg during microflare 3's rise time (Figure 3.5) while reaching temperatures up to 7 MK and having an equivalent GOES class of A0.1. The instantaneous thermal energy, 6.5×10^{27} erg (Table 3.3), is larger than the estimated non-thermal energy which indicates not all non-thermal emission is seen or that the non-thermal component may not be solely responsible for all of microflare 3's heating. However, we note that the method used to calculate microflare 3's volume give a conservative upper limit; therefore, by using a different method, or by utilising a filling factor, the thermal and non-thermal energy disagreement may become insignificant or disappear. We compare this microflare to the A5.7 (scaled: A8.1, see Section 2.2) microflare, the subject of Glesener et al. (2020), where we find microflare 3 to have considerably weaker non-thermal emission due to a higher break energy and steeper spectrum. Microflare 3 then becomes the faintest non-thermal AR X-ray microflare thus far observed by NuSTAR.

In this chapter, we discussed evidence supporting the presence of:

1. ten equivalent GOES sub-A class microflares with thermal energies between 10^{26}–10^{28} erg, reaching temperatures up to ~10 MK, comparable temperatures to those linked to hot transient coronal loops and microflare impulsive stages (Testa et al. 2020; Testa & Reale 2020);

2. one of the weakest non-thermal X-ray microflares in literature with an equivalent GOES class of A0.1; and

3. mixed polarity photospheric magnetic fields in the proximity of observed X-ray microflare footpoints, potentially explaining repeated microflare heating from the same general loop structure, with cancellation rates similar to those found in other microflares (Chitta et al. 2020).

There are many microflares in this data-set; however, one is of particular interest because of its incredibly faint emission and occurrence in isolation. We, therefore, investigate microflare 4 further in the following chapter and in Cooper et al. (2020).

Table 3.2: The isothermal fit parameters of 8 microflares on 2018 September 9–10 where the brightest two have been excluded. The estimated loop volume, thermal energy, and GOES equivalent class is displayed for each microflare. The FPM(s) used for spectral analysis along with the flare phase(s) is also indicated. All parameters are displayed with 1-σ equivalent uncertainties obtained from MCMC analysis. This table is taken from Cooper et al. (2021).

Microflare	Phase(s) of Microflare	FPM Used	Background/Pre-flare (T [MK], EM [$\times 10^{46}$ cm^{-3}])	Excess/Microflare (T [MK], EM [$\times 10^{44}$ cm^{-3}])	Loop Volume [cm^3]	Inst. Thermal Energy [erg]	GOES Class
1	All	B	$3.59^{+0.14}_{-0.20}$, $1.07^{+0.42}_{-0.19}$	$8.14^{+0.40}_{-1.68}$, $0.14^{+0.59}_{-0.05}$	2.2×10^{26}	$1.87^{+1.51}_{-0.26}\times10^{26}$	A0.001
2[a]	Rise/Peak	B	$3.17^{+0.07}_{-0.13}$, $2.99^{+0.42}_{-0.71}$	$5.17^{+0.61}_{-0.75}$, $4.24^{+9.37}_{-2.85}$	2.2×10^{26}	$6.53^{+3.47}_{-2.36}\times10^{26}$	A0.01
4[b]	All	B	3.20, 1.74	$6.66^{+0.35}_{-0.26}$, $0.80^{+0.26}_{-0.20}$	1.9×10^{25}	$1.08^{+0.11}_{-0.10}\times10^{26}$	A0.005[c]
5	All	A&B	$4.06^{+0.14}_{-0.48}$, $1.16^{+0.66}_{-0.10}$	$6.53^{+0.60}_{-0.39}$, $5.40^{+5.11}_{-3.18}$	7.9×10^{25}	$5.59^{+1.75}_{-1.67}\times10^{26}$	A0.03[c]
6	Rise/Peak	B	$3.38^{+0.07}_{-0.09}$, $3.32^{+0.83}_{-0.50}$	$8.72^{+0.91}_{-1.70}$, $0.16^{+0.51}_{-0.07}$	7.9×10^{25}	$1.30^{+0.83}_{-0.20}\times10^{26}$	A0.002[c]
7	All	B		$3.37^{+0.04}_{-0.03}$, $262.16^{+21.36}_{-25.28}$	5.2×10^{26}	$5.15^{+0.16}_{-0.20}\times10^{27}$	A0.2[c]
8	Decay	A		$4.14^{+0.03}_{-0.02}$, $89.22^{+5.56}_{-6.18}$	4.4×10^{27}	$1.07^{+0.03}_{-0.03}\times10^{28}$	A0.1
9	Peak	A&B		$4.30^{+0.04}_{-0.03}$, $76.38^{+4.57}_{-5.11}$	4.4×10^{27}	$1.03^{+0.02}_{-0.03}\times10^{28}$	A0.1

[a] Background component was constrained between 3–4 MK.

[b] Pre-flare parameters were fixed components during the microflare fitting.

[c] Scaled GOES equivalent class is A0.007, A0.04, A0.003, and A0.3 for microflare 4, 5, 6, and 7, respectively (see Section 2.2).

Table 3.3: Microflare 3's spectral fit parameters from FPMA and B simultaneous fits. The photon index below E_B is fixed at 2 and the pre-flare model was fixed during all microflare times. This table is taken from Cooper et al. (2021).

Phase of Microflare	Thermal Model (T [MK], EM [$\times 10^{44}$ cm^{-3}])	Non-thermal Model (E_B [keV], γ, Norm [ph keV^{-1} cm^{-2} s^{-1} at 1 keV])	Loop Volume [cm^3]	Inst. Thermal Energy [erg]	Non-Thermal Energy [erg]	GOES Class
Pre-flare	$4.09^{+0.02}_{-0.02}$, $63.00^{+3.00}_{-3.00}$					
Rise	$6.84^{+0.12}_{-0.09}$, $5.50^{+0.33}_{-0.42}$	$6.21^{+0.21}_{-0.43}$, $8.29^{+0.83}_{-1.52}$, $0.80^{+0.25}_{-0.19}$	3.5×10^{27}	$3.93^{+0.07}_{-0.09}\times10^{27}$	$1.27^{+0.66}_{-0.42}\times10^{27}$	A0.04[b]
Peak	$6.65^{+0.03}_{-0.03}$, $15.97^{+0.41}_{-0.44}$		3.5×10^{27}	$6.53^{+0.06}_{-0.06}\times10^{27}$		A0.1
Decay	$5.75^{+0.08}_{-0.06}$, $23.65^{+1.24}_{-1.41}$		4.9×10^{27}	$8.14^{+0.12}_{-0.14}\times10^{27}$		A0.1
Whole Time (Rise–Decay)	$6.54^{+0.02}_{-0.02}$, $11.60^{+0.20}_{-0.20}$	6.21, 8.29, 0.27[a]	4.9×10^{27}	$6.50^{+0.04}_{-0.04}\times10^{27}$	1.27×10^{27}	A0.07[b]

[a] Rise time broken power-law components are fixed and the normalisation parameter is scaled to the full time interval spectrum.

[b] Scaled GOES equivalent class is A0.06 for the rise phase and A0.1 for the whole time (see Section 2.2).

Table 3.4: Microflare 10's spectral fit parameters from FPMA and B simultaneous fits. The pre-flare model was fixed during all microflare times. This table is taken from Cooper et al. (2021).

Phase of Microflare	Thermal Model 1 (T [MK], EM [$\times 10^{46}$ cm^{-3}])	Thermal Model 2 (T [MK], EM [$\times 10^{44}$ cm^{-3}])	Loop Volume [cm^3]	Inst. Thermal Energy [erg]	GOES Class
Pre-flare	$3.05^{+0.04}_{-0.35}$, $1.70^{+1.99}_{-0.08}$	$6.60^{+0.20}_{-0.61}$, $0.38^{+0.40}_{-0.07}$			
Init. Rise	$4.15^{+0.12}_{-0.89}$, $0.16^{+0.40}_{-0.02}$	$8.14^{+0.06}_{-0.95}$, $0.58^{+0.71}_{-0.06}$	9.9×10^{26}	$2.98^{+1.25}_{-0.13}\times10^{27}$	A0.03[a]
Cont. Rise	$4.38^{+0.30}_{-0.30}$, $0.45^{+0.16}_{-0.10}$	$8.15^{+0.21}_{-0.70}$, $1.64^{+1.45}_{-0.38}$	1.4×10^{27}	$6.08^{+0.69}_{-0.43}\times10^{27}$	A0.1
Plateau	$4.30^{+0.17}_{-0.26}$, $1.29^{+0.32}_{-0.20}$	$8.08^{+0.13}_{-0.98}$, $3.58^{+4.84}_{-0.64}$	3.6×10^{27}	$1.59^{+0.19}_{-0.08}\times10^{28}$	A0.3[a]

[a] Scaled GOES equivalent class is A0.04 for the initial rise phase and A0.4 for the plateau phase (see Section 2.2).

4

NuSTAR Observation of a Minuscule Microflare in a Solar Active Region

The work in this chapter is based on the scientific research presented in Cooper et al. (2020).

We further investigate the weakest X-ray AR microflare currently in literature, microflare 4 from Chapter 3. This event exhibits NuSTAR's ability to examine some of the weakest X-ray microflares, helping us probe flare physics at incredibly weak scales; for example, we can obtain the temperatures reached and investigate the potential for non-thermal emission over the course of these events.

4.1 Time Profile and Imaging

Initially, this microflare was discovered by inspecting NuSTAR's FOV integrated light curve (Figure 3.1), shown in greater detail in Figure 4.1 (top left panel). A small 'bump' is observed in the 2.5–4 keV X-ray range at ~11:04 UTC which becomes very evident in the higher 4–7 keV energy range. The microflare is not clearly observable in any native SDO/AIA channel but becomes slightly clearer when producing the Fe XVIII proxy channel (Section 2.5.1). The SDO/AIA 94 Å image does show a loop structure; however, it is difficult to distinguish this microflaring loop from surrounding sources and only the Fe XVIII proxy emission shows the presence of plasma heating (Figure 4.1, bottom middle and left panels). Due to the position of AR12721 during microflare 4 (Figure 4.1, top right panel), we find that only NuSTAR's FPMB can be used for detailed analysis as the AR is obscured by FPMA's chip-gap which can only be used qualitatively to corroborate the microflare in time and space.

We note that the higher energy range plotted for NuSTAR (4–7 keV, Figure 4.1) rises before the lower energy emission in 2.5–4 keV and Fe XVIII which is expected flaring behaviour. This could be indicative of either hotter plasma being present early in the microflare's evolution or the presence of an accelerated electron distribution producing non-thermal emission.

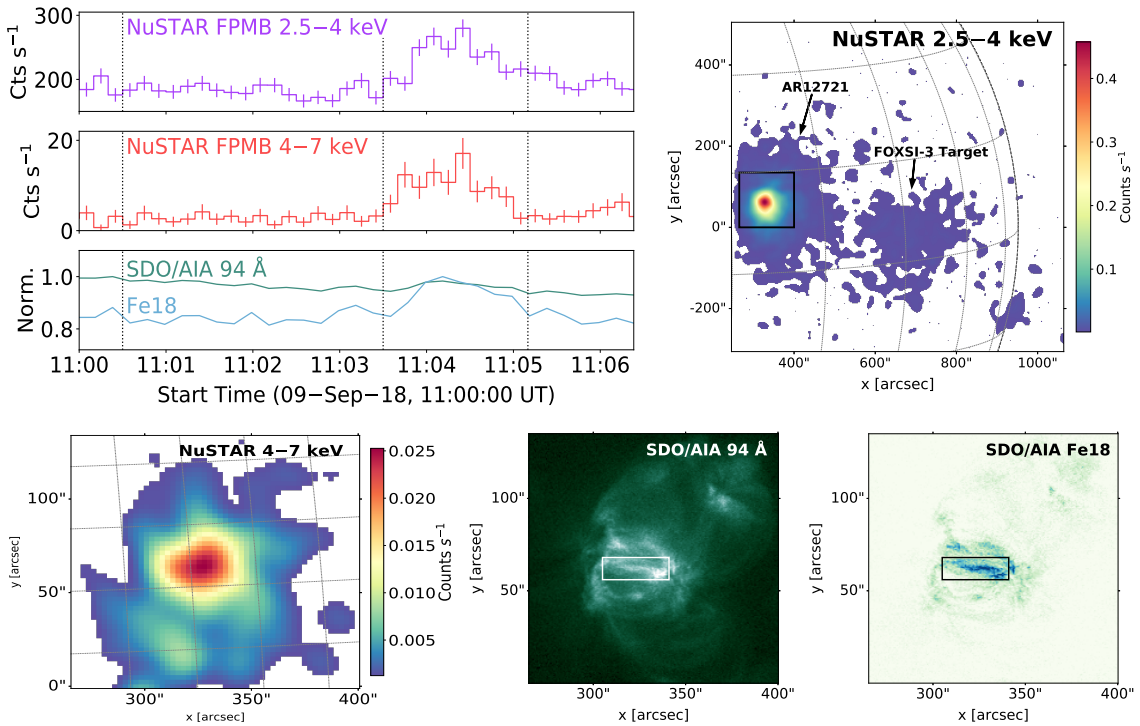

Figure 4.1: NuSTAR livetime corrected 2.5–4 keV (purple) and 4–7 keV (red), and SDO/AIA 94 Å (green) and Fe XVIII (blue) time profiles of the 11:04 UTC microflare (top left panel). The black box in the top right panel, and the image in the bottom left panel, shows the region used to create the NuSTAR light curves while the area used to produce the maximum normalised SDO/AIA 94 Å and Fe XVIII time profiles is shown as a white and black box in the bottom middle and right panels, respectively. The pre-flare (11:00:30 to 11:03:30 UT) and microflare (11:03:30 to 11:05:10 UT) times are shown with vertical dotted lines in the light curve panel where all images are integrated over the microflare time. Both NuSTAR maps have been livetime corrected and are smoothed with a Gaussian filter with a FWHM of ~15″. This figure is taken from Cooper et al. (2020).

A single pointing correction is all that is required across all relevant times for the NuSTAR image to become co-aligned with SDO/AIA, found using the Fe XVIII proxy map shown in Figure 4.1 (bottom right panel) through cross-correlation. As with most NuSTAR-SDO/AIA co-alignment, we still assume a conservative shift uncertainty of 10″ due to emitting structures being less defined in X-rays.

We show the shifted NuSTAR emission during the pre-flare and microflare times on time averaged Fe XVIII backgrounds via absolute contour levels (Figure 4.2). We use the same energy ranges and a consistent colour scheme as those used in Figure 4.1 and the NuSTAR emission has been deconvolved using the Richardson-Lucy method (see Section 1.4.1 and Richardson 1972). The lower energy, 2.5–4 keV, and higher energy, 4–7 keV, X-ray ranges are deconvolved with 20 and 10 iterations, respectively. The lower NuSTAR energy range shows some general AR emission during the pre-flare time which then increases in brightness during the microflare time. The 4–7 keV range shows a significant increase in emission during the microflare time at the same location as the lower energy channels.

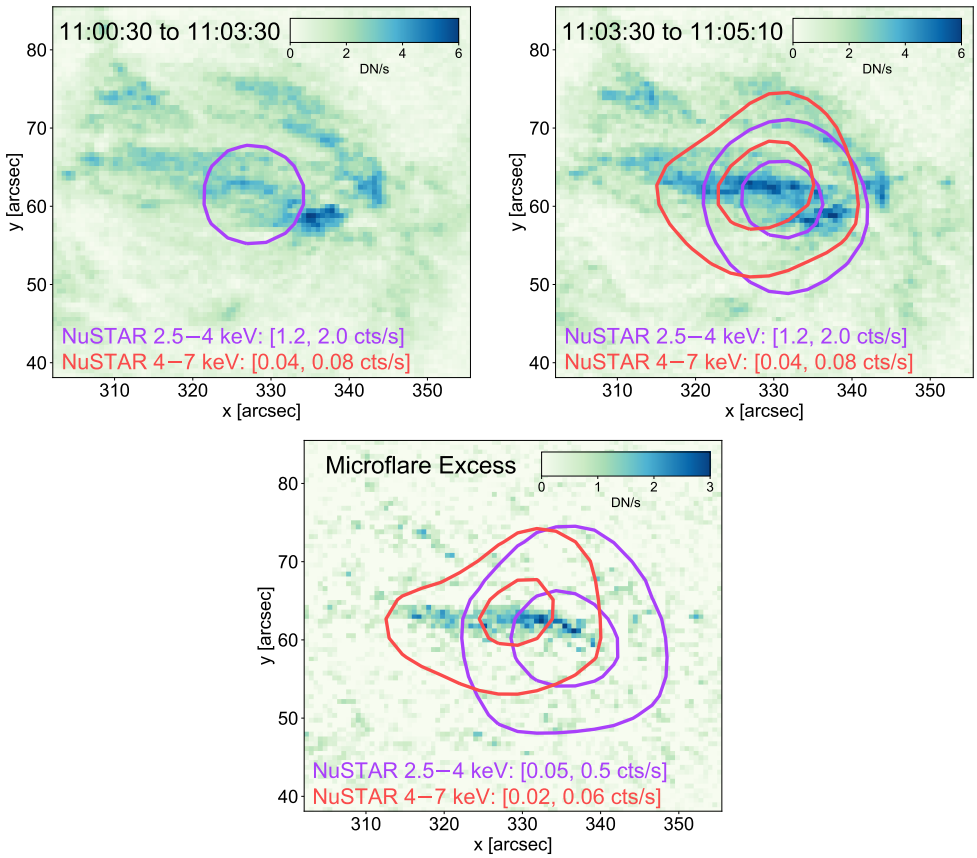

Figure 4.2: Pre-flare (top left panel) and microflare (top right panel) SDO/AIA Fe XVIII average emission with NuSTAR absolute contour levels for 2.5–4 keV (purple) and 4–7 keV (red) at the levels shown. The bottom panel shows the microflare excess and is the result of subtracting the pre-flare emission from that seen during the microflare time (top right minus top left panel). This figure is taken from Cooper et al. (2020).

Investigating the heating that takes place during the microflare time, we subtract the pre-flare from the microflare emission to estimate the microflare excess—the emission solely produced by the microflare—which is shown in Figure 4.2 (bottom panel). This illuminates an elongated loop structure which is clear in Fe XVIII and the 2.5–4 keV X-ray range. In addition, the 4–7 keV source is similarly elongated and is situated slightly to the left of the 2.5–4 keV source. This difference in location may not be significant, however, as it is within NuSTAR's spatial resolution (Grefenstette et al. 2016). Note that even though 4–7 keV emission is present during the pre-flare time, contours for this energy range are not created in Figure 4.2 (top left panel) due to the X-ray count rate being lower than the smallest valued contour.

All first steps in analysis—time profiles in Figure 4.1, images in Figure 4.2, and the spectral fitting results shown in Figure4.4—show undoubtedly the existence of a small and weak event. However, the

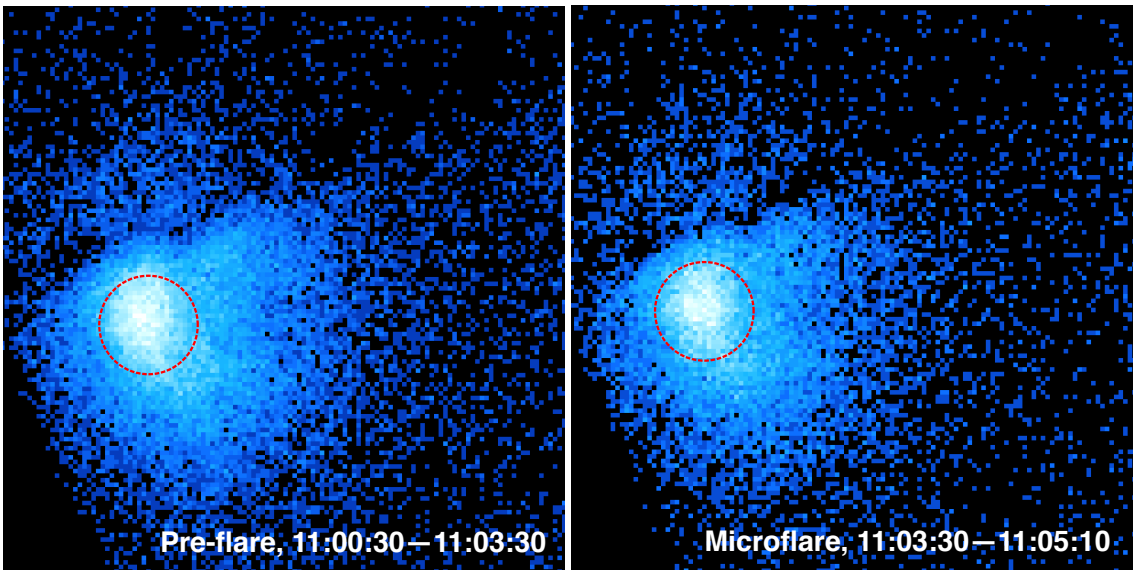

Figure 4.3: SAOImageDS9 images of the pre-flare (left panel) and microflare (right panel) times as seen by NuSTAR in RA-Dec coordinates. The red dashed circle in each panel (radii of 26.5″) indicates the region used for the spectral fitting and the colourmap is plotted on a log scale. The detector's edge and the chip-gap can be seen during both times to the left of and above the source, respectively.

physical loop length of 20″ estimated from Figure 4.2 is not uncommon for brighter X-ray microflares (Glesener et al. 2017; Hannah et al. 2008).

4.2 NuSTAR Spectral Fitting

Utilising NuSTAR's spectroscopic capabilities, we fit the FPMB microflare spectrum with a single APEC thermal model (see Section 5.3.2). Figure 4.3 shows a red circular region of radius 26.5″ centred on the brightest emission for the pre-flare (left panel) and microflare (right panel), produced using SAOImageDS9[1]. The area enclosed in the red circle is used to produce each spectrum depicted in Figure 4.4 via the NuSTAR data pipeline (Section 2.1.4).

We fit the pre-flare spectrum (11:00:30 to 11:03:30 UTC) with one thermal model using the XSPEC software (see Section 5.3.2 and Arnaud 1996). We find a temperature of 3.2 MK and an emission measure of 1.7×10^{46} cm^{-3} (Figure 4.4, left panel). As described in Chapter 3, these are typical non-flaring AR values measured by NuSTAR (Wright et al. 2017; Glesener et al. 2017; Hannah et al. 2019). This fitted pre-flare model was then used as a fixed component when fitting the microflare time (11:03:30 to 11:05:10 UTC, Figure 4.4, right panel). We then apply an additional thermal model beyond

[1]An image display tool for astronomical data, https://sites.google.com/cfa.harvard.edu/saoimageds9

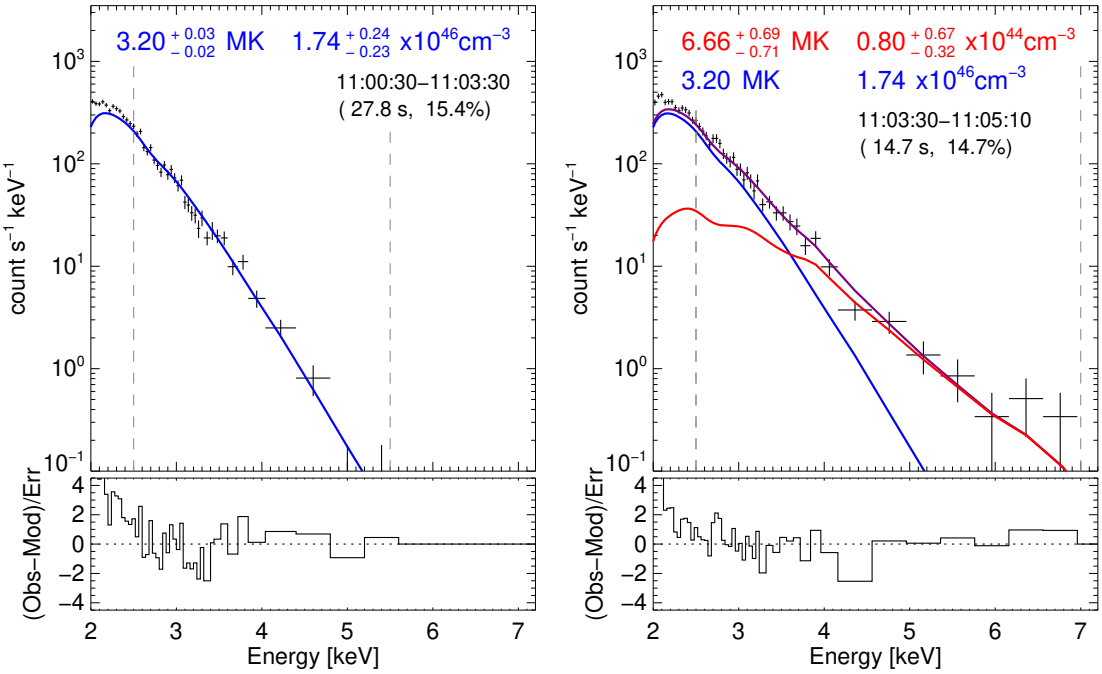

Figure 4.4: Pre-flare (left panel) and microflare (right panel) thermal model spectral fits (microflare 4, Section 3.4.2). The pre-flare spectrum is fitted with a single thermal model (blue) which is fixed during the microflare fitting while a second thermal model (red) fits the microflare excess. Both the fixed pre-flare and the microflare model are summed to produce the total purple model. The temperatures and emission measures of both models are shown along with the time range, duration, and livetime of the observation. The fitting range is shown with vertical dashed lines and the errors shown indicate a 90% confidence region from MCMC analysis. This figure is taken from Cooper et al. (2020).

the pre-flare one to fit the microflaring emission. The observation of this microflare has significantly higher livetime than most previously studied NuSTAR microflares, with a livetime of ~15% compared to 1–5%. The higher livetime results in better sensitivity to the hotter material emission, helping to better constrain the hot microflare excess.

The microflare produces a harder spectrum when compared to the pre-flare with the hot excess dominating >4 keV, corroborating the analysis performed in Section 4.1. We find the microflaring emission is well represented with a temperature of 6.7 MK which is hotter than or similar to other NuSTAR microflares but with an order of magnitude smaller emission measure of 8.0×10^{43} cm^{-3} (see Figure 1.4 and Glesener et al. 2017; Hannah et al. 2019). The temperature of the microflare excess is also consistent with the SDO/AIA Fe XVIII emission observed.

The spectral fitting result for the microflare time does not change significantly when varying the fixed pre-flare component to the extremes of its small uncertainty range. In addition, we investigate how the NuSTAR spatial region selection affects the resulting spectrum. We find that the region used is

small enough so that the surrounding AR emission does not dominate over the microflare and large enough to encompass the higher energy counts produced by the event. Again we note the fact that the temperature and emission measure are inversely related when fitting meaning that the upper/lower temperature range corresponds to the lower/upper emission measure range.

Using the same IDL routine, `goes_flux49.pro`, described in Section 2.2 and 3.3 with the temperature and emission measure, we calculate an estimated GOES flux in the 1–8 Å waveband of 5×10^{-11} W m^{-2} (scaled: 7×10^{-11} W m^{-2}, see Section 2.2). This produces the previously stated equivalent GOES class from Section 3.4.2 on the order of a thousandth of an A class flare (A0.005 or scaled to A0.007).

4.3 Thermal Energy

We use Equation 1.15 from Section 1.4.4 in combination with the temperature and emission measure excess fit to estimate the instantaneous thermal energy release of the microflare. Using the higher spatially resolved SDO/AIA Fe XVIII images in Figure 4.2 (bottom panel), initially we estimate the microflaring loop to be 22″ by 2″ or approximately 1.6×10^{9} cm by 1.3×10^{8} cm. We then utilise a half-torus geometry giving a volume of 3.2×10^{25} cm^{3}.

With a microflare excess temperature and emission measure of 6.6 MK and 8.0×10^{43} cm^{-3}, respectively, we obtain a thermal energy of $1.4^{+0.3}_{-0.2}\times10^{26}$ erg. We note that this volume and energy is larger than the final one stated in Section 3.4.2 for microflare 4. This is due to the volume estimate here being undoubtedly an upper limit as the Fe XVIII excess (Figure 4.2) subtly suggests that the microflare emission is concentrated to the right of the defined loop structure. We find further support for a smaller loop size when performing additional analyses, described in Section 4.5. This volume then produces a thermal energy upper limit especially when a filling factor is not considered for this event.

Even if the energy is significantly overestimated, the thermal energy is still lower than the previous smallest microflare observed by NuSTAR (Hannah et al. 2019), which obtained a higher emission measure but was cooler with a GOES class of A0.02 (scaled: A0.03, see Section 2.2). Energies on the order of 10^{26} erg have been obtained from EUV observations of heated magnetically braided loops (Cirtain et al. 2013); however, this was from plasma up to 4 MK.

4.4 Non-thermal Limits

We do not see any evidence that an additional model is required above those already used in Figure 4.4 (right panel) to account for emission from a non-thermal source. So, we use the approach described in Wright et al. (2017) to obtain upper limits on any non-thermal emission present in the NuSTAR microflare's spectrum. We do this by simulating a NuSTAR spectrum from the fitted thermal models and the instrument spectral response.

4.4.1 Exploring the Thick Target Parameter Space

We perform this process in OSPEX (Section 5.3.1) as we require the thick target model `f_thick2.pro`[2] (see Section 1.4.5) to represent the potentially hidden non-thermal electron emission. We start by assuming the total microflare thermal model shown in Figure 4.4 (right panel) perfectly represents the emitting material that NuSTAR is detecting. We then take the temperatures and emission measures from the XSPEC fitted APEC thermal models and recreate the total model in OSPEX using two `f_vth.pro`[3] thermal models in photon-space. We then use NuSTAR's spectral response matrix (Section 2.1.5) to transform our model into count-space with units counts s^{-1} keV^{-1}. See Chapter 5 for more detail on this transformation and the complications it poses.

We then take the count rate model (R) and find the total number of counts (N) from the overall model via

$$N = round\left(\sum_i R_i \times \delta E_i \times \delta t \times livetime\right) \quad \text{[counts]}, \tag{4.1}$$

where i indicates the energy bin with width δE_i, δt is the true exposure of the observation, and the $livetime$ being the livetime fraction (ontime divided by δt, see Section 2.1.2). The term $\delta t \times livetime$ gives the effective exposure, we use NuSTAR's native energy binning where all bins have a width of 0.04 keV, and *round* indicates the total should be rounded to the nearest integer. The total number of counts is then kept consistent between all iterations of the synthetic spectra about to be created from the double thermal total model.

A maximum normalised cumulative sum of the count spectrum ($R \times \delta E \times \delta t \times livetime$) gives a probability distribution function between 0 and 1. This cumulative probability distribution is then randomly sampled N times, building up a discretised synthetic NuSTAR spectrum for the microflare; we label this spectrum A (see Figure 4.5, left panels).

[2] https://hesperia.gsfc.nasa.gov/ssw/packages/xray/idl/f_thick2.pro
[3] https://hesperia.gsfc.nasa.gov/ssw/packages/xray/idl/f_vth.pro

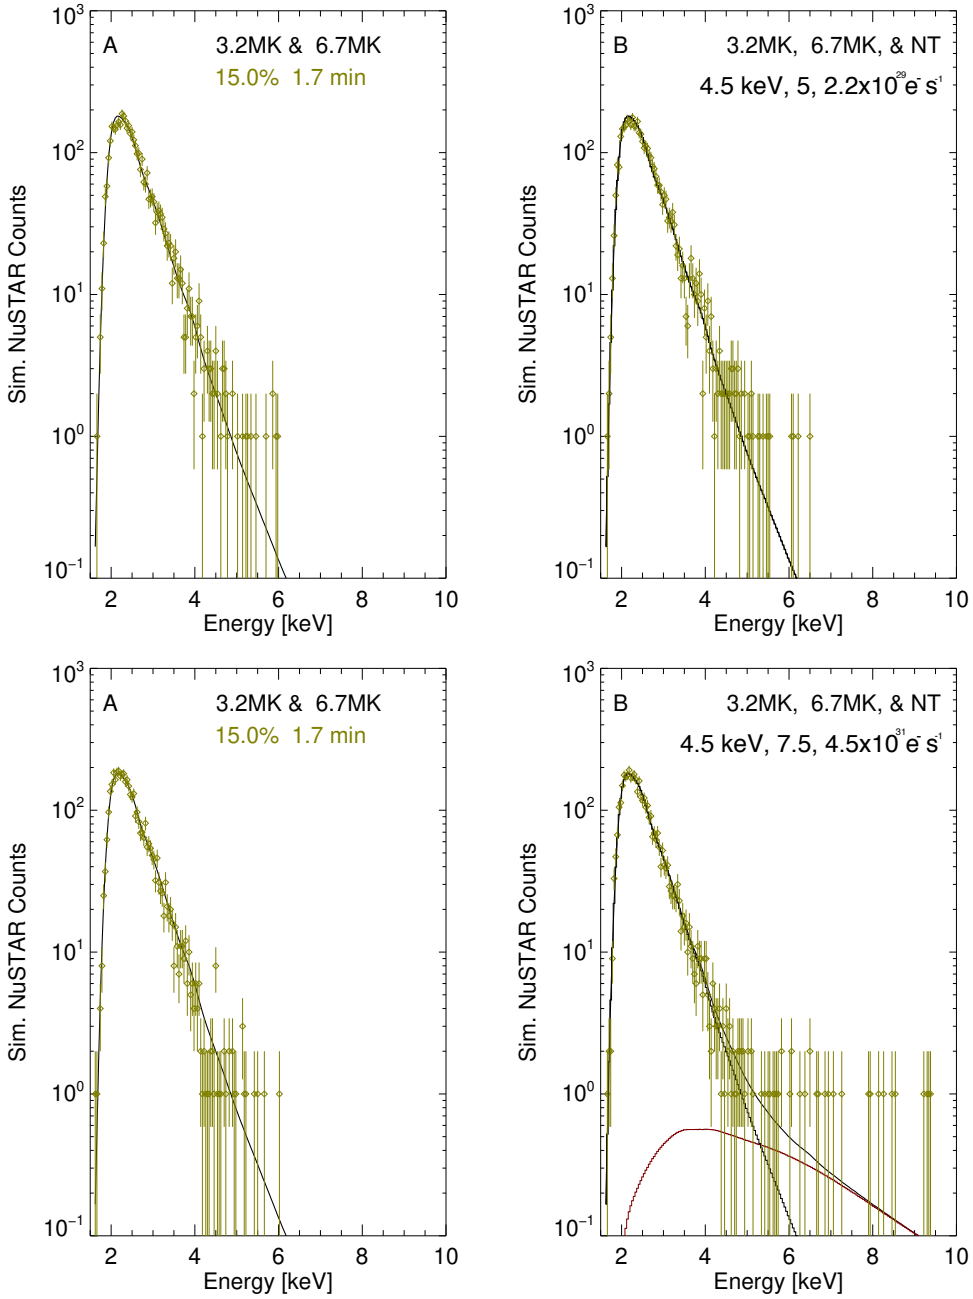

Figure 4.5: Two iterations of a NuSTAR synthetic counts spectrum for microflare 4 (left panels) produced from the microflare time fitted parameters displayed in Figure 4.4. The top right panel shows the synthetic spectrum made from the same thermal fitted parameters but also includes a thick target (NT) component with an electron index of 5, low-energy cut-off of 4.5 keV, and an electron flux of 2.2×10^{29} e$^-$ s^{-1} where the bottom right synthetic spectrum shows an electron index of 5, low-energy cut-off of 7.5 keV, and an electron flux of 4.5×10^{31} e$^-$ s^{-1}. The former introduced non-thermal model is consistent with a null detection, whereas, the latter is not.

We then perform the same process except we include the thick target model with the two thermal ones resulting in another synthetic spectrum (B, Figure 4.5, right panels). The goal is to compare both spectra and assess if the thick target emission with a specific electron flux, electron index, and low-energy cut-off (see Section 1.4.5) could have been produced by the source but completely hidden in the statistical noise of the measured counts.

We determine spectrum B to be possible if its additional non-thermal counts remained within A's Poissonian uncertainty $\leqslant$7 keV (i.e., B contributes no more than $\sqrt{A_{counts}}$ over the span of where the real spectrum has counts in Figure 4.4) and gave fewer than 4 counts $>$7 keV, being consistent with a null detection to 2-σ (Gehrels 1986).

We ensure to search over a wide range thick target parameter space to find the limit of possible B spectra while staying in a sensible range relative to the microflare spectrum. The ranges chosen are 10^{-29}–10^{-33} e^- s^{-1}, 5–9, and 4.5–12 keV for the electron flux, electron index, and low-energy cut-off in steps of $10^{0.05}$ e^- s^{-1} (multiplicative), 1, 0.5 keV, respectively.

Once the successful parameter combinations are found we then calculate their associated non-thermal power to investigate whether energy input from accelerated electrons is sufficient to power the microflare.

4.4.2 Non-thermal Energy from Single Power-law Electron models

We have already discussed a method to obtain a non-thermal power estimate from a photon broken power-law assumed to be produced by accelerated particle emission (Section 1.4.5); however, since OSPEX has the thick target model in its library, we can calculate the non-thermal power directly from the fitted electron distribution parameters.

The non-thermal power from a single power-law of electrons is calculated with

$$P(\geqslant E_C) = 1.6 \times 10^{-9} \frac{\delta - 1}{\delta - 2} N_N E_C \quad [\text{erg s}^{-1}], \tag{4.2}$$

where δ is the electron index, N_N is the non-thermal electron flux, and E_C is the low-energy cut-off for the electron distribution (Section 1.4.5 and 1.4.5). Therefore, we calculate the associated power with the parameters capable of producing emission consistent with a null NuSTAR detection for the microflare using Equation 1.31.

Figure 4.6 (left panel) shows the distribution of non-thermal powers obtained from the thick target parameters consistent with a null detection where we find a considerable number showing consistency with the heating required over the microflare time ($\sim 10^{24}$ erg s^{-1}). We then plot the upper limits of the

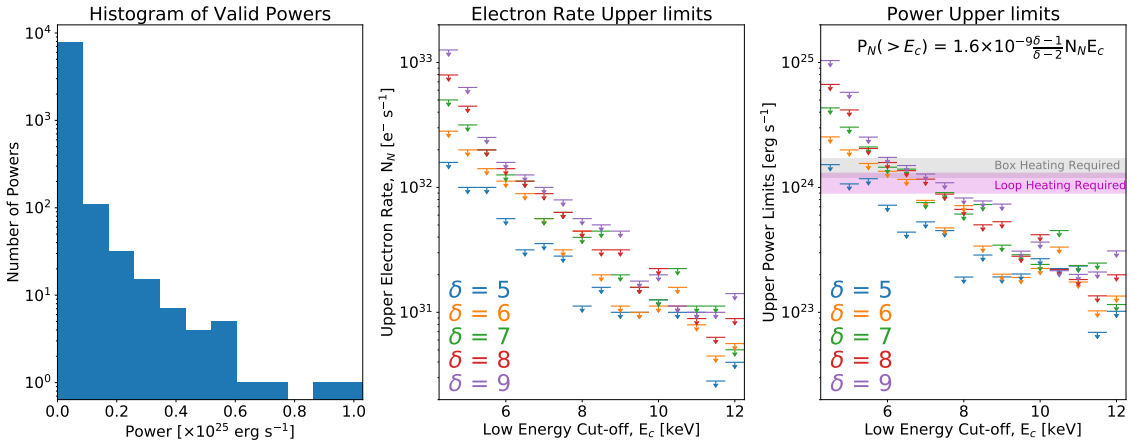

Figure 4.6: The non-thermal model parameters that met the null detection criteria. A histogram of single power-law electron distribution powers obtained from electron power-law indices (δ), total electron fluxes (N_N), and low-energy cut-offs (E_C) that are consistent with a null detection for the microflare (left panel). Upper limits found on the electron rate (middle panel) and power (right panel) at each low-energy cut-off and electron index investigated. The different colours indicate the electron index the limits are associated to and the non-thermal power equation is included in the right panel. The heating required for the microflare thus far is indicated by the 'Box Heating Required' grey region (right panel) where the magenta, 'Loop Heating Required' area shows the required heating for a smaller microflare volume considered in Section 4.5.

energies for a range of electron power-law indices (δ) over the low-energy cut-off range investigated (right panel); here we find that the required heating, shown by the grey shaded region, is provided from low-energy cut-offs below 6 keV with a power-law index of $\geqslant$6.

We find lower upper limits for this microflare than those upper limits found for more energetic microflares in Wright et al. (2017), which is to be expected. Our highest power here ($\sim 10^{25}$ erg s^{-1}) is below the lower end of non-thermal powers calculated in a previous study of RHESSI microflare energetics (Table 1 in Hannah et al. 2008).

The electron power-law index and the low-energy cut-off limits are consistent with those obtained in previous NuSTAR microflare studies, Glesener et al. (2020) and Wright et al. (2017), with only the total electron flux being significantly smaller. This behaviour may be expected as the X-ray peak emission in this microflare is orders of magnitude smaller than ones previously discussed.

However, these results do not appear to be consistent with the events presented in Testa et al. (2014), who investigated a nanoflare heating model and coronal loop footpoint brightenings in UV. Their model required low-energy cut-offs at ~10 keV, higher than found here, to match observation. This potentially suggests the microflare studied here and the modelled nanoflares in Testa et al. (2014) are different types of events and may not be comparable. The microflare electron distribution is not consistent with energies as high as the proposed nanoflare model despite the microflare being an

order of magnitude larger in energy release. Indicated by Figure 4.6, a similar electron power-law index range has been used between both analyses (δ =5–10) but only low-energy cut-offs below 6 keV are consistent with the microflare.

4.5 Multi-thermal Microflare Analysis

Microflares heat plasma to a range of temperatures with NuSTAR being sensitive to the hottest ones present which is why the NuSTAR X-ray spectrum is well fitted by an isothermal model. We can investigate the multi-thermal nature of the pre-flare and microflare emission by calculating the emission measure distribution.

4.5.1 Differential Emission Measure from Regularised Inversion

In order to calculate the differential emission measure[4] (DEM, see Section 1.4.4), from a single pixel or integrated region we use the approach described in Hannah & Kontar (2012)[5] which can also be applied to full image maps as described in Hannah & Kontar (2013)[6].

An observable, g, for a given instrument, i, is

$$g_i = \int_T K_i(T)\xi(T)dT + \delta g_i, \tag{4.3}$$

where $K_i(T)$ is the temperature response of the instrument (see Chapter 2), $\xi(T)$ is the differential emission measure of the plasma (see Section 1.4.4), and δg_i is the error in the observable.

To solve Equation 4.3, we can consider the problem of

$$\chi^2 := \left\| \frac{g - K\xi(T)}{\delta g} \right\|^2 = min, \tag{4.4}$$

where g and K contain all instrument observables and temperature responses, respectively, and $\|v\|$ is the Euclidean (L2) norm, defined $\|v\|^2 = v^\intercal v = \sum v_i^2$ (Judge et al. 1997). Equation 4.4 does not provide a unique solution for $\xi(T)$ and is an ill-posed inverse problem (Jefferies et al. 1972; Craig & Brown 1977) which is made even more difficult when including various observation and instrument uncertainties.

[4]Determined by the number of free electrons and electron density of a plasma that is in thermal equilibrium and optically thin, differentiated with respect to temperature (Phillips et al. 2012, pp. 94 and 144).
[5]Original IDL code available here https://github.com/ianan/demreg/tree/master/idl_org
[6]Original IDL code available here https://github.com/ianan/demreg/tree/master/idl with a Python version at https://github.com/ianan/demreg/tree/master/python

Hannah & Kontar (2012) then go on to use Tikhonov regularisation (Tikhonov 1963) and apply linear constraints to obtain a solution for $\xi(T)$. This modifies Equation 4.4 into

$$\chi^2 + \Theta = \left\| \frac{g - K\xi(T)}{\delta g} \right\|^2 + \lambda \left\| L\xi(T) \right\|^2 = min, \tag{4.5}$$

where the 'smoothing function' (Θ) helps reduce noise amplification and control the regularisation by subjecting Equation 4.5 to $\Theta \leqslant constant$. The terms λ and L represent the regularisation parameter and the constraint matrix, respectively.

The constraint matrix, L, is used to control the degree of regularisation. If L is proportional to the identity matrix then the solution is a minimised $\xi(T)$ that explains the observation, this is zeroth-order regularisation (L_0). The order of L can also be increased to L_1 or L_2 which are related to the first and second derivatives with respect to ξ which will favour flatter or smoother solutions, respectively.

We now find that Equation 4.5 does have a possible and unique solution unlike Equation 4.4. This allows for ξ to be expressed in terms of λ where Generalised Singular Value Decomposition is used (GSVD; Hansen 1992).

A suitable choice of regularisation parameter λ can be obtained using the Morozov's discrepancy principle (Morozov 1967), controlling χ^2 in data-space[7]. Therefore, when using real data with observable uncertainties of δg, we expect the optimal λ value to produce a DEM, ξ_λ, that meets the requirement

$$\left\| g - K\xi_\lambda(T) \right\|^2 - \rho \left\| \delta g \right\|^2 = min, \tag{4.6}$$

where ρ is a customisable parameter, ideally with a value of 1. Using GSVD, the leftmost expression ($\left\| g - K\xi_\lambda(T) \right\|^2$) is solved by optimising for λ with respect to Equation 4.6; therefore, allowing the DEM to be calculated for a given set of observations.

We use this DEM reconstruction method (Hannah & Kontar 2012) in the analysis of the incredibly weak microflare. This allows us to investigate the multi-thermal nature of the microflare and the excess hot material above the quiescent AR it produces.

4.5.2 Emission Measure Distribution

Because of the microflare's small scale, different instrument pixel size, co-alignment, and lack of defined features, we simplify the analysis and opt to calculate the DEM over an integrated area instead

[7] https://github.com/ianan/demreg/blob/master/notes/demreg_maths.pdf

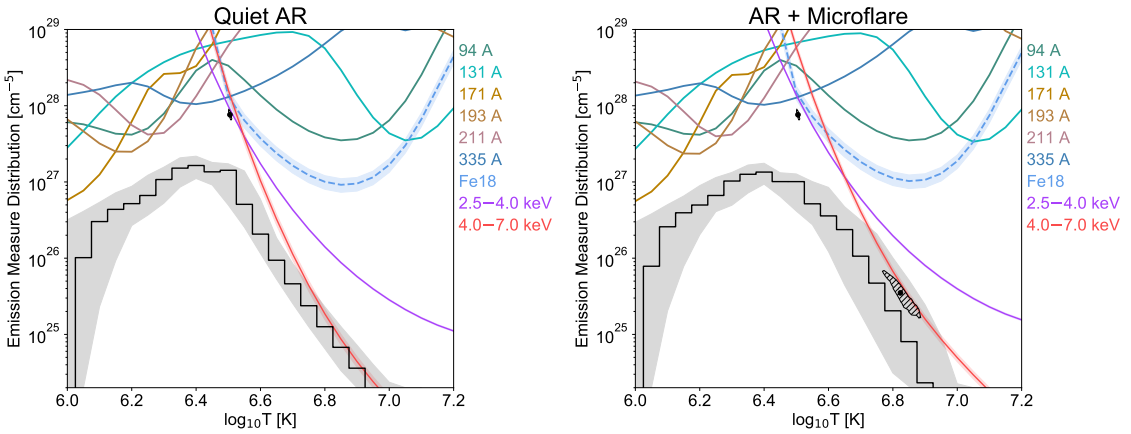

Figure 4.7: Calculated emission measure distributions (black) for the pre-flare time (left panel) and microflare time (right panel) with NuSTAR and SDO/AIA loci curve plotted. The uncertainty range for the NuSTAR loci curves (purple and red), the Fe XVIII curve (dashed, blue), and the emission measure distribution (grey) are all indicated with their respective shaded regions. The boxed region shown in Figure 4.1 (bottom right panel) was used to calculate the SDO/AIA and NuSTAR instrument loci curves. The spectral fit values are indicated with their hatched 90% confidence regions. This figure is taken from Cooper et al. (2020).

of on a pixel by pixel, full map basis. Therefore, we use the method described in Hannah & Kontar (2012).

In order to gain information of the emission measure at different temperatures the emission measure distribution (EMD) is calculated using the regularized inversion approach discussed in Hannah & Kontar (2012). The differential emission measure (DEM) is calculated and then multiplied by the temperature bin width. This produces an EMD in units of cm^{-5}. The EMD can be seen in Figure 4.7 as a black histogram with its uncertainty range in grey.

Figure 4.7 shows instrument emission measure (EM) loci curves, coloured, which are calculated by dividing the flux observed with the detector's temperature response. Figure 4.7 (left panel) shows that there is good agreement between the NuSTAR and Fe XVIII loci curves and the spectral fit temperature and emission measure from Figure 4.4 at the pre-flare time, depicting a more isothermal scenario (Section 1.4.4). This indicates that it is similar emission seen by both NuSTAR and SDO/AIA (Hannah et al. 2019). However, during the microflare time (Figure 4.7, right panel) the loci curves separate due to the higher temperature component introduced by the event which is corroborated in Figure 4.4 (right panel).

The EMD (Figure 4.7) and the spectral fit values from Figure 4.4 are consistent with each other. In the pre-flare/quiet AR time the EMD peaks at the spectral fit value and falls off quickly with increasing temperature (Figure 4.7, left panel). Whereas, during the microflare time the EMD has a slower decrease due to the presence of hotter material (Figure 4.7, right panel). This is similar behaviour as to

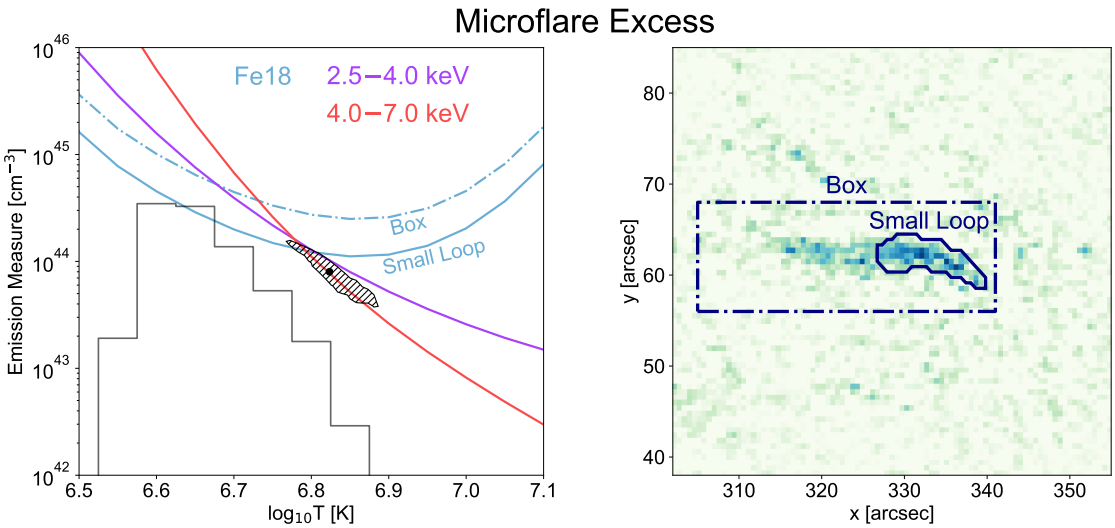

Figure 4.8: Microflare excess NuSTAR (purple, red) and SDO/AIA Fe XVIII (blue) EM loci curves in comparison to the microflare excess EMD (grey). SDO/AIA Fe XVIII loci were obtained for the 'Box' region (the same region used in Figure 4.1, 4.6, and 4.7) and the 'Small Loop' region shown in the right panel. The spectral fit value for the microflare excess is shown in the loci plot with a 90% confidence region where errors on all other quantities are omitted due to their unavoidable large magnitude. This figure is taken from Cooper et al. (2020).

what is seen in other X-ray, and EUV, microflares (Athiray et al. 2020). This analysis provides another important reason to have higher energy X-ray spectroscopic observations of the Sun as they are able to help constrain material at hotter temperatures (Reale et al. 2009; Schmelz et al. 2009; Ishikawa et al. 2017) since Figure 4.7 shows that it is the NuSTAR loci curves that are significantly constraining the EMD at log(T)>6.7 (>5 MK).

In order to gain more understanding of the microflare plasma present in the previously defined 'Box' region we subtract the quiet AR EMD (Figure 4.7, left panel) from the microflare EMD (Figure 4.7, right panel) to isolate the microflare heated plasma, shown in the left panel of Figure 4.8. The loci curves become more consistent with each other and the excess spectral fit values (6.7 MK and 8.0×10^{43} cm^{-3}), which implies that similar emission is again being seen by NuSTAR and the Fe XVIII proxy.

The Fe XVIII loci curve from the 'Small Loop' area, shown in the right panel (blue, solid), is far more consistent with the NuSTAR loci curves and the spectral fit values than the 'Box' region (blue, dashed-dotted). Therefore, it is likely that the microflare emission came from the 'Small Loop' region with the pre-flare/quiet AR emission was composed over a larger area of the AR.

More evidence for this conclusion can be seen when calculating the expected Fe XVIII flux from a plasma with the microflare spectral fit values (Figure 4.4, right panel) and comparing them to the observed Fe XVIII flux (Section 1.4.4). When using the area of the 'Box' we obtain a synthetic Fe XVIII flux

of $1.20^{+0.11}_{-0.09}$ DN s^{-1} pix^{-1} and an observed flux of 1.95±0.06 DN s^{-1} pix^{-1} for the quiescent AR time; therefore, ~62% of the emission seen by SDO/AIA is also seen by NuSTAR. However, NuSTAR only sees ~30% of the SDO/AIA 'Box' emission during the microflare (synthetic flux: $0.07^{+0.06}_{-0.04}$ DN s^{-1} pix^{-1}, observed: 0.23±0.09 DN s^{-1} pix^{-1}). When calculating the synthetic flux for the microflare time with the 'Small Loop' we find the values to be much more consistent with ~69% agreement with a synthetic flux of $1.14^{+1.03}_{-0.57}$ DN s^{-1} pix^{-1} and observed Fe XVIII flux of 1.66±0.16 DN s^{-1} pix^{-1}. Errors on the observed Fe XVIII proxy emission are propagated from the uncertainties obtained from its constituent native SDO/AIA channels via the SolarSoft IDL routine `aia_bp_estimate_error.pro`[8] and the errors on the synthetic values are propagated from the spectral fit uncertainties.

Therefore, it is concluded that the microflare does indeed take place in the 'Small Loop' region shown in Figure 4.8. Thus, the new loop volume is calculated to be 1.9×10^{25} cm^3 (approximate density of 2×10^9 cm^{-3}, Equation 3.2) providing an energy estimate of $1.1^{+0.2}_{-0.2}\times10^{26}$ erg, lowering an already small thermal energy upper limit. We now have a lower heating requirement, finding this new energy to still be consistent with a non-thermal heating scenario, as shown with a magenta shaded region in Figure 4.6 (right panel).

4.6 Summary and Conclusions

In combination with SDO/AIA, we use NuSTAR to further investigate the weakest microflare from the 2018 September data-set discussed in Chapter 3, finding it to be the weakest thermal X-ray microflare to be detected in an AR. We find typical pre-flare/quiescent AR temperatures and emission measures of approximately 3 MK and 10^{46} cm^{-3}, respectively. The tiny microflare is found to reach temperatures of 6.7 MK with an emission measure of 8.0×10^{43} cm^{-3}, a hotter temperature and lower emission measure than previously investigated NuSTAR microflares (Hannah et al. 2016, 2019; Glesener et al. 2017; Wright et al. 2017).

The microflare discussed in this chapter is not the most spatially compact flare, in fact ~10″ in length; however, with a value of $1.1^{+0.2}_{-0.2}\times10^{26}$ erg, the microflare produces the smallest instantaneous thermal energy release seen in an AR X-ray microflare. Small brightenings observed in high-resolution EUV data of magnetically braided loops achieve comparable energies (Cirtain et al. 2013). This would suggest that NuSTAR is able to detect the X-ray emission from the many observed small EUV brightenings and is starting to approach those scales beyond microflares, towards nanoflare energies.

[8] https://hesperia.gsfc.nasa.gov/ssw/sdo/aia/idl/response/aia_bp_estimate_error.pro

We find no obvious direct evidence of non-thermal emission in the microflare NuSTAR spectrum (Figure 4.4, right panel); however, we find indirect signatures of electron acceleration from other sources. Some potential evidence is provided by the more impulsive nature of the higher X-ray energy range (4–7 keV) time profile compared to the lower energy emission in Fe XVIII and the 2.5–4 keV range. We also find a reasonable region in the non-thermal parameter space which could produce the observed microflare heating while still being consistent with null detection.

Even though the microflare is incredibly weak it is very clear in the NuSTAR X-ray data and more difficult to locate in the SDO/AIA EUV images. This emphasises the need for sensitive X-ray telescopes capable of observing flares on the Sun. This event could be used to help identify others of a similar nature, through careful study of X-ray emission, and further the investigation the flare frequency distribution of small to large flares (Crosby et al. 1993; Hudson 1991; Hannah et al. 2011, see Section 1.3).

Throughout the previous chapter (Chapter 3) and this chapter, we have demonstrated the need for spectral fitting with respect to high energy spectroscopy. In Chapter 5, we elaborate on spectral fitting approaches with particular emphasis on the abilities of available X-ray spectral fitting tools.

5

X-ray Spectral Fitting Approaches

In this chapter, we will cover fitting models to solar X-ray spectra. In Section 5.1 we will recap the specific difficulties when fitting X-ray spectra (compared to say EUV or optical) and in Section 5.2 we will cover appropriate fitting statistics and methods to characterise X-ray spectra. Then in Section 5.3 will summarise the current software in solar and astrophysics for fitting X-ray spectra, highlighting their strengths and weaknesses, before introducing a new solar python package in Section 5.4, Sunxspex. For this package I have developed code for it to work with a variety of solar X-ray data, interface to existing python fitting packages, and successfully fitted a range of solar X-ray spectra. This will be introduced in Section 5.4, with more detailed examples given in Chapter 6.

5.1 Fitting X-ray Spectral Data

Spectral fitting with X-ray data presents specific difficulties when compared to fitting EUV or optical data. There are three main issues which mean X-ray spectral fitting can be trickier than fitting other wavelength ranges. These three issues are that:

1. the models are computed in photon-space and need to be converted to count-space, which is determined by the instrument specific spectrum response matrix (SRM), in order to be compared to the observation;

2. a detected count with a given energy can be measured due to a photon at a higher energy which is described by the non-diagonal entries in that SRM; and

3. the emission greater than a few keV is dominated by bremsstrahlung continuum emission from either thermal or non-thermal electron distributions, with the resulting X-ray spectrum falling off as a power-law with increasing photon energy; therefore, we need to derive physical

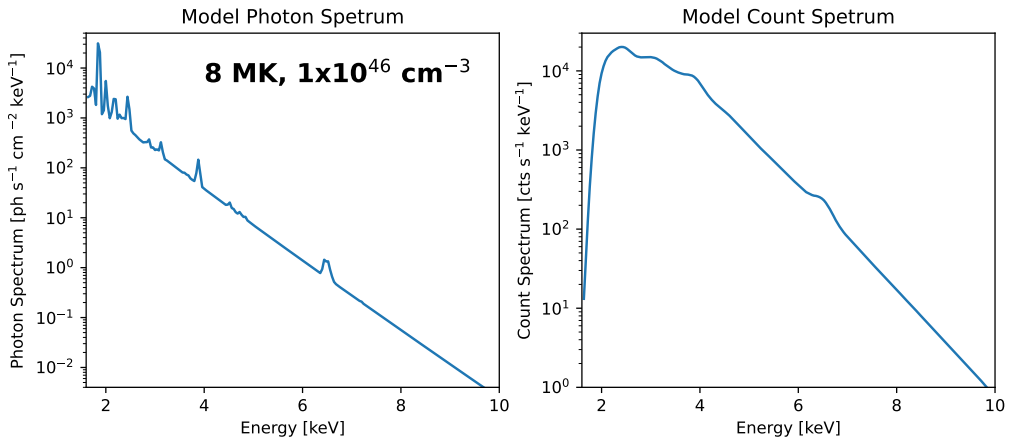

Figure 5.1: Thermal photon model, produced using the Sunxspex Python package (Section 5.4), for a plasma at 8 MK with an emission measure of 10^{46} cm^{-3} (left panel) and the corresponding NuSTAR count model spectrum (right panel). Clear spectral lines visible in the photon model spectrum (left panel) are smoothed out by NuSTAR's instrument non-diagonal SRM (right panel) and the steep count spectrum component below ~2 keV is caused by the instrumental response and effective area and helps avoid throughput being dominated by low energy counts.

parameters from an X-ray spectrum that has both very bright and orders of magnitude fainter (and very noisy) components.

Figure 5.1 shows an example of a photon model (left panel), a thermal model with a temperature and emission measure of 8 MK and 10^{46} cm^{-3}, respectively. The right panel in Figure 5.1 then shows the resulting count model from the photon model after it is convolved with NuSTAR's SRM. We aim to obtain the physical parameters from a given X-ray spectrum by forward-fitting photon models to the X-ray data. However, given the bright and noisy components, and the possible inclusion of a spectrum representative of a background, careful consideration of fit statistics is needed and will be discussed in Section 5.2.

5.2 Fitting Methods and Statistical Approaches

There are many fitting methods and statistical approaches that can be taken when fitting models to data. This section will cover those specific to X-ray spectral analysis with emphasis on low count regimes.

The fitting approach discussed and used in all chapters is a forward-fitting approach. This involves producing a photon spectrum, by modelling an electron distribution and convolving with

the bremsstrahlung cross-section (Section 1.4.5 and 1.4.4), and folding this through the SRM (see Section 2.1.5 for NuSTAR's SRM) as shown in Figure 5.1. This is done using

$$C = I \# SRM \quad [\text{counts s}^{-1}\ \text{keV}^{1}], \tag{5.1}$$

where C is the modelled count spectrum array, I represents the modelled photon spectrum array (photons s^{-1} cm^{-2} keV^{-1}), and SRM is the spectral response matrix of the instrument in units of counts $photon^{-1}$ cm^2 (see Figure 5.1) with # representing matrix multiplication. The modelled count spectrum can then be compared to the observed count spectrum via a fit statistic. The parameters that result in the minimum statistic value are taken to best represent the emitting observed plasma under the assumptions of the initially defined models.

A different approach could be taken by attempting to invert the observed count data through the SRM into photon-space then comparing this directly with the photon models. One advantage of this approach is that it would attempt to recover the original photon spectrum which may reveal spectral features that would otherwise be undetectable when forward-fitting with already defined parametric models (Piana 1994; Piana & Brown 1998; Kontar et al. 2004). However, this is an ill-posed inverse problem (Bertero et al. 1985) which is difficult and expensive to solve. Brown et al. (2006) discusses a similar example when recovering the underlying electron distribution from a simulated photon spectrum.

Within the forward-fitting approach there are numerous aspects to consider. There are a limitless number of fit statistics, and many different optimisation methods to utilise each fit statistic, to comprehend when exploring the model parameter-space. In addition, Bayesian techniques can be applied to provide robust understandings of the posterior distributions of parametric fits to data giving values and error margins that are more representative and meaningful of the fitting process. While fitting, we can also consider multiple data sources or spectra simultaneously when considering the fit statistic carefully (see Section 5.2.5).

5.2.1 Likelihoods and Fit Statistics

Through any fitting process, a fit statistic is required in order to evaluate how well the data is fitted by the model in question. In the case of a likelihood this is better expressed as a measure of the probability that the data was produced from such a model given a source distribution (Sivia & Skilling 2006, pp. 15). Therefore, a fit statistic is a 'goodness-of-fit number', whereas a likelihood is tied to the fundamental nature of the data and gives the probability that the data was produced by the scenario described by the model (see Section 3.5 of Sivia & Skilling 2006).

In practice, the fit statistic is calculated across every dimension of the data being fitted and then combined in a multiplicative fashion. To simplify this process in the case of a likelihood, we often take the logarithmic form of the expression converting the multiplication into a summation and any exponent into a factor. The logarithmic form also avoids increasingly small numbers being computed since summation continuously increases the final absolute value, whereas multiplication can lead to vanishingly small numbers.

We describe fit metrics associated with two statistical regimes crucial to astronomical data analysis, Gaussian and Poissonian.

Gaussian Likelihood

The Gaussian log-likelihood takes the form

$$ln(\mathcal{L}_{Gaussian}) = -\frac{1}{2}\sum_{i=1}^{N}\left(ln(2\pi\sigma_i^2) + \left(\frac{D_i - \mu_i}{\sigma_i}\right)^2\right), \tag{5.2}$$

where $\mathcal{L}_{Gaussian}$ is the likelihood, N is the number of data-points, D_i is the observed data in counts, and σ_i represents the errors for each data-point i (see Section 2.3 of Sivia & Skilling 2006). In addition, μ_i represents the modelled data (also in counts) which is equivalent to the i^{th} entry in $C \cdot \delta t \cdot \delta E$ where δt and δE are the time and energy binning, with respect to Equation 5.1.

The normal distribution is described by two values, the peak value and standard deviation which is a measure of its width. The peak value is also equal to the mean, median, and mode of the distribution. Therefore, if we assume a Gaussian source distribution of the data D with standard deviation σ then Equation 5.2 provides the probability that the data came from the source model μ with a maximum value resulting in the best-fit model.

We note that Equation 5.2 has a maximum of $-\frac{1}{2}ln(2\pi\sigma_i^2)$. We find that this term is a constant as it only relies on the data. Therefore, it is trivial to produce a fit statistic from the Gaussian log-likelihood, only including those terms relevant to the varying model fit to the data.

Chi-squared Fit Statistic

The chi-squared fit statistic, χ^2 (see Section 3.5 of Sivia & Skilling 2006), is ubiquitous throughout many areas of science and statistics due to its intuitive nature and importance to the widely applicable

Gaussian distribution. The statistic is

$$\chi^2 = -\sum_{i=1}^{N} \left(\frac{D_i - \mu_i}{\sigma_i} \right)^2, \tag{5.3}$$

which is the last term in Equation 5.2. This fit statistic removes the expression that does not rely on μ in Equation 5.2, only comparing the squared, uncertainty weighted differences between the data and the calculated model.

Since all constants are removed, we find that the maximum value of Equation 5.3 is 0 which would result from a perfect fit of the model to every data-point. Forms of the χ^2 test are also routinely used even when the data does not pertain to the Gaussian regime due to its easily interpreted meaning. For example, the residuals shown from spectral fitting in Chapters 3, 4, and 8 utilises the statistic in their respective residual plots even though much of the data is well within the Poissonian regime.

Poissonian Likelihood

Gaussian statistics (Equation 5.2 and 5.3) are useful when dealing with a large number of independent and random events, such as detected counts from a bright source (see Section 2.4.1 and 3.1 of Sivia & Skilling 2006). However, within a solar X-ray spectrum, we are likely to have a low-count component within our spectra (Section 5.1); therefore, a different approach should be taken.

Poisson statistics should be used when events are measured in a finite time range where the event rate is constant and each event is independent of the others with no restriction on the number of events present. This distribution is a special case of the binomial distribution which describes the probability of n sequential independent events (see Section 5.1.1 of Sivia & Skilling 2006). Therefore, we initially start with a description of the binomial regime.

First, we express the probability of obtaining a certain outcome; i.e., a certain combination of 'events' and 'not events' or say obtaining k events from n measurements. Therefore, the probability of this specific outcome happening once is

$$P(\text{outcome}) = P(\text{event})^k \times P(\text{no event})^{n-k} = P(\text{event})^k \times (1 - P(\text{event}))^{n-k}, \tag{5.4}$$

where P represent the probability of the indicated desire between 0 and 1. However, when performing n measurements, this specific outcome could occur multiple ways (see Chapter 4 of Kurt 2019).

This is calculated by establishing that n objects can be arranged $n!$ ways; however, if we only pick k objects from n then we have $n \times (n-1) \times ... \times (n-k+1)$ ways. This is then divided by the number

of ways this list of k objects can be arrange, which is $k!$. The expression for this is the mathematical choose formula,

$$C(n,k) = \frac{n!}{(n-k)!} \times \frac{1}{k!}, \tag{5.5}$$

where n and k have the same meanings as Equation 5.4. Therefore, to obtain the probability of getting a specific desired outcome of events, we multiply Equation 5.4 with Equation 5.5 producing

$$\mathcal{L}_{Binomial} = \frac{n!}{k!(n-k)!} \times P(\text{event})^k \times (1 - P(\text{event}))^{n-k}, \tag{5.6}$$

giving an expression for the binomial likelihood.

In many astrophysical cases, we are in a regime where a large number of measurements, n, can be made and we may or may not detect an event. We may not know the exact number of measurements; however, if we assume the event rate is, on average, constant we may write

$$\lambda = n \times P(\text{event}), \tag{5.7}$$

where λ is the success rate over n measurements. Therefore, if we rearrange for $P(\text{event})$, substitute into Equation 5.6, and take the limit as $n \to \infty$ we arrive at

$$\mathcal{L}_{Poissonian} = \frac{\lambda^k e^{-\lambda}}{k!}, \tag{5.8}$$

which is the expression for the Poissonian likelihood (see Section 5.3.4 of Sivia & Skilling 2006). Expressing this in logarithm form and changing nomenclature to be consistent with Equation 5.2 we obtain

$$ln(\mathcal{L}_{Poissonian}) = \sum_{i=1}^{N} \left(D_i ln(\mu_i) - \mu_i - ln(D_i!)\right). \tag{5.9}$$

Note that no errors are included in the Poissonian log-likelihood. This is because the source distribution would be assumed to be Poissonian and, therefore, its variance would be completely determined by its value; e.g., a value of N has variance N and standard deviation $\sqrt{N}$. This is different to other distributions where the value and variance can be independent.

Although Equation 5.3 is commonly used in astronomical data analysis, Humphrey et al. (2009) shows that a poissonian approach should be adopted while fitting. Even when working with a large number of counts, the number of counts in every individual energy bin also has to be sufficiently large in order for Equation 5.3 to be utilised.

When using Equation 5.9 it can be cumbersome to continuously calculate the factorial term, which only relies on the data and does not change when performing model fitting.

Cash Fit Statistic

In order to create a fit statistic from the Poissonian log-likelihood in Equation 5.9, we can remove the data dependent terms that do not change when fitting the same data with varying models. This is similar to the creation of χ^2 (Equation 5.3) from the Gaussian log-likelihood in Equation 5.2. Cash (1979) derives the form

$$ln(\mathcal{L}_{Cash}) = \sum_{i=1}^{N} (D_i ln(\mu_i) - \mu_i)\,, \tag{5.10}$$

where the factorial term is entirely dropped from Equation 5.9. However, we may wish to keep the likelihood nature of Equation 5.9 while avoiding the factorial calculation.

C-stat Fit Statistic

Another approach to removing the factorial term from Equation 5.9 is used in XSPEC, the astronomical X-ray spectral fitting package (Section 5.3.2), and is represented as

$$ln(\mathcal{L}_{C-stat}) = \sum_{i=1}^{N} \left(D_i \left(ln\left(\frac{\mu_i}{D_i}\right) + 1 \right) - \mu_i \right), \tag{5.11}$$

which is obtained by altering the $ln(D_i!)$ term in Equation 5.9. We replace this factorial with

$$ln(x!) \approx x ln(x) - x, \tag{5.12}$$

termed Stirling's approximation (Sivia & Skilling 2006, pp. 114). Equation 5.11 retains the information of the original Poisson likelihood (Equation 5.9). Therefore, Equation 5.11 will approach the same behaviour of the Gaussian and χ^2 statistic for data with large count numbers in each data bin when optimised over the model parameter-space (Humphrey et al. 2009). Note that, since Stirling's approximation does not hold for a value of zero ($x = 0$), the original Poisson likelihood is used (Equation 5.9) in place of Equation 5.11 when zero count scenarios ($D_i = 0$) are encountered.

5.2.2 Optimisation Methods

In combination with a given fit statistic, we have access to a plethora of optimisation techniques. These methods control the parameter-space exploration, attempting to locate the maximum or minimum fit statistic which corresponds to the best fitting model to the data. Additional checks may be required to increase confidence that the optimised statistic represents the global, and not the local, best-fit value. The statistics presented in Section 5.2.1 all have a maximum value when the best fitting model

is obtained. For some processes we are required to adapt the fit statistics to reach a minimum value for the best-fit; therefore, each are multiplied by a factor of -2 for this to be achieved. This factor is used to remove the $1/2$ constant in Equation 5.2 and is only used for minimising the statistic, *not* for MCMC or nested sampling analysis which are discussed in Section 5.2.3 and Chapter 7, respectively.

The Nelder-Mead Method

Through personal experimentation, one optimisation method that works well when hunting for the best-fit in X-ray spectral fitting is the 'Nelder-Mead' method (Gao & Han 2012). The Nelder-Mead method was used for the fitting displayed in Figure 8.11, 8.12, and 8.13. This method utilises a simplex[1] in the n-dimensional parameter-space which consists of $n + 1$ vertices. The fit statistic at each vertex is the calculated and the 'centre-of-mass' coordinate is calculated. The worst valued vertex is then reflected through the centre-of-mass. The simplex undergoes one of three operations depending on the value of the reflected point in comparison to all other points.

If the reflected point has a statistic value between the best and second worst value of all original vertices then it replaces the worst vertex. If the reflected point has the best statistic thus far then the simplex is expanded and either the reflected or new expanded point replaces the worst original vertex. However, if the reflected value is less desirable than the original vertex values (excluding the worst one) then the simplex is contracted (Gao & Han 2012).

The contraction position depends on whether the original worst vertex is better or worse than the reflected point statistic. If the former is the case then the new contraction point is from the original worst vertex value towards the centre-of-mass, whereas if it is the former then the contraction point is created between the reflected point and the centre-of-mass. The contraction point is accepted if it has a more desirable fit statistic than the coordinate from which it contracts away. If the contraction fails (i.e., the contraction point is worse than any value produced) then a shrinking process occurs where all vertices are replaced using a function dependent on the best vertex (Gao & Han 2012). A termination criteria is made based on the standard deviation of the calculated fit statistic from the vertices.

The Nelder-Mead method is a direct search method and so does not need any gradient information to be utilised. This also means it can be used on non-continuous functions and those that are not differentiable. This nature allows the method to be easily applied when fitting X-ray spectral data with highly complex models, and is made available in Python via Scipy's `minimize` method[2]. However,

[1] The simplest geometrical shape able to be created in an n-dimensional space. For example, a point in 0-dimensions, a line segment in 1-dimension, a triangle in 2-dimensions, etc.

[2] `https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.minimize.html`

we note that many other fitting methods are available that may be more efficient depending on the specific problem at hand.

Optimisation tools are crucial when performing spectral fitting; however, finding the best-fit value may not always be what is required from the analysis. We describe Markov chain Monte Carlo (MCMC) methods, which helps explore the parameter-space and map its behaviour, in Section 5.2.3.

5.2.3 Markov Chain Monte Carlo Methods

Markov chain Monte Carlo (MCMC) methods make use of techniques in order to sample the posterior distribution, producing a probability density. This is generally accomplished by creating a chain of samples where each probability depends on the proceeding value; this chain is then built up through a stochastic, Monte Carlo process which randomly selects the next parameter-space step with some acceptance criteria. Therefore, this creates a chain of samples. There are many Monte Carlo methods to produce chains, one of which is the relatively simple Metropolis-Hastings algorithm (Hastings 1970). This method's fundamental and intuitive approach is used in some way throughout many other MCMC approaches. MCMC methods are made available in many Python packages, one of which being Emcee (Foreman-Mackey et al. 2013).

As mentioned, the MCMC process samples the posterior distribution which is the probability of a model given the data. Here we wish to investigate the posterior of the parameters θ for a given data-set D and model M using

$$P(\theta|D,M) = \frac{P(D|\theta,M)P(\theta|M)}{P(D|M)}, \tag{5.13}$$

where $P(D|\theta,M)$ is the likelihood, $P(\theta|M)$ is the parameter priors, and $P(D|M)$ is the evidence. When using Equation 5.13 in relation to the MCMC process we generally neglect the evidence term on on the denominator and just consider the numerator. This is because the evidence in Equation 5.13 does not change when fitting the same model and only acts as a normalising constant; however, the evidence $P(D|M)$ is very significant when comparing the fits from two different models and can be investigated via *nested sampling* (see Chapter 7).

The parameter prior can take many forms; however, we focus on using a uniform distribution between sensible boundaries giving a probability of $P = 1$ ($ln(P) = 0$) at or inside the boundaries and $P = 0$ ($ln(P) = -\inf$) elsewhere. For example, an incredibly conservative temperature range would be between 5–15 MK for microflaring plasma. The likelihood takes an appropriate form relative to the data being analysed as described in Section 5.2.1.

Therefore, once the likelihood expression is chosen (perhaps one from Section 5.2.1), the MCMC process can be applied to the data. Initially, points are chosen throughout the target probability density distribution, the posterior, at N θ_{step} coordinates. These points are termed 'walkers'. A 'jumping distribution' is also defined (e.g., a Gaussian centred on the current walker position) which helps determine the new trial walker position, θ', from the previous location to be assigned as the next step θ_{step+1} (see Section 9.4.4 of Sivia & Skilling 2006).

The ratio of posteriors is taken between the new proposed point and the old point, θ'/θ_{step}. A random number, α, is then generated between 0–1. If the ratio is greater than or equal to α then the new walker step is accepted ($\theta_{step+1} = \theta'$); however, if the ratio is less than α then the new point is rejected ($\theta_{step+1} = \theta_{step}$). This process continues for a set number of steps for each defined walker.

The results of an MCMC can be displayed in a corner plot, as shown in Figure 3.6 and 3.7. These plots show the projection of all the accepted MCMC walker samples throughout the process, displaying all correlations in density distribution between each parameter, with the mode of this distribution showing the most probable model parameters (the *maximum a posteriori*). The diagonal panels show the integrated walker density for each model parameter.

For example, MCMC analysis can be performed on fitting a single thermal model to a given count spectrum, similar to the pre-flare spectrum in Figure 4.4. From Equation 5.13, we have

$$P(T, EM|C, M) \propto P(C|T, EM, M)P(T, EM|M), \tag{5.14}$$

where T and EM are the temperature and emission measure for the thermal model M with C being the count spectrum. Therefore, we sample the posterior with different values of T and EM providing us with a robust measure of how well constrained these parameters are with reliable error estimates.

5.2.4 Error Estimation

When the model fitting is performed, we need to gain an understanding of confidence in the fitted parameters. There are multiple methods to attain uncertainty ranges for the fitted parameters. We can obtain crude 1-σ errors if we assume normally distributed data measurement errors and that the parameters are linear over the uncertainty range. In this case, the parameter distribution is given by

$$P(\theta) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{\theta - \mathrm{m}}{\sigma}\right)^2}, \tag{5.15}$$

with θ, m, and σ being the parameter, mean value, and standard deviation of the Gaussian distribution. From this we can define an objective function, $\mathcal{J}(\theta)$, by using the minimisation form[3] as discussed in Section 5.2.2. Therefore, we find

$$\mathcal{J}(\theta) = -2ln(P(\theta)) = ln(2\pi\sigma^2) + \left(\frac{\theta - \mu}{\sigma}\right)^2. \tag{5.16}$$

We can then calculate the local curvature around the mean value μ (i.e., the fitted value) via the Hessian matrix (Press 1996, pp. 676–678). The Hessian matrix of $\mathcal{J}(\theta)$ is then

$$\mathcal{H}(\mathcal{J}(\theta)) = \frac{\partial^2 \mathcal{J}(\theta)}{\partial \theta^2} = 2\sigma^{-2}, \tag{5.17}$$

where σ is the covariance matrix. Therefore, the covariance matrix can be found from inverting $\mathcal{H}$ and multiplying by 2. The Hessian matrix is calculated by

$$\mathcal{H}(\mathcal{J}(\theta)) = \begin{bmatrix} \frac{\partial^2 \mathcal{J}}{\partial \theta_0^2} & \cdots & \frac{\partial^2 \mathcal{J}}{\partial \theta_0 \partial \theta_j} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 \mathcal{J}}{\partial \theta_i \partial \theta_0} & \cdots & \frac{\partial^2 \mathcal{J}}{\partial \theta_{i=j}^2} \end{bmatrix}, \tag{5.18}$$

where in practice the differentiation is performed from first principle with an appropriate $\delta\theta$. A disadvantage of this method, beyond the simplifying assumptions, is that we only gain estimates of the 68% (1-σ) uncertainty range.

However, there are more robust ways of obtaining a more complete picture of the distribution uncertainty ranges. This is essentially provided for free when performing MCMC analysis (Section 5.2.3) where any uncertainty range can be found through the probability density shown by the accepted walker positions. For example, the 68% (1-σ equivalent) uncertainty ranges are shown with contours for each parameter in Figure 3.6 and 3.7. An additional benefit to the MCMC process is that no assumptions have to be made about the underlying probability density hyper-surface shape.

Ireland et al. (2013) shows both of these approaches applied to X-ray spectral observations as well as two other methods, χ^2-mapping and the 'Monte Carlo' method. The former involves minimising the χ^2 (Equation 5.3) over all parameters to get χ^2_{min}, then one parameter is stepped through a range of values while minimising over the rest to get a range of χ^2_{steps} with $\delta\chi^2 = \chi^2_{steps} - \chi^2_{min}$ giving an estimate of the confidence range[4]. The Monte Carlo method involves generating a new 'observed' spectrum such that the new counts in each energy bin are drawn from a Poisson distribution with a mean value of the true observed spectrum from the same bin, the new spectrum is then fitted. This

[3]The log-likelihood multiplied by −2.
[4]A $\delta\chi^2 < 1$ occurring ~68% of the time, $\delta\chi^2 < 4$ occurring ~95% of the time, etc.

process is then repeated for a large number of iterations giving a distribution of parameter values centered on the parameters that produced χ^2_{min}. An uncertainty range on each parameter can then be estimated from the resulting distributions.

They found that all methods, compared to the more complete density distributions provided by the Bayesian MCMC approach, appeared to underestimate the true error range and were less informative (Ireland et al. 2013). Therefore, the ability to be able to easily perform MCMC analysis on any fitting process is essential, especially when the parametric model becomes increasingly complex while fitting multiple data-sets.

5.2.5 A Note On Simultaneous Fitting

Once a model is chosen, along with a relevant fit statistic and fitting method, it can be straightforward to forward-fit a single given spectrum. However, we may have multiple spectra of the same emitting source to fit. This is very common with NuSTAR's two FPMs, but we may also have spectra from completely different observatories. To incorporate all the available spectral information into the same optimisation process we can simultaneously fit all available spectra. This can be useful as each spectrum may be noisy or difficult to fit individually and, therefore, each additional spectrum helps constrain any fitting process.

Previous RHESSI spectral analysis show another way to incorporate all data into the fitting process, since RHESSI (Section 2.3) has nine detectors with six usually being optimal for spectroscopy. The fitting performed in OSPEX (Section 5.3.1) has required the six spectra to be summed and the responses averaged; therefore, performing the fit to a single spectrum. However, this ignores systematic differences between the detectors and so simultaneous fitting is preferable (see later in Section 6.2.2).

When including additional spectra, we can think of this as adding in an extra dimension to calculate the fit statistic over. Therefore, just like summing the fit statistic over the energy bins of a given spectrum, we are also able to sum over all spectra.

A simple simultaneous fitting example is given in the case of microflare 3's pre-flare spectrum in Section 3.4.1. Using XSPEC (Section 5.3.2), we fitted both NuSTAR's FPMA and B with a single thermal model with parameters T and EM. We then calculate the fit statistic for both spectra from a single thermal model created using a single temperature and emission measure. A scalar parameter (C) is used for FPMB's spectrum to account for any systematic offset between NuSTAR's instruments. Both of these fit statistics are then summed; it is this final statistic that is then minimised by varying the C, T, and EM parameters.

Simultaneous fitting is a powerful method as it allows all available data to constrain any model or model parameter and is also useful for determining systematic differences between instruments. For example, the total model can be kept the same between all spectra and optimised or, under a more likely scenario, certain parameters can be kept the same (e.g., T and EM) with other parameters varying specifically to their spectrum (e.g., the C scalar for FPMB's spectrum).

5.2.6 Event Background Handling

When performing analysis on spectral data, it is often beneficial to define a background spectrum which represents the surrounding or underlying emission to the event, or counts from sources completely independent of the solar source, i.e. instrument background. There are multiple ways in which this background can be incorporated into the fitting to either represent the background directly or help constrain the background model components.

Method 1: Fixed Background Model

One method (method 1) has been used when including any pre-flare component in Chapter 3, 4, and 8, where the pre-flare is used to represent the surrounding non-flaring AR. We fit the defined background (or pre-flare) using parametric models first, this provides a model that can then be kept constant throughout any successive event spectra. This subtly changes the fit statistic as Equation 5.8 takes the form

$$\mathcal{L}_{Poissonian,\,\mathrm{M1}} = \prod_{i=1}^{N} \frac{(\mu_i + b_i)^{D_i} e^{-(\mu_i + b_i)}}{D_i!}, \tag{5.19}$$

where b is the counts from the fixed background model. If a pre-flare time is not available to determine a fixed background model, a freely varying model can be introduced to fit the low-energy emission instead. This can be taken as fitting a 'background' component which can be corroborated if the fitted parameters are similar to those found by the surrounding non-flaring emission at other times. For example, this is the case for microflare 5 and 6 in Chapter 3.

The analysis for a large number of NuSTAR microflares has used this method in order to preserve statistical information due to low-count observations (Wright et al. 2017; Glesener et al. 2017; Hannah et al. 2019; Duncan et al. 2021).

Method 2: Simultaneously Fitted Background

A second approach we can take (method 2), if the background can be fit well with a parametric model, is to simultaneously fit the background model component in both the background and event spectrum. Similar to method 1, this changes the form of the log-likelihood

$$\mathcal{L}_{Poissonian,\,\mathrm{M2}} = \prod_{i=1}^{N} \frac{(\mu_i + Cb_i)^{D_i} e^{-(\mu_i + Cb_i)}}{D_i!} \frac{b_i^{B_i} e^{-b_i}}{B_i!}, \tag{5.20}$$

where C is a ratio scaling factor to account for the background spectra potentially being sampled over different time ranges or spatial regions and B represents the observed background counts. For example, if the background and data counts were obtained over different time ranges (t_B and t_D, respectively) then C would become $\frac{t_D}{t_B}$.

This has the advantage that the background model fit will be influenced by all the relevant data it is attempting to represent; however, this method is a more computationally expensive process and may not provide drastically different results to the previous method.

Method 3: Background Data Plus Model

Both method 1 and 2 require the background spectrum to be well represented by parametric models; however, this is not always appropriate. Therefore, we discuss another method that utilise the background spectrum data more directly instead of a fitted representation. Method 3 takes the background spectrum data and combines it with the event model, therefore, fitting the full observed event spectrum with an event model plus the background data. Thus, again using the Poisson case as an example, the log-likelihood changes to

$$\mathcal{L}_{Poissonian,\,\mathrm{M3}} = \prod_{i=1}^{N} \frac{(\mu_i + CB_i)^{D_i} e^{-(\mu_i + CB_i)}}{D_i!}, \tag{5.21}$$

which is similar to Equation 5.19 but with the model background b being replaced with the observed background counts B.

Method 4: Non-parametric Background Model

It is possible to obtain a maximum likelihood model for the background even if we cannot use a parametric model and is an available method in XSPEC (Section 5.3.2). In the Poissonian regime, we

can utilise method 2's expression for this purpose. We take the logarithmic form of Equation 5.20 and then locate the log-likelihood maximum with respect to the background model b via

$$\frac{\partial ln(\mathcal{L}_{Poissonian,\,\mathrm{M2}})}{\partial b} = 0. \tag{5.22}$$

This can then be solved analytically to give a background count model of

$$b_i = \frac{D_i + B_i - \left(1+\frac{1}{C}\right)\mu_i + \sqrt{\left(\left(1+\frac{1}{C}\right)\mu_i - D_i - B_i\right)^2 + 4\left(1+\frac{1}{C}\right)B_i\mu_i}}{2(C+1)} \quad \text{[counts]}, \tag{5.23}$$

where C is the ratio scaling factor describing differences in which the data and background are chosen as described for method 2.

Method 5: Background Subtraction

The final method to be discussed is arguably the simplest and eliminates the need to account for the background during the fitting process. Method 5 denotes the process of just subtracting the background counts from the event counts. The resulting 'background removed' spectrum's likelihood is

$$\mathcal{L}_{Poissonian,\,\mathrm{M5}} = \prod_{i=1}^{N} \frac{\mu_i^{(D_i - CB_i)} e^{-\mu_i}}{(D_i - CB_i)!} = \prod_{i=1}^{N} \frac{\mu_i^{S_i} e^{-\mu_i}}{S_i!}, \tag{5.24}$$

where S represents the counts solely detected from the source of activity being investigated. Therefore, the new spectrum, S, is fitted with models detailing the event. The only surviving background information in this process with regards to the fitting is included in the error for the new spectrum. The new error is given by

$$\sigma_S = \sqrt{D + BC^2} = \sqrt{S + BC + BC^2}, \tag{5.25}$$

where C keeps the same meaning as in Equation 5.20.

This method can only be valid while the data is in the Gaussian regime. This is not appropriate for a small number of counts (e.g., when in the Poissonian regime) as mathematical properties are not preserved through subtraction. This method is used in many situations and is the default method used in fitting software like OSPEX when including a background spectrum, discussed in Section 5.3.1.

5.3 Available Spectral Fitting Software

As has been shown from Section 5.1 and 5.2, X-ray spectral fitting can quickly become complex. Specialised spectral software needs to exist to specifically handle solar X-ray spectroscopic data and its intricate nature; from the fitting taking place in count-space, meaning the models need to be folded through a non-diagonal instrument response, to the data being simultaneously bright and noisy.

A number of fitting programs exist and are commonly used in scientific data analysis. Two widely used programs for solar X-ray spectroscopic analysis are OSPEX and XSPEC and will be detailed in Section 5.3.1 and 5.3.2, respectively. A sample of other fitting tools, one general and two specific to X-ray spectroscopic analysis, are then discussed in Section 5.3.3. We focus on the presence and absence of features of each software, identifying those that would be convenient to be included in a new spectral fitting software.

5.3.1 OSPEX

The Object Spectral Executive (OSPEX), the successor to the Spectral Executive (SPEX), is a commonly used package for fitting solar X-ray spectra, particularly with RHESSI data (Schwartz et al. 2002)[5]. The OSPEX fitting software is distributed through SolarSoft (SSWIDL) and, therefore, requires an IDL license; however, it is optimised for solar data as it supports a wide range of solar instruments.

OSPEX has been utilised in almost every published RHESSI X-ray spectrum, in statistical microflare studies (Hannah et al. 2008), in multi-microflare investigations (Saqri et al. 2022), and in single flare scenarios (Caspi & Lin 2010). It has also been used to help investigate phenomena related to flaring activity (Kuhar et al. 2017; Mulay et al. 2018). Therefore, the OSPEX package is well established in the solar X-ray community and can be used both through a well-designed and intuitive graphical user interface or the command line.

Many crucial solar models are supplied through OSPEX such as several thermal and non-thermal photon models based on the physics discussed in Section 1.4.4 and 1.4.5. However there remains some significant limitations when performing spectral fitting with certain data products or observations.

Few Fitting Statistics

OSPEX makes a Gaussian assumption on the data as the fitting involves minimising the χ^2 fit statistic (Section 5.2.1). As OSPEX was initially designed to work with RHESSI data this assumption was

[5] https://hesperia.gsfc.nasa.gov/ssw/packages/spex/doc/ospex_explanation.htm

reasonable given the high count rates over much of the X-ray spectra. Therefore, this Gaussian assumption is unlikely to cause many issues when fitting data with a large number of counts, if certain conditions are met (Humphrey et al. 2009); however, this is certainly not the recommended approach when fitting spectra that have energy bins with very few counts. In this case, many of the data bins will be in the Poisson regime meaning that the errors in the measurements with fewer counts will be overestimated and potentially neglected during the fitting process.

It is important to treat these low count regimes rigorously as much of the higher energy channels tend to have the fewest counts. These energies hold the most information on the hottest plasma or can hold signatures of any non-thermal emission that may occur during the event.

The default error estimates for OSPEX also takes the Hessian approach of Equation 5.17. Again, this utilises Gaussian assumptions of the data and may not be appropriate for many use cases. However, there is the ability for more rigorous error estimates through MCMC analysis within OSPEX but, again, this relies on the χ^2 statistic. In addition, OSPEX also limits its parameter-space search to around the best fitted values, assuming them to be correct; therefore, it is unlikely to locate other potential solutions beyond the current local χ^2_{min}.

Single Spectrum Fitting

OSPEX also lacks the ability to fit multiple spectra simultaneously. Therefore, in the case for RHESSI spectral analysis, spectral fitting was performed by summing the spectra and averaging the responses over RHESSI's detectors in order to perform one fit to a spectroscopic observation. As discussed in Section 5.2.5, this feature is not only desired so that we might include all available data of a phenomenon and maintain all the statistical information, but also crucial when we have noisy spectra with very weak signal. Being able to simultaneously fit several spectra is an incredibly important feature, especially under the conditions described, and so any spectral fitting package able to accomplish this has a huge advantage.

5.3.2 XSPEC

XSPEC is an X-ray spectral fitting program (Arnaud 1996) and has been in use since the 1980s[6]. It is a very powerful tool capable of simultaneously fitting multiple spectra from a variety of astrophysical sources with a variety of fit statistics and likelihoods. XSPEC is available across many MacOS and

[6] https://heasarc.gsfc.nasa.gov/xanadu/xspec/manual/node9.html

Linux versions and comes with the HEASoft software[7] which is a unified release of the FTOOLS and XANADU software packages (Blackburn 1995).

XSPEC is primarily designed for astrophysical data analysis; however, can be used in a solar data context (Wright et al. 2017; Duncan et al. 2021). XSPEC has even been used alongside OSPEX for microflare analysis in an attempt to utilise the benefits that both softwares possess (Glesener et al. 2020; Mondal et al. 2021), for example by fitting a spectrum with XSPEC's Poissonian statistics and then using the fitted parameters and the solar models available in OSPEX to search for non-thermal upper limits (Cooper et al. 2020).

When performing the fitting process, XSPEC allows a large variety of fit statistics and likelihoods as well as multiple methods for the parameter exploration. The default error estimates made for the parametric fits from XSPEC are slightly more robust than the default approach used in OSPEX as it somewhat explores the parameter-space to obtain uncertainty ranges instead of the Hessian matrix approach which can only give an estimate of the 68% uncertainty range (see Section 5.2.4). Here, XSPEC uses a similar method to the χ^2-mapping described in Ireland et al. (2013). However, XSPEC allows easy implementation of MCMC analysis in order to properly map out confidence ranges. However, although XSPEC boasts quite an impressive suite of analytical tools, we discuss some disadvantages associated with its use.

Not Optimised for Solar Data Products

One disadvantage to XSPEC compared to the solar optimised OSPEX is the limited support for solar-specific data products. That is not to say fitting solar instrument data is not possible, but that it is not trivial to load and fit these data products, all produced in slightly different formats.

Additionally, XSPEC's interface is not ideal for solar spectroscopy as it is optimised for astrophysical data analysis. In solar spectroscopy, there can be numerous X-ray spectra of different time ranges and spatial regions, whereas many astrophysical sources have relatively few, long integration spectra. Therefore, the XSPEC interface can become convoluted very quickly when fitting the numerous solar X-ray spectra.

Substantial HEASoft Install

No license is required to use XSPEC; however, there is a very complex installation process. It requires many manual command line inputs with specific C, C++, Fortran, Perl, and Python compilers; the

[7] https://heasarc.gsfc.nasa.gov/docs/software/heasoft/

relative versions of the compilers can also have an effect on the install. In addition, XSPEC has a Python interface called PyXSPEC; however, this still requires the full installation.

Limited Solar Models

XSPEC does not contain the non-thermal models commonly used in solar X-ray spectral analysis which are available in the OSPEX software; however, it does have equivalent thermal models. One thermal model is termed an 'APEC' (Astrophysical Plasma Emission Code) thermal model (see Chapter 3, 4, and 8). The APEC model utilises the Astrophysical Plasma Emission Database (APED)[8], as opposed to the `CHIANTI` atomic database (Section 1.4.2) that is primarily used in solar physics and OSPEX, which means that the models from both fitting softwares are not expected to produce the exact same results. In addition, recent work by Mondal et al. (2021) has developed the XSPEC isothermal model `chisoth` which does utilise `CHIANTI`.

5.3.3 Other Fitting Programs

In addition to OSPEX and XSPEC, there are a number of other fitting programs available. SPEX (and PySPEX) was created to analyse high resolution X-ray spectral data obtained from cosmic sources (Kaastra et al. 1996)[9]. Being a program mainly written in Fortran, it boasts relatively fast fitting while being able to handle many free parameters. Similar to XSPEC, there also exists a Python interface for spectral fitting.

Another spectral fitting software to be mentioned is Sherpa (Doe et al. 2007)[10]. It was originally developed to work with data from the Chandra X-ray observatory (Weisskopf et al. 2000) with the intention for it to be easily generalised. Sherpa is the first X-ray spectral fitting software discussed in the chapter that is primarily in Python although it can be built to be able to use the models in XSPEC[11]. However, it does require the installation of other software[12]. It broadly emulates the XSPEC program and has many of the same capabilities; however, this means that it is optimised for astrophysical data products and not solar.

The list of spectral fitting programs discussed thus far is not exhaustive but provides a good representation of the tools available. We summarise the capabilities of OSPEX, XSPEC, and Sunxspex

[8] `http://www.atomdb.org/faq.php`
[9] `https://www.sron.nl/astrophysics-spex`
[10] `https://sherpa.readthedocs.io/en/4.14.1/index.html`
[11] `https://sherpa.readthedocs.io/en/4.14.1/install.html#xspec`
[12] `https://cxc.harvard.edu/ciao/`

Table 5.1: Comparison of the OSPEX, XSPEC, and Sunxspex X-ray spectral fitting programs described throughout Section 5.3 and 5.4.

Spectral Fitting Software	OSPEX	XSPEC	Sunxspex
Optimised for solar data products	Yes	No	Yes
Installation	IDL license required	Complex HEASoft install	Trivial Python package install
Solar specific models	Yes	Limited	Yes
Simultaneous fitting	No	Yes	Yes
Fitting Statistics	Assumes Gaussian Statistics	Multiple	Multiple
Fitting methods	One fixed method	Multiple	Multiple
Sampling methods	MCMC	MCMC	MCMC, nested sampling, any other tool available in Python

(described in Section 5.4) in Table 5.1. However, it is clear that not one fitting package incorporates all the required features necessary for modern solar X-ray spectroscopy.

5.4 Sunxspex, A Python Spectral Fitting Package

We present the Sunxspex spectral fitting Python package[13]. This new X-ray spectral fitting tool is made to incorporate all the necessary and convenient features for solar spectroscopic data products while the pure Python composition allows for a simple installation process and close integration with the other, very well maintained and active, solar python package sunpy. Therefore, we aim for Sunxspex to have all those very necessary tools that are present in OSPEX (e.g., solar models) and also XSPEC (e.g., numerous statistical fitting approaches and simultaneous fitting).

Prior to my involvement in the Sunxspex project, non-thermal models were already available for use. I have contributed to the completion of the thermal model that is now present. In addition to this, I have also created the entire fitting infrastructure. I have since been able to fit a range of solar data which will be discussed in more detail in Chapter 6.

[13] https://github.com/sunpy/sunxspex

5.4.1 Data Support and Interface

All that is required for X-ray spectral fitting is the observed counts spectrum at the measured count bin energies and an SRM with defined photon and count bin energies. All other information is specific to the observation, like the observation effective exposure, effective area, observation livetime, and distance to the source. Therefore, it is crucial to accurately read and load the appropriate observed counts and corresponding SRM from an instrument correctly for any spectral fitting session.

Instrument support is an obvious necessity when spectral fitting and so I have ensured Sunxspex easily handles data products from RHESSI and NuSTAR with improving support for STIX. As well as instrument specific support for solar observatories, I have made it trivial to incorporate custom spectral data into the fitting process. This means that there is little difficulty in analysing available data even if it is from an instrument that is not supported or, if the user wishes, to investigate processes using completely synthetic spectra.

The pure Python, object orientated nature of Sunxspex allows the spectral fitting sessions to be dynamic and interactive; this allows any fitting process to be incredibly flexible and customisable to the user's needs. Therefore, unlike programs like XSPEC, there is no need for multiple similar batch scripts located in several directories dependent on the data location. The object orientated construction also allows many complicated procedures to be performed in a single line of code like fitting the spectrum or sampling the parameter-space.

Additionally, similar to OSPEX, I have made the fitting session able to be saved and re-loaded. This helps avoid large amounts of analysis having to be re-run as the work can just be continued by loading the saved environment.

Sunxspex is easily altered and improved for the community's needs, from just allowing the user more control over their individual fitting environment or making permanent global changes to the software's fundamental workings.

Models and Model Creation

Sunxspex is designed to be solar optimised and, therefore, includes the common solar X-ray spectral models that are included in the OSPEX software (Section 5.3.1). These are primarily the thermal and non-thermal bremsstrahlung models where the non-thermal, thin-target model (Section 1.4.5) has been used in publication (Chen et al. 2021).

Similar to OSPEX and XSPEC, I have ensured models can be defined by providing a mathematical expression as a string of characters. Additionally, I allow more complicated models to be provided as

Python functions. Therefore, the models being forward-fitted to the loaded spectra can be incredibly simple or exceedingly complex. This diversity again gives the user a large degree of customisation and usability.

The Fitting and Parameter Exploration

Sunxspex can also make use of the extensive mathematical libraries available in Python. This allows many powerful tools to be easily swapped or compared to each other throughout the spectral fitting environment. For example, this allows a number of fit statistics to be defined and optimisation or sampling methods to be used.

The default optimisation method used for the spectral fitting is the Nelder-Mead method described in Section 5.2.2 via Scipy's `minimize` method[14]. In addition, the default Sunxspex sampler is the popular MCMC Python package Emcee (Foreman-Mackey et al. 2013) which uses a method proposed by Goodman & Weare (2010) incorporating part of the Metropolis-Hastings algorithm (Section 5.2.3) and properties from the Nelder-Mead method (Section 5.2.2). Using Sunxspex, I show an example of fitting an X-ray spectrum with an isothermal model using Scipy's `minimize`, and MCMC analysis on the same spectrum, in Section 6.1.1. I then show increasingly more complex examples of spectral fitting with Sunxspex in the rest of Chapter 6. Sunxspex also includes another powerful tool in its arsenal called nested sampling which will be discussed in detail in Chapter 7.

The error estimation in the fitting method is calculated via the Hessian approach described in Section 5.2.4; however, it is recommended to perform MCMC analysis to gain more appropriate error margins for the most probable fits. All methods, except method 4, for including background spectra in the fitting can be implemented in Sunxspex. Method 4 is not implemented yet due to its complex and specific nature. Sunxspex can also fit multiple spectra simultaneously (Section 5.2.5).

5.5 Summary and Conclusions

We describe fundamental properties of X-ray spectral fitting software optimised for solar data products. We have discussed reading in the solar X-ray spectral data, being able to access models describing solar plasma, and fitting and parameter exploration methods. We indicated the importance of supporting a wide range of instruments and variety in fitting approaches.

We discussed several useful fit statistics and their appropriate regimes as well as some common parameter exploration and sampling methods. In addition, a number of error estimations are presented

[14] https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.minimize.html

in order to constrain the degree of confidence on the fitted result. We describe the process of simultaneous spectral fitting and how a background spectrum can be incorporated into the fit.

Current X-ray spectral fitting software is also discussed. We put particular emphasis on two well used programs, OSPEX and XSPEC, that are used regularly for solar X-ray spectral fitting while briefly including a description of other spectral fitting software; we highlight available and unavailable features in each. A summary of this in presented in Table 5.1.

Finally, we present the new Python X-ray spectral fitting software Sunxspex—where I have created the entire fitting infrastructure and helped produce the available isothermal model—created to combine many features included in the other disparate discussed programs. This new tool is also included in Table 5.1. This new Sunxspex fitting software is described further in Chapter 6 where we detail many of its features using examples with real observations from published works.

6

X-ray Spectral Fitting in Python

In Chapter 5 we introduced Sunxspex, a new X-ray spectral fitting software that is written in Python and optimised with solar X-ray spectroscopic data in mind. We indicated its capabilities in comparison to features in other programs (Table 5.1) and presented the need for a spectral fitting tool of its nature.

In this chapter, we discuss results from Sunxspex and discuss the underlying methods being used. We initially present the results from fitting a single NuSTAR spectrum, indicating many core and peripheral features that can be utilised during this single fit. We then compare fitting results of NuSTAR spectra with Sunxspex in increasing degrees of complexity to published works (Section 6.1). We also return to microflare 3, analysed in Section 3.4.1, to fit the non-thermal component in the rise time spectrum with the thick target model, something unable to be performed in XSPEC at the time (Section 6.1.5).

In addition, we also discuss fitting results from the RHESSI and STIX observatories, showing novel approaches to their analysis through simultaneous fitting, something unable to have been performed in the OSPEX fitting tool (Section 6.2). Lastly, we present some of the first results from simultaneously fitting multiple solar X-ray spectroscopic data-sets from completely different observatories (Section 6.3).

Prior to my involvement, Sunxspex only contained non-thermal models and partial work on an isothermal model. My work, presented in this chapter, has added the ability to actually fit a range of solar X-ray data, with me creating data loaders for NuSTAR, RHESSI, and STIX data, wrappers and functions to interface the models with the data and with python fitting packages (Scipy and Emcee), and plotting and analysis tools for these results. As well as the examples shown in this chapter, I have created Sunxspex fitting example notebooks available on GitHub[1,2,3].

[1] https://github.com/sunpy/sunxspex/blob/master/examples/fitting_NuSTAR_spectra.ipynb
[2] https://github.com/sunpy/sunxspex/blob/master/examples/fitting_RHESSI_spectra.ipynb
[3] https://github.com/sunpy/sunxspex/blob/master/examples/fitting_custom_spectra.ipynb

6.1 Sunxspex Application to NuSTAR Spectra

NuSTAR data comes in the form of an event-list as it does not operate under set exposure times, or 'frames', but is continuously open to detection, unless it is within the 'deadtime' of already reading a recorded count. Therefore, we pre-process the NuSTAR data using the NuSTAR data pipeline, producing the spectral file (PHA file), the response matrix file (RMF file), and the ancillary response function file (ARF). The PHA, RMF, and ARF files are all needed to perform spectral fitting on NuSTAR data, with the counts data being stored in the PHA file and the spectral response matrix being created from both the RMF and ARF files. See Section 2.1 for more details regarding NuSTAR data recording and processing.

The capabilities of the Sunxspex spectral fitting software are vast and so we present several examples of fitting NuSTAR spectra in increasingly complex scenarios.

6.1.1 An Isothermal Example

We choose a very representative, simple isothermal spectrum in order to introduce the first Sunxspex spectral fit; therefore, we choose microflare 3's pre-flare time (10:26:50–10:28:30 UTC) from Section 3.4.1. Fitting just FPMA's emission from NuSTAR with a single thermal `f_vth` model in Sunxspex, we find a temperature of 4.7 MK and emission measure 3.3×10^{45} cm^{-3} (Figure 6.1, left panel).

Note that the Sunxspex result in Figure 6.1 cannot be directly compared to the XSPEC results shown in Figure 3.3 (bottom left panel) from Section 3.4.1 as we only fit one NuSTAR FPM here and, in addition, we use different atomic databases for the Sunxspex and XSPEC isothermal models and the fitted energy range is different. However, the values found from the XSPEC fit in Section 3.4.1 (temperature of 4.1 MK and emission measure 6.3×10^{45} cm^{-3}) are not too different compared to the Sunxspex result since, although the XSPEC temperature is lower, the emission measure is correspondingly higher to produce a similar fit.

This spectrum was found to be well located on a single detector for both NuSTAR FPMs—although we just use one here—and fitted well with a single thermal model (Figure 3.3, bottom left panel, and 6.1). Note, the thermal model name in Sunxspex is currently called `f_vth` and is identical to the `f_vth`[4] thermal model used in OSPEX. Both work by interpolating across look-up tables produced from the same `CHIANTI` database (see Section 1.4.4) which includes information on the continuum and line emission of the emitting plasma.

[4] https://hesperia.gsfc.nasa.gov/ssw/packages/xray/idl/f_vth.pro

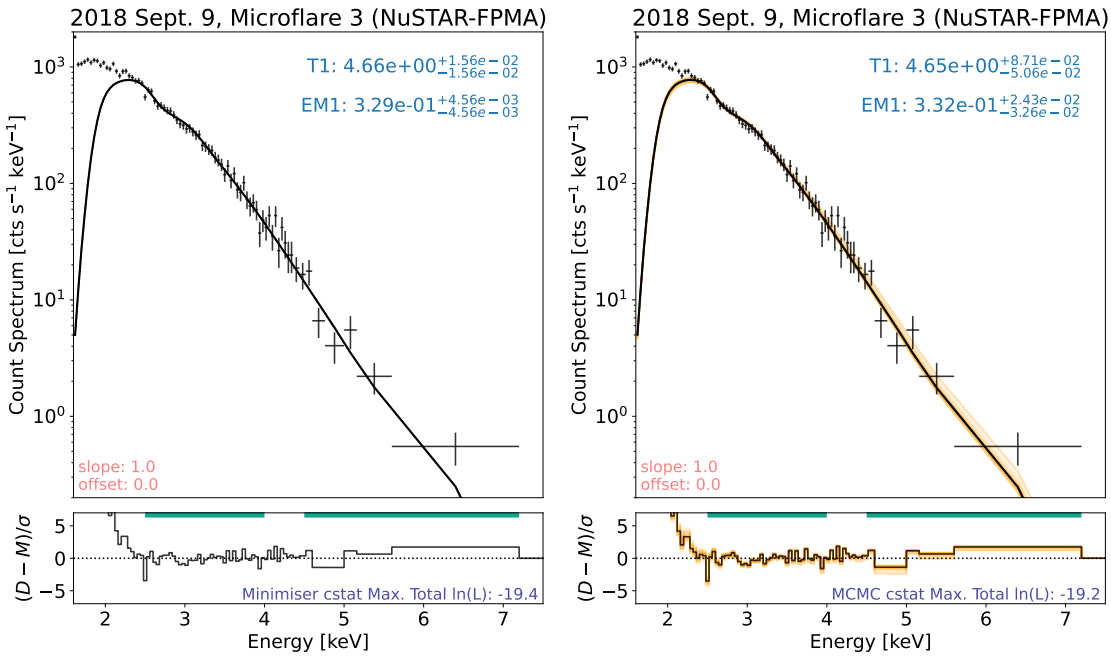

Figure 6.1: NuSTAR FPMA spectrum (black data-points), from Figure 3.3, fitted in Sunxspex with a single thermal model (black line, left panel) and utilising MCMC analysis (right panel) with the most probable fit result (the *maximum a posteriori*, black) and one hundred random samples in orange. The fitted thermal parameters are displayed in blue (units are MK and 10^{46} cm^{-3}), showing a temperature ($T1$) of 4.7 MK and emission measure ($EM1$) of 3.3×10^{45} cm^{-3}. The spectral data has been resampled for plotting such that each bin has a minimum of 10 counts to simplify the plot. The 'slope' and 'offset' values refer to the gain response parameters (see Section 2.1.5); they keep their default values here. The residuals panel shows the deviation of the fit from the data normalised by the error and the green lines show the 2.5–4 keV and 4.5–7.2 keV energy ranges the spectrum was fitted over. Text in the residuals panel shows whether the result is from fitting the model to the data (minimised) or sampled (MCMC) with the best C-stat (Section 5.2.1), logarithmic value.

In Figure 6.1, we show a simple feature of Sunxspex by fitting the NuSTAR spectrum over two energy ranges, 2.5–4 keV and 4.5–7.2 keV, ignoring any effect the bins between 4–4.5 keV have on the fitting process. This can be utilised to avoid including bins that are deemed to be problematic for any reason by the user or to exclude an energy range that is not appropriate for the current model being fitted. Figure 6.1 also shows the ability for Sunxspex to resample the count spectrum such that each data-bin displayed contains a minimum number of counts. The resampling can be performed to the data being fitted and applied when plotting (much like in OSPEX and XSPEC discussed in Section 5.3); however, we have only resampled for plotting purposes in Figure 6.1.

Sunxspex is able to accept all three NuSTAR fits files, the PHA, ARF, and RMF files (Section 2.1). As the NuSTAR processed files have a standard naming format, the names for all the files (PHA, ARF, and RMF) can either be given or just those for the PHA files and Sunxspex will auto-load the ARF and RMF files.

The model `f_vth` is assigned to the object that loaded in the NuSTAR PHA, ARF, and RMF files with the model attribute. This creates a parameter table where every model parameter in combination with every loaded spectrum, one in this case, is able to be set with a status, value, and parameter search bounds. We set the status of each parameter, whether it should be free to fit, fixed at a value, or tied to another value. All spectral parameters for the first loaded spectrum are free while any additional loaded spectra have their parameters tied to the first. The value and bounds of each parameter can be edited by the user while the values are automatically updated with the new resulting value from the fitting or sampling process.

In addition to fitting the spectrum via Sunxspex's default minimiser (Scipy's `minimize` function with the Nelder-Mead method), we can easily perform MCMC analysis (Section 5.2.3) on the loaded spectrum over the free parameters (Sunxspex's default is Emcee's method described in Goodman & Weare (2010)). Figure 6.1 (right panel) shows the resulting MCMC analysis to microflare 3's pre-flare time. The MCMC analysis is performed with four walkers with 2400 steps each. The prior is a uniform distribution between 2.5–8 MK and 1–8$\times10^{45}$ cm^{-3} for the temperature and emission measure, respectively. We find similar parameters from the *maximum a posteriori* to those found by the fitting process, temperature of 4.7 MK and emission measure 3.4$\times10^{45}$ cm^{-3}, with the fitted values being well within the MCMC results 1-σ uncertainty range.

The confidence range is visualised using a corner plot of the MCMC analysis and is shown in Figure 6.2, as discussed in Section 5.2.3. We show the log-probability (log-posterior) chain plot of all the MCMC walkers throughout the sampling (Figure 6.2, left panel) and the resulting corner plot of the walker parameter positions in relation to each other (right panels). We discard 20 samples which are clearly in the 'burn-in' phase of the MCMC, shown by a dotted orange line in the chain plot. The confidence area of the two-parameter MCMC sampling is shown in the middle-left corner plot panel with the temperature and the emission measure for the loaded spectrum ('spectrum1') on the x- and y-axis, respectively. We find that the marginalised posterior samples for each parameter (diagonal panels) show Gaussian-like distributions. The bottom row in the right panel shows how the non-normalised posterior value, the multiplication of the defined likelihood and prior, is distributed with respect to the walker positions.

6.1.2 Two Telescopes with a Multi-Thermal Model Example

To further show the capabilities of Sunxspex we now fit microflare 10's pre-flare spectra in the same fashion as the analysis performed in Section 3.7. The spectrum of FPMA & B are fitted well with two parametric thermal models. We fit both FPMA & B spectra obtained over the same time and energy ranges (Figure 6.3) as is common with previous XSPEC analysis discussed (Chapter 3). We load in

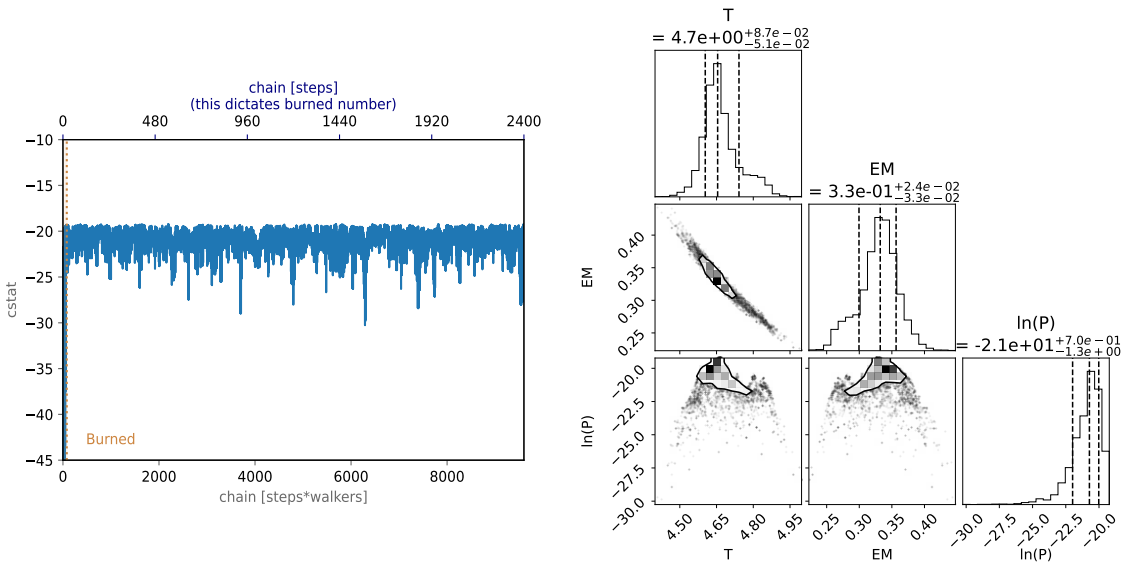

Figure 6.2: The walker log-probability chain (left panel) and the resulting corner plot (right panel) corresponding to Figure 6.1's right panel produced in Sunxspex. The chain plot shows the behaviour of the walkers throughout their random sampling and indicates the convergence of the walkers in the posterior-space. The corner plot is produced from the walker samples, showing the walker probability density for the temperature and emission measure posterior. The corner plot diagonal panels show the marginalised walker density distributions.

both NuSTAR FPMA and FPMB files and assign the model to `C(f_vth+f_vth)`. This model includes two thermal photon sub-models and a scalar `C`, again building a parameter table. There are two loaded spectra, labelled 'spectrum1' and 'spectrum2', where each get assigned all model parameters; for example, both get a $T1$, an $EM1$, and so on. In order to fit these spectra simultaneously we tie all of spectrum2's parameters to those of spectrum1. However, the C scalar is fixed to a value of 1 for spectrum1 and allowed to vary for spectrum2.

The fitting then takes place by searching over $T1$, $EM1$, $T2$, $EM2$ and spectrum2's C. Each model is then determined by their own parameters from the table and a fit statistic is calculated for both spectra, it is the sum of these fit statistics that is then optimised during the parameter search.

Since both spectrum1 and spectrum2 are similar instruments, with the same energy resolution, it is trivial to produce a plot of the average emission and fit between telescopes (Figure 6.3, right panel). It is this style of plot shown in Figure 3.15 (bottom left panel) for the pre-flare spectrum of microflare 10. Qualitatively, we find similar features in the residuals between both plots even though different methods have been used to resample the count spectrum.

As well as visual corroboration, we also find good quantitative agreement between the fits. Table 6.1 shows values obtained from XSPEC (Cooper et al. 2021) and those obtained from the Sunxspex

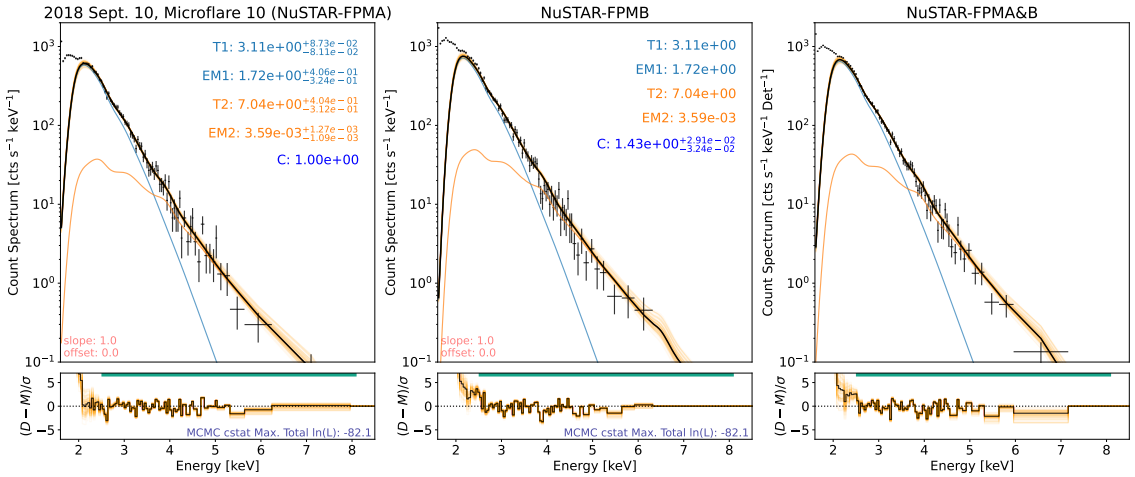

Figure 6.3: NuSTAR pre-flare spectra from Figure 3.15. Each panel is in the same format as described in Figure 6.1 with the addition of a scalar C which accounts for any systematic offset between both instruments. The resulting fit for FPMA, FPMB, and FPMA & B is shown in the left, middle, and right panel, respectively.

Table 6.1: Spectral fitted parameters from XSPEC and sunxspex corresponding to Figure 3.15 and 6.3. We include the minimised Sunxspex fit results in the middle column.

Model Parameter	XSPEC (Cooper et al. 2021)	Sunxspex (Fit)	Sunxspex (MCMC)
Temperature 1 [MK]	$3.05^{+0.04}_{-0.35}$	3.26 ± 0.01	$3.11^{+0.09}_{-0.08}$
Emission Measure 1 [cm^{-3}]	$1.70^{+1.99}_{-0.08}\times10^{46}$	$1.18\pm0.01\times10^{46}$	$1.72^{+0.41}_{-0.32}\times10^{46}$
Temperature 2 [MK]	$6.60^{+0.20}_{-0.61}$	6.92 ± 0.04	$7.04^{+0.40}_{-0.31}$
Emission Measure 2 [cm^{-3}]	$3.80^{+4.00}_{-0.70}\times10^{43}$	$3.77\pm0.05\times10^{43}$	$3.59^{+1.27}_{-1.09}\times10^{43}$
C	$1.43^{+0.03}_{-0.03}$	1.42 ± 0.02	$1.43^{+0.03}_{-0.03}$

analysis performed in this section. The relatively large value for the scalar C is likely due to the source being closer to FPMA's detector chip-gap.

From Table 6.1, we find some values are slightly different from the XSPEC values compared to the Sunxspex fitted ones, with values within or almost within uncertainty margins; however, it is important to note that XSPEC and Sunxspex utilise different atomic databases, but that difference should be small. In addition, we find, for a similar isothermal fit, the temperature can decrease/increase if the emission measure increases/decreases. Therefore, fitting two isothermal models allows for more variation. However, the XSPEC confidence range was obtained through MCMC analysis; therefore, we compare with the Sunxspex equivalent process. Here we find very good agreement between XSPEC and Sunxspex values with all parameters having a significant overlap in their 1-σ (68%) confidence

ranges. The difference between the fitted values and the MCMC values stress the dependence of the fit on the initial parameter 'guesses' within Sunxspex and also the underestimation of the error margins via the Hessian approach (Section 5.2.4).

6.1.3 Comparisons to Glesener et al. (2020)

We have discussed many of the abilities that Sunxspex boasts compared to other spectral fitting tools; therefore, we will continue by re-analysing more published works, this time from Glesener et al. (2020). We start by performing the analysis for the left panel of Figure 4 in Glesener et al. (2020), the result of which is shown in Figure 6.4.

We choose to re-analyse the spectral fitting results with the same data used in Glesener et al. (2020) as this spectrum is one of the weakest AR microflares that has very strong evidence for non-thermal emission. Despite its weak nature relative to many other studied flares, the non-thermal spectral signal is strong with many counts, even in the higher energy ranges ($>$7 keV). In Glesener et al. (2020), this event needed to be fitted in OSPEX to make use of the non-thermal thick target models. However, using OSPEX meant that half of the spectral data available had to be neglected, a scenario in which Sunxspex becomes particularly useful.

Figure 6.4 shows the recreation of the left panel of Figure 4 in Glesener et al. (2020). Starting with fitting the NuSTAR FPMB spectrum with a thermal model and the cold thick target we find a reasonable fit, when looking at the residual panel, of these two sub-models.

From Table 6.2, we find that the fit from Glesener et al. (2020) and the equivalent fit from Sunxspex produce parameter values within uncertainty ranges of each other. This shows that Sunxspex is able to re-produce similar results to the well-established fitting software, OSPEX. In addition, we take the analysis one step further and utilise the data provided by FPMA's observation to perform simultaneous fitting which could not be done in OSPEX.

Figure 6.5 shows the result of including the FPMA spectrum and the right column of Table 6.2 shows the resulting fitted parameters. We find, again, that the majority of parameters share an uncertainty region with the published OSPEX values with the exception of the temperature which does not exactly overlap. This could be due to the use of Poisson statistics in Sunxspex compared to Gaussian in OSPEX. Moreover, the addition of the FPMA spectrum in the Sunxspex fitting, which could not be included in the OSPEX result, will help to further constrain the value of these parameters.

It should be noted, however, that the errors on the Sunxspex values are produced from the Hessian method and so are likely to be underestimating the true uncertainty range, as discussed in Section 5.2.4.

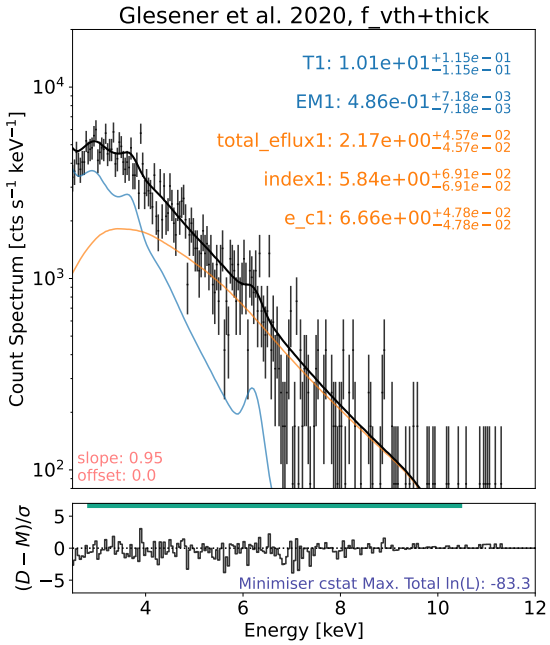

Figure 6.4: Spectrum from Glesener et al. (2020) fitted with an isothermal and cold thick target model. The units for the temperature ($T1$) and emission measure ($EM1$) are the same as described in Figure 6.1. The total electron flux (total_eflux) and the low-energy cut-off being in units of 10^{35} e$^-$ s^{-1} and keV, respectively, with the electron spectral index (index) being dimensionless.

We also set the gain slope value to 0.95 for all fits and the scalar C value in the simultaneous fitting to 1.1 as described in Glesener et al. (2020). The gain correction is discussed further in Section 2.1.5.

6.1.4 Comparisons to a Duncan et al. (2021) Microflare

We further increase the complexity of analysis being performed in the new Sunxspex fitting software by comparing with a result obtained with XSPEC in Duncan et al. (2021) using the same data. This example shows a similar number of free parameters being fitted to a NuSTAR spectrum compared to the example in Section 6.1.3; however, one of the parameters is an instrument response parameter and not a model parameter (Section 2.1.5). This response parameter, or gain, helps define how the input photon energies to the detector are converted into counts energies; the gain slope is a multiplicative factor in this conversion. The response parameters are created and handled similarly to the model parameters; however, they are maintained in a separated response parameter table within the python spectral object.

The spectrum we choose is used to illustrate the need for a livetime dependent gain correction to be applied to NuSTAR spectral analysis as detailed in Duncan et al. (2021). In addition, both FPMs

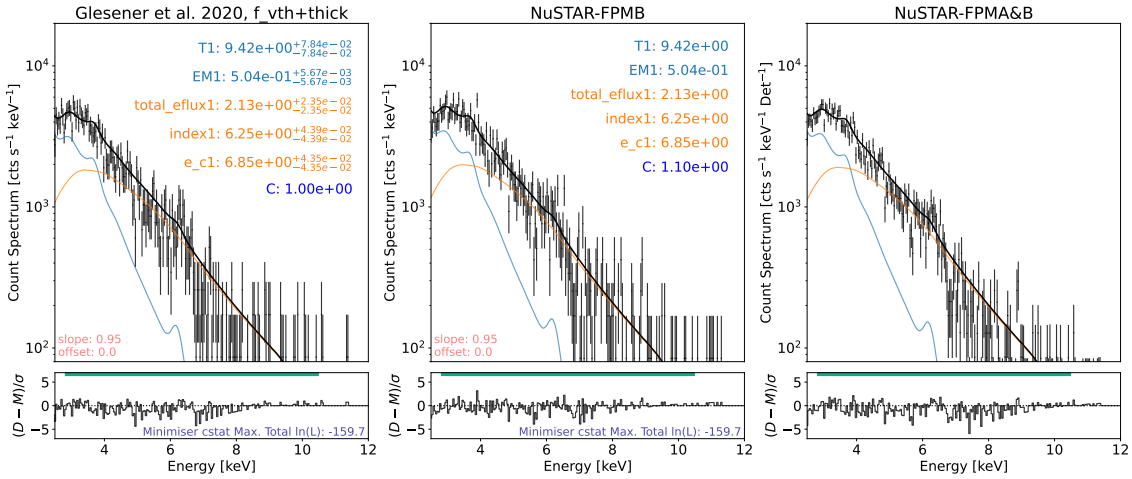

Figure 6.5: NuSTAR FPMA (left panel) and B (middle panel) spectra from Glesener et al. (2020) simultaneously fitted in Sunxspex with an isothermal and cold thick target model. The units for each model are shown in the appropriate colours and the units are the same as those described in Figure 6.4. The right panel shows the fit on the mean FPMA & B spectrum.

are used and simultaneously fitted with two thermal models. We replicate this fitting by first fitting over the lower energy ranges (3–6 keV) with one thermal model the higher energy range (6–10.8 keV) with the second, allowing the gain parameter to vary throughout all fitting. We then fit the total, double-thermal model over the whole energy range.

Figure 6.6 shows the 'May1618' microflare from Duncan et al. (2021) analysis performed in Sunxspex. From visual inspection, we find the fit has been performed well and does not show any obvious deviation from the published spectral fit of Figure 10 in Duncan et al. (2021). In Figure 6.6, we show

Table 6.2: Thermal and cold thick target fit results for the microflare in Glesener et al. (2020) using OSPEX and Sunxspex .

Model Parameter	OSPEX (FPMB) (Glesener et al. 2020)	Sunxspex (FPMB)	Sunxspex (FPMA&B)
Temperature [MK]	10.3±0.7	10.1±0.1	9.4±0.1
Emission Measure [cm^{-3}]	5.0±1.3×10^{45}	4.9±0.1×10^{45}	5.04±0.01×10^{45}
Electron Flux [e^{-} s^{-1}]	2.1±1.2×10^{35}	2.17±0.05×10^{35}	2.13±0.02×10^{35}
Electron Spectral Index	6.2±0.6	5.8±0.1	6.25±0.04
Low-energy Cut-off [keV]	6.2±0.9	6.66±0.05	6.85±0.04

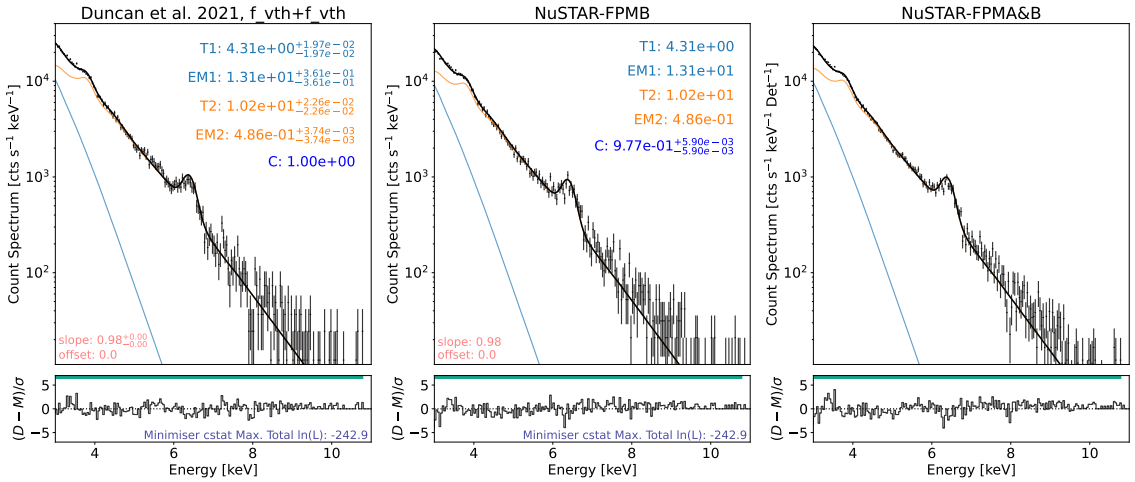

Figure 6.6: The May1618 microflare from Duncan et al. (2021) fitting in Sunxspex. NuSTAR FPMA and B fitted with two thermal models with the gain slope response parameter optimised during the fitting due to a known livetime dependent NuSTAR effect (Section 2.1.5). All parameters and plot details the same as those defined in Figure 6.3.

the results of the simultaneous fit with FPMA (left panel), FPMB (middle panel), then with the average FPMA and B emission (right panel).

We present the results of the Sunxspex fitting, in comparison to the published XSPEC result, in Table 6.3. The results from the published XSPEC results (Duncan et al. 2021) and those obtained using Sunxspex are incredibly similar.

Note that the fits from these different spectral analysis tools are not entirely equivalent due to different atomic databases being used for the thermal model calculations; however, it is clear this does not cause a significant difference. In addition, different parameter search methods are used between XSPEC and Sunxspex when optimising the fit with the uncertainty ranges also being calculated differently.

We have shown in Section 6.1.2–6.1.4 that Sunxspex is able to perform analysis from published works. Moreover, in Section 6.1.5, we progress the analysis of Section 3.4.1's microflare 3 using Sunxspex, beyond what is achievable with OSPEX or XSPEC.

6.1.5 Return to the Rise of 2018 September's Microflare 3

We return to the rise time of microflare 3 (Section 3.4.1) where we were unable to sufficiently investigate the presence of non-thermal emission with the thick target model because of the limitations of both XSPEC (no thick target model) and OSPEX (unable to simultaneously fit spectra or utilise Poisson statistics).

Table 6.3: Double-thermal fit results using XSPEC and Sunxspex for the May1618 microflare in Duncan et al. (2021).

	OSPEX (FPMA&B) (Duncan et al. 2021)	Sunxspex (FPMA&B)
Model Parameter		
Temperature 1 [MK]	$4.1^{+0.2}_{-0.1}$	4.31 ± 0.02
Emission Measure 1 [cm^{-3}]	$1.4^{+0.6}_{-0.4}\times10^{47}$	$1.3\pm0.4\times10^{47}$
Temperature 2 [MK]	$10.00^{+0.03}_{-0.03}$	10.20 ± 0.02
Emission Measure 2 [cm^{-3}]	$4.6^{+0.1}_{-0.2}\times10^{45}$	$4.86\pm0.04\times10^{45}$
Response Parameter		
Gain Slope	0.977 ± 0.002	0.977 ± 0.001

Figure 6.7 shows the result of fitting microflare 3's rise time spectra in Sunxspex with a pre-flare thermal component and a microflare thermal component—as in Section 3.4.1 with XSPEC—with the photon broken power-law model being replaced with the thick target model (Section 1.4.5). This allows us to gain insight into the electron distribution which produced the non-thermal photon spectrum, giving more reliable estimates of the non-thermal power.

We performed MCMC analysis, utilising C-stat (Section 5.2.1), on the pre-flare spectrum of microflare 3 (Figure 3.3, bottom left panel) using Sunxspex and find a most probable temperature and emission measure of $4.6^{+0.1}_{-0.1}$ MK and $3.7^{+0.5}_{-0.4}\times10^{45}$ cm^{-3}, respectively. This gives a slightly higher temperature and lower emission measure parameters compared to the XSPEC results. We keep this pre-flare constant for the rise phase fit.

Fitting the rise stage of microflare 3 with a thermal and thick target model, on top of the fixed pre-flare (Method 1 in Section 5.2.6), we find that the impulsive phase produces a characteristic temperature of $7.1^{+0.2}_{-0.2}$ MK and emission measure of $5.6^{+0.8}_{-0.7}\times10^{44}$ cm^{-3}. In the case of this hotter thermal component, we do find good agreement between the Sunxspex here and the XSPEC values in Section 3.4.1; i.e., they share an overlap in the parameter uncertainty margins.

Fitting in Sunxspex has allowed us to properly fit the non-thermal emission identified in Section 3.4.1 >7 keV in both FPMs with the thick target model. We find a total electron flux value of $8.4^{+3.3}_{-2.7}\times10^{32}$ e^{-} s^{-1}, an electron spectral index of $9.6^{+1.4}_{-1.0}$, and a low-energy cut-off of $6.8^{+0.4}_{-0.3}$ keV. This results in a non-thermal power of $1.04^{+0.35}_{-0.30}\times10^{25}$ erg s^{-1}, compared with the more crudely estimated $7.03^{+3.67}_{-2.32}\times10^{24}$ erg s^{-1} from the photon broken power-law representation in Figure 3.5. These values are consistent with each other and so our conclusions from Section 3.4.1 do not change

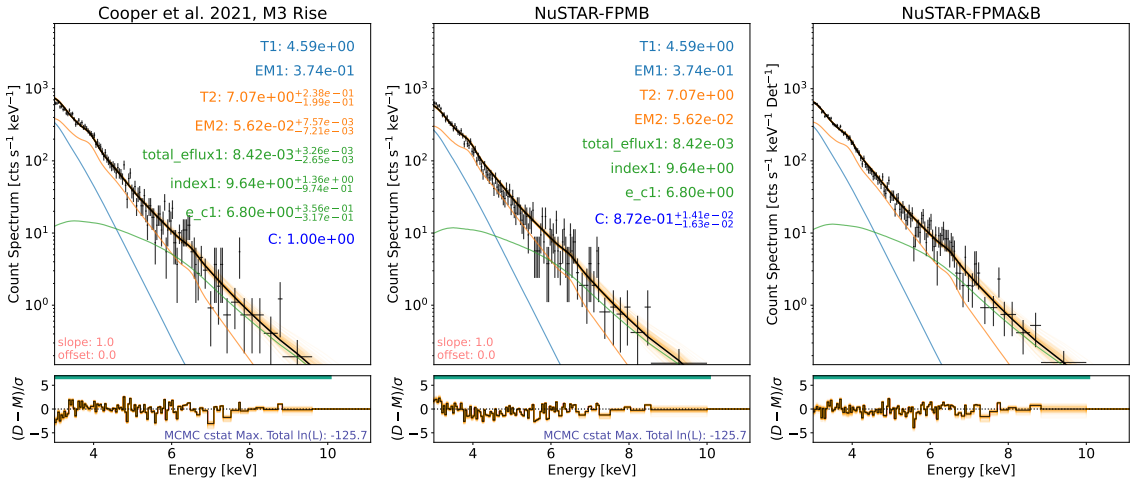

Figure 6.7: Microflare 3's rise spectra (from Section 3.4.1) analysed using Sunxspex. The same format as Figure 6.5 with the gain slope response parameter being unchanged from its default value of 1. The spectra are resampled such that every count energy bin contains at least 2 counts for clarity.

significantly; however, the more robust non-thermal power gives a slightly higher total non-thermal energy, bring it closer to the upper thermal energy estimate for microflare 3.

Figure 6.8 shows the corner plot for the MCMC analysis used to produce Figure 6.7. It shows the sampled walker positions throughout the random walking process which results in a probability density of the posterior parameter-space. The marginalised walker distributions (the diagonal panels) all show normal-like shapes with the bulk of the probability distribution sampled well within the prescribed priors. This shows the robustness of the resulting fits to the data.

The result from Sunxspex corroborates the published XSPEC findings in Cooper et al. (2021) and gives us an insight into the non-thermal electron distribution which produced the measured non-thermal emission. Note, we also find the same unreasonable behaviour in Sunxspex when using another thermal model for this emission as we did in the XSPEC case discussed in Section 3.4.1 (Figure 3.4, 3.5, 3.6, and 3.7). Being able to use these solar specific models as well as utilise all spectral data and appropriate fit statistics showcases the power and need for a spectral fitting software such as Sunxspex.

6.2 Application to RHESSI and STIX Spectra

The RHESSI and STIX processed spectrum files contain count, or count flux, spectra vs time and energy from the whole solar disk. This means that it is far simpler to define an event and possibly a background time with that being applied to the data more dynamically.

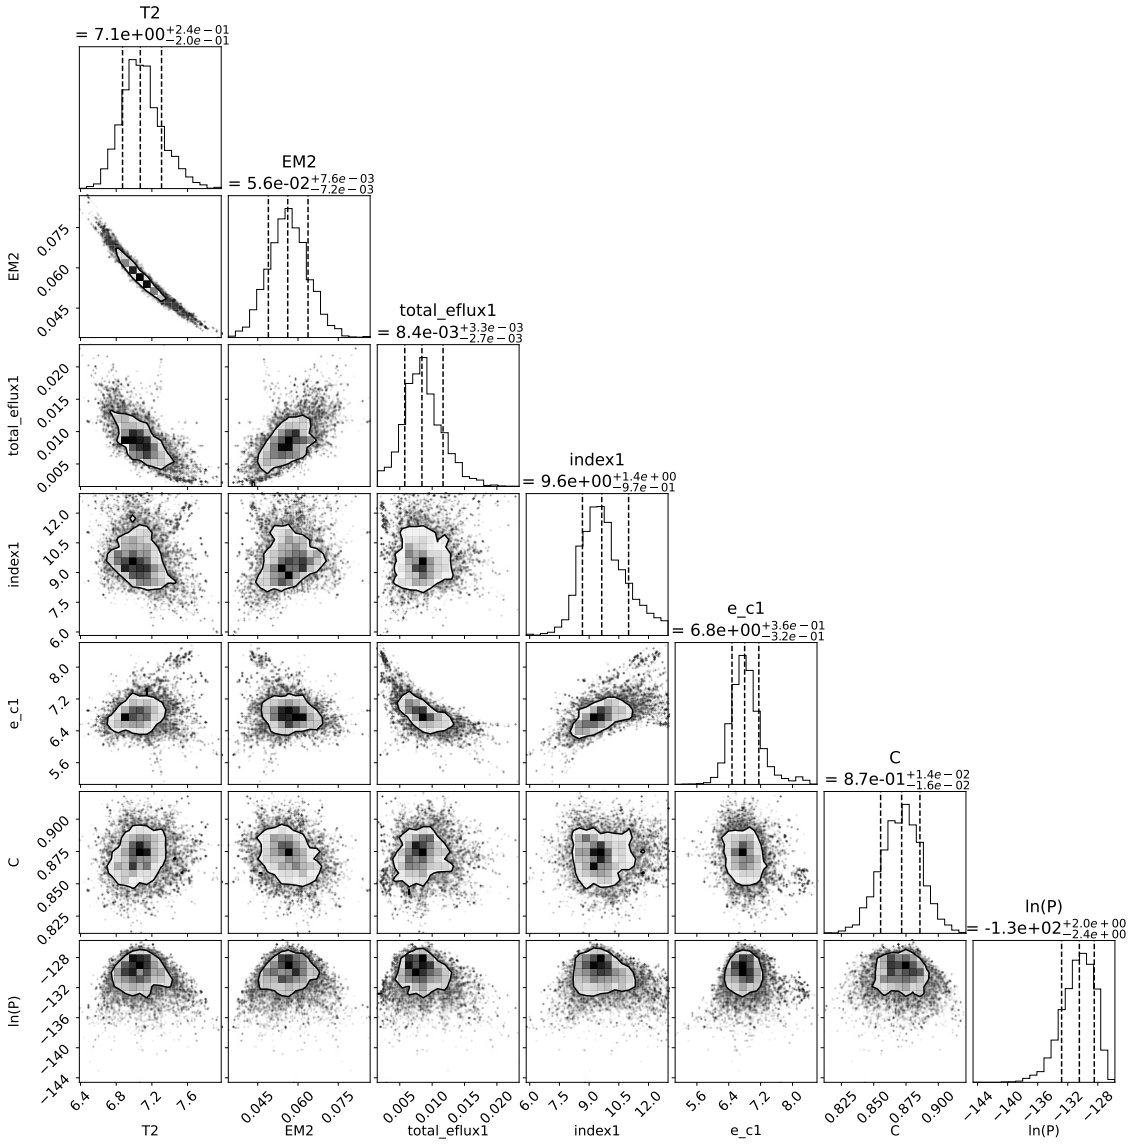

Figure 6.8: Corner plot of the MCMC analysis used to produce Figure 6.7 performed in Sunxspex. The contours in the off-diagonal panels show 68% (1-σ equivalent) of the walker density and how each parameter relates to every other one, while the diagonal panels show the marginalised distributions of each parameter. The most probable values, with the 1-σ equivalent uncertainty range, are shown above each diagonal panel.

We present several examples of using RHESSI and STIX data products in Sunxspex. We perform spectral analysis from published work involving a RHESSI spectrum summed across several detectors (Section 6.2.1). We then investigate that same observation but, instead we simultaneously fit all used RHESSI detectors (Section 6.2.2). Finally, we show an example of fitting a STIX spectrum in Section 6.2.3.

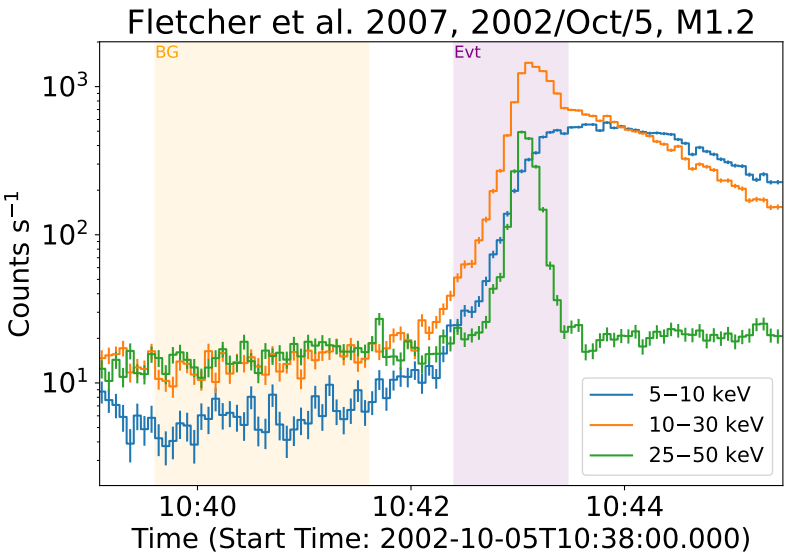

Figure 6.9: Time profiles for three RHESSI energy ranges during the Fletcher et al. (2007) 2002 October 5 flare. The background spectrum is integrated over the orange region and the flare time is integrated over the magenta shaded region.

6.2.1 A RHESSI Spectrum in Sunxspex

We choose to investigate a published RHESSI flare that occurred on 2002 October 5 (Fletcher et al. 2007) and achieved a GOES class of M1.2 (scaled: M1.7, see Section 2.2). We follow the approach described in the previous analysis, performed in OSPEX, and define the background time to be 10:38:32–10:40:32 UTC and the flare time to be 10:41:20–10:42:24 UTC. This is shown in Figure 6.9 where the time profiles of three RHESSI energy ranges, produced in Sunxspex, are depicted. The background time range is indicated by the orange region and the flare time by the magenta region.

Figure 6.10 shows this flare is fitted well with an isothermal model between 6–15 keV with clear indication of non-thermal emission at 15–45 keV; therefore, we utilise a thermal model along with the thick target model to represent the flare spectrum (10:41:20–10:42:24 UTC) which has been summed over six RHESSI detectors. Here we show two approaches to fitting a spectrum with a background, the left panel shows the RHESSI background-subtracted spectrum being fitted (Method 5 in Section 5.2.6) while the right panel shows the result of fitting the total flare time spectrum with the models combined with the background (Method 3 in Section 5.2.6). We see that, in this case, there is no significant disagreement of values which is likely due to the RHESSI spectrum having a large number of counts in each bin and, therefore, being in a Gaussian statistical regime.

Fletcher et al. (2007) finds a temperature of 24.1 MK and an emission measure of 1.6×10^{47} cm^{-3} from fitting the background-subtracted spectrum. The non-thermal emission had to be approximated

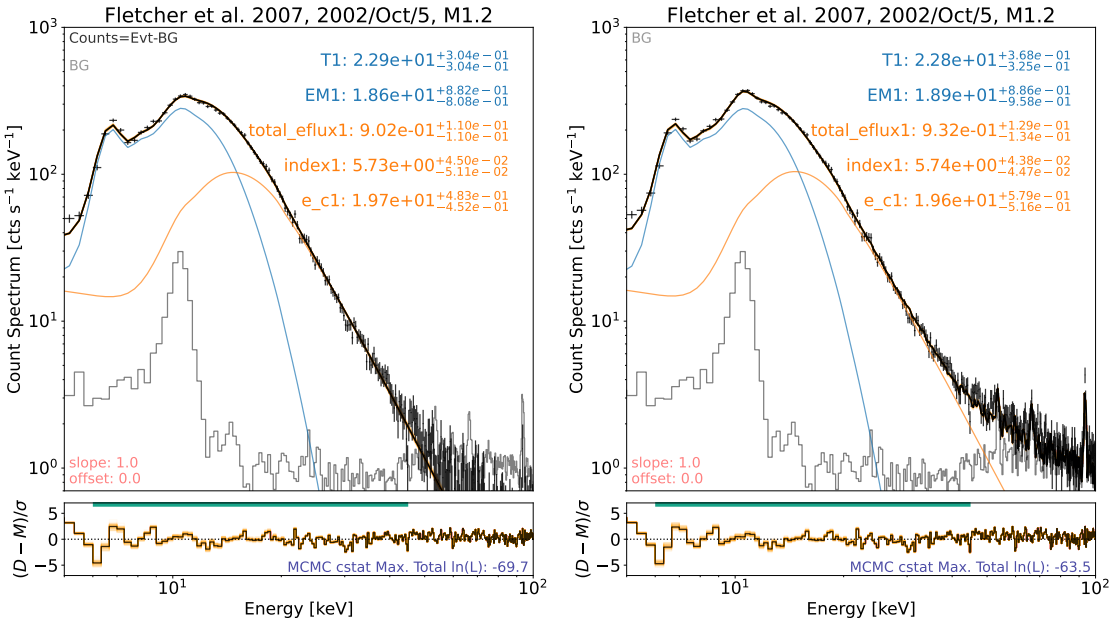

Figure 6.10: RHESSI Fletcher et al. (2007) 2002 October 5 flare spectral MCMC results with Sunxspex. The background-subtracted spectrum is fitted with a thermal (blue) and non-thermal (orange) model (left panel) and the full flare spectrum is fitted with both models plus the background (right panel). The former and latter are examples of Method 5 and 3 in Section 5.2.6, respectively. Different background handling methods only appear obvious at the higher energies when the background spectrum (grey) dominates >50 keV. The spectrum being fitted (which is the result of summing over RHESSI detectors 1, 3, 4, 6, 8, and 9) is shown with black data-points and the total model being fitted is shown with a black line. The parameters values shown have the same meaning as those in Figure 6.4.

with a photon broken power-law model. From the MCMC analysis in Figure 6.10, we find similar parameters to those in Fletcher et al. (2007), but they are not consistent within the 68% confidence ranges. However, this fit would have been performed prior to important improvements to OSPEX and RHESSI calibration; for example, the error calculation for the background-subtracted spectrum was not complete until 2010[5]. This could account for the discrepancy as the fit depends on the data-point uncertainties due to the Gaussian statistics used in OSPEX.

Therefore, to compare with the RHESSI spectral analysis here, the background-subtracted spectrum has been fitted again in OSPEX[6] giving a temperature of 22.5 MK and an emission measure of 2.0×10^{47} cm^{-3}. However, now we have used the thick target instead of a photon broken power-law model approximation giving a total electron flux of 9.8×10^{34} e^{-} s^{-1}, a low-energy cut-off of 19.4 keV, and an electron power-law index of 5.8.

From the equivalent Sunxspex fit, the background-subtracted analysis shown in Figure 6.10 (left

[5]OSPEX change-log (2010 September 29).
[6]Private communications, Iain Hannah.

Table 6.4: Thermal and cold thick target fit results using OSPEX and Sunxspex for the 2002 October 5 RHESSI flare. The temperature and emission measure from (Fletcher et al. 2007) are stated as 24.1 MK and 1.6×10^{47} cm^{-3}, respectively; however, the OSPEX values presented here come from a more robust and recent fit. The spectrum being fitted is indicated as either the background-subtracted spectrum, 'Flare-BG', or the whole flare spectrum, 'Flare'.

Model Parameter	OSPEX (Flare-BG)	Sunxspex (Flare-BG)	Sunxspex (Flare)
Temperature [MK]	22.5	$22.9^{+0.3}_{-0.3}$	$22.8^{+0.4}_{-0.3}$
Emission Measure [cm^{-3}]	2.0×10^{47}	$1.9^{+0.1}_{-0.1}\times10^{47}$	$1.9^{+0.1}_{-0.1}\times10^{47}$
Electron Flux [e^{-} s^{-1}]	9.8×10^{34}	$9.0^{+1.1}_{-1.1}\times10^{34}$	$9.3^{+1.3}_{-1.3}\times10^{34}$
Electron Spectral Index	5.77	$5.73^{+0.05}_{-0.05}$	$5.74^{+0.04}_{-0.04}$
Low-energy Cut-off [keV]	19.4	$19.7^{+0.5}_{-0.5}$	$19.6^{+0.6}_{-0.5}$

panel), we find a temperature of 22.9 MK, an emission measure of 1.9×10^{47} cm^{-3}, a total electron flux of 9.0×10^{34} e^{-} s^{-1}, a low-energy cut-off of 19.7 keV, and an electron power-law index of 5.7. Therefore, we find excellent agreement between the Sunxspex and OSPEX result, shown in Table 6.4. The agreement between the OSPEX χ^2 result and the Sunxspex C-stat result corroborates the fact that the RHESSI data is in the Gaussian regime.

We, therefore, show the ability to re-analyse published RHESSI spectra using Sunxspex. However, we can easily perform simultaneous spectral fitting of the detectors which contribute to the spectrum in Figure 6.10 instead of those detectors, each with their own inherent differences, being summed.

6.2.2 Spectra From Multiple RHESSI Detectors in Sunxspex

The RHESSI spectrum from Section 6.2.1 is the summed combination of RHESSI detectors 1, 3, 4, 6, 8, and 9. With OSPEX, the count information from each of these detectors have to be summed into one spectrum, with one average spectral response matrix, to be fitted. This can mask systematic differences between the detectors; however, this summing is not a necessity in Sunxspex.

Figure 6.11 shows the result of fitting all six RHESSI detectors simultaneously. We use the same energy fitting range defined in Section 6.2.1 for each spectrum as well as the same background and event time ranges. The same component models, comprised of a thermal and non-thermal model, are fitted to all full spectra (*not* background-subtracted) with only a relative scaling factor (C) allowed to vary between spectra to account for any systematic differences.

Fitting the detector spectra simultaneously, we find similar parameters to those found from fitting the summed spectrum in Figure 6.10 (right panel). We find a temperature of 22.8 MK and an emission

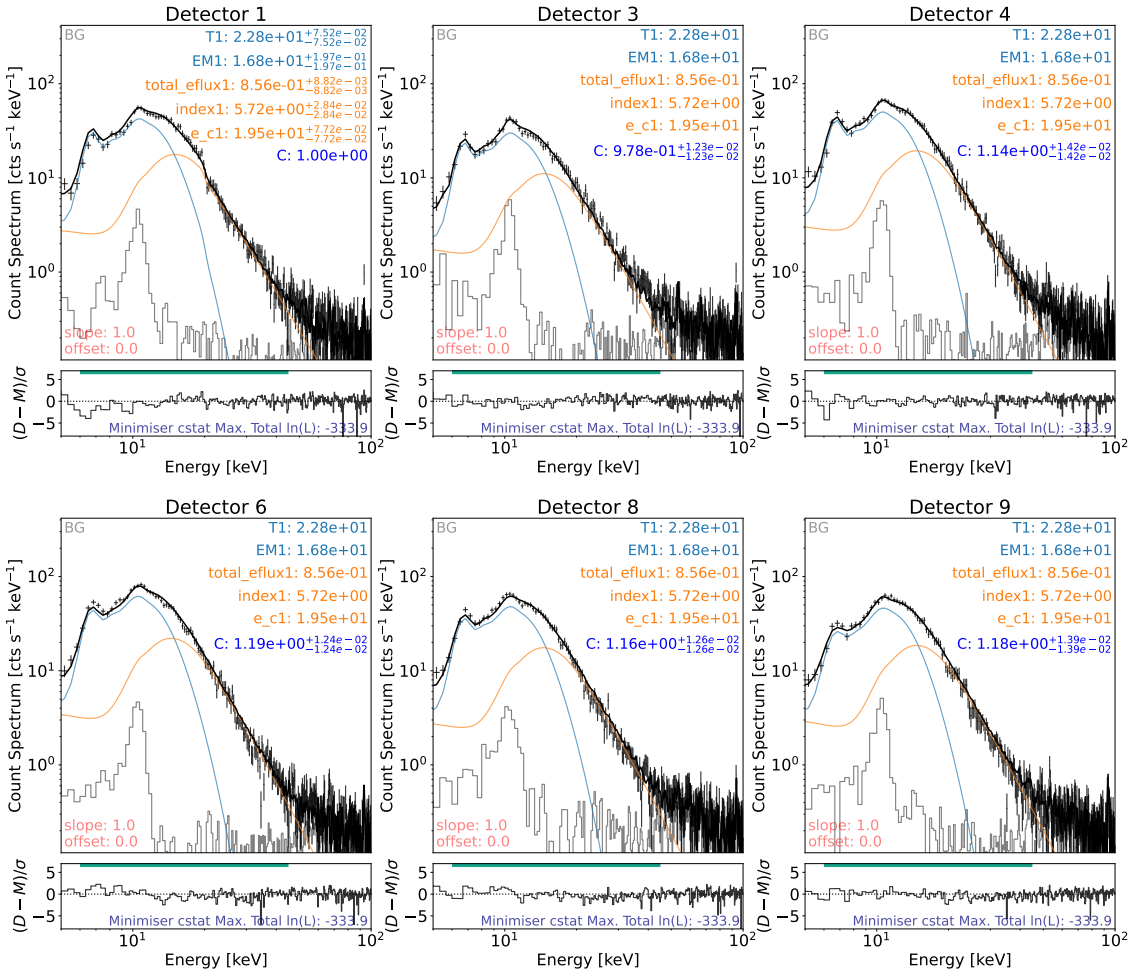

Figure 6.11: RHESSI spectrum of M-class flare per detector in Sunxspex. The individual RHESSI detector spectra, labelled, that are summed to produce the spectrum displayed in Figure 6.10. The fitting process, method, and model is done in the same fashion as the summed RHESSI spectrum except with all detectors being fitted simultaneously. The same combination of a thermal and non-thermal model is used with only a scalar factor, C, allowed to vary between spectra to account for any systematic uncertainty.

measure of 1.7×10^{47} cm^{-3} with a total electron flux, a low-energy cut-off, and an electron power-law index of 8.6×10^{34} e^{-} s^{-1}, 19.5 keV, 5.7, respectively.

From Table 6.5, there is very good agreement between the fitted parameters found from the summed spectrum and the simultaneously fitted one. This is expected to be the case if all spectra respond similarly to the same incoming emission; however, we ultimately lose statistical information in the Poissonian regime and any oddities within each detector spectrum fit due to additively combining data-sets. Therefore, simultaneous fitting allows us to robustly handle the statistical fitting and gives

Table 6.5: Results from the simultaneous fitting of six RHESSI detectors shown in Figure 6.11. The summed spectrum fitted parameters come from the 'Sunxspex (Flare)' columns in Table 6.4. Note that the simultaneously fitted parameter uncertainties are calculated using Gaussian assumptions and, therefore, are likely to underestimate the true 1-σ error.

Model Parameter	Summed Detector Fitting	Simultaneously Fitted Detectors
Temperature [MK]	$22.8^{+0.4}_{-0.3}$	22.8 ± 0.1
Emission Measure [cm^{-3}]	$1.9^{+0.1}_{-0.1} \times 10^{47}$	$1.68 \pm 0.02 \times 10^{47}$
Electron Flux [e^{-} s^{-1}]	$9.3^{+1.3}_{-1.3} \times 10^{34}$	$8.6 \pm 0.1 \times 10^{34}$
Electron Spectral Index	$5.74^{+0.04}_{-0.04}$	5.72 ± 0.03
Low-energy Cut-off [keV]	$19.6^{+0.6}_{-0.5}$	19.5 ± 0.1

Table 6.6: Relative to detector 1, the scaling C factors obtained to allow for systematic offsets when fitting six RHESSI detectors (Figure 6.11). All factors have an estimated uncertainty of ± 0.01.

	Detector 3	Detector 4	Detector 6	Detector 8	Detector 9
Scalar Factor C	0.98	1.14	1.19	1.16	1.18

us a method to investigate the relative response and state of each detector across any margin of time, energy, or space.

The scaling factor, C, is measured relative to the first loaded spectrum, detector 1's spectrum, with all values shown in Table 6.6. Using this factor, we can monitor the performance of each detector compared to one another as well as inspect for potential signs of unexpected systematic differences and the presence of significant or insignificant spectral features in count-space between the spectra. We find, in this example, that the scaling factors across the detectors only vary by ~20% with respect to one another.

Being able to fit RHESSI spectra in a similar fashion to what can be done in OSPEX with additional capabilities is of great benefit to the solar X-ray spectroscopic community within itself. We continue to perform spectral fitting with Sunxspex on Solar Orbiter's STIX in Section 6.2.3.

6.2.3 A STIX Spectrum in Sunxspex

Solar Orbiter's STIX, the instrument from the newest mission discussed in this chapter, stores its data in the form of counts against time and energy produced from the level 1 STIX data files via SSW or

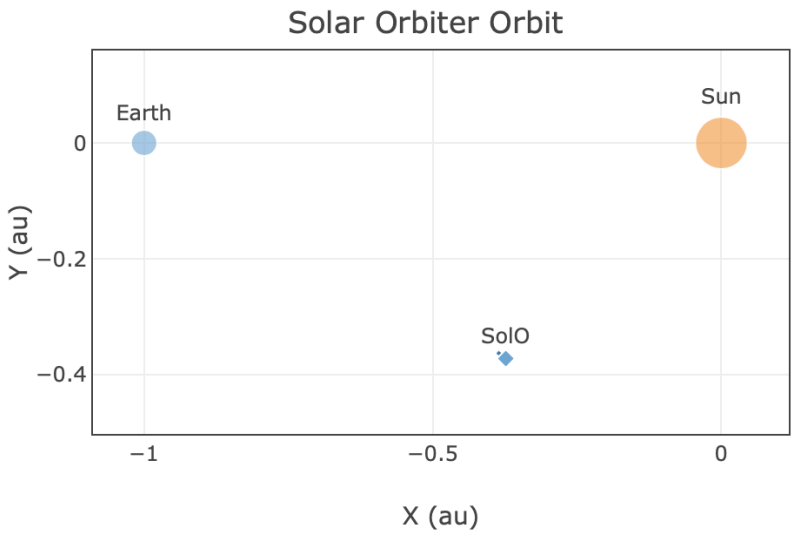

Figure 6.12: The location of Solar Orbiter and, therefore, STIX on 2020 June 6 in relation to the Earth and the Sun. Figure obtained using the STIX Data Center's Ancillary Data Browser.

STIXpy[7], similar to RHESSI data. However, analysing STIX data requires additional care due to its varying distance to the Sun (see Section 2.4.1).

Figure 6.12, obtained using the STIX Data Center's Ancillary Data Browser[8], shows the location of Solar Orbiter on 2020 June 6 and, therefore, STIX's position in relation to the Earth and the Sun. At the time of this microflare we see that STIX is a little over half an au from the Sun compared to other instruments like NuSTAR and RHESSI which are located at 1 au. Therefore, will see emission from the Sun sooner and with a greater flux than those instruments based around Earth.

STIX observed a microflare, SOL2020-06-06T19:47, at 19:47 UTC (time corrected to 1 au) on 2020 June 6. We choose to analyse this microflare as it was also observed by NuSTAR, we will revisit the NuSTAR spectra for this microflare in Section 6.3.1. Figure 6.13 shows the microflare time with the shaded magenta region (19:46:31–19:48:58 UTC) and a background time range shown with the orange shaded region (19:40:00–19:46:30 UTC). The microflare profile is prominent in the the 5–7 keV energy range (blue) with a relatively weaker, but still present signal in the 7–10 keV energy range.

Figure 6.14 shows the fit to the STIX spectrum using the MCMC approach. We fit all counts from the microflare time with a thermal model (black line) and the background (grey) across 5–10 keV. At the time of observation STIX was 0.53 au from the Sun; therefore, the photon model spectrum is increased by a factor of 3.6 before being convolved with the spectral response matrix and compared to

[7] https://github.com/samaloney/stixpy
[8] https://datacenter.stix.i4ds.net/view/ancillary

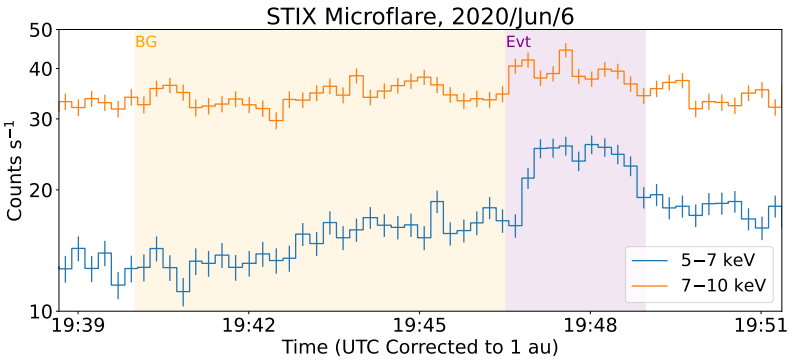

Figure 6.13: Time profiles for two SO/STIX energy ranges covering a flare on 2020 June 6. The background spectrum for spectral analysis is integrated over the orange region and the flare time is integrated over the magenta shaded region. The time axis has been corrected to 1 au.

the observed count spectrum. Using Sunxspex, the fitted parameters give a temperature of $9.4^{+0.6}_{-0.5}$ MK and emission measure of $7.0^{+4.1}_{-2.7} \times 10^{45}$ cm^{-3} (Figure 6.14).

We note that, in Figure 6.14, the STIX background dominates the emission observed from the microflare time. The background is consistent with what has been observed in previous analysis (Battaglia et al. 2021).

Fitting STIX spectra is relatively new with the first spectral fitting results being performed in OSPEX (Battaglia et al. 2021). Therefore, although the Sunxspex tools specific to the STIX instrument are not yet complete, we are able to perform spectral fitting once we obtain the relevant files without issue.

In addition, this microflare occurred during a NuSTAR solar observation campaign within the instruments $12' \times 12'$ field of view. Therefore, using Sunxspex, we can perform the first ever simultaneous fitting between STIX and NuSTAR spectra.

6.3 Multi-instrument Observations

We have shown examples of simultaneous spectral fitting of data from different detectors from the same observatory in Section 6.1.2–6.1.5, and 6.2.2. However, we may also perform simultaneous fitting with spectra from any instrument, regardless of observatory, as long as there is a significant overlap in the instrument responses; i.e., they are responsive to similar plasma temperatures.

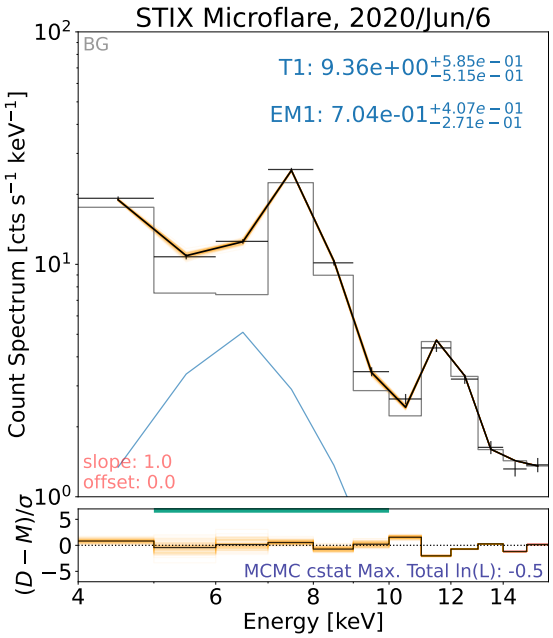

Figure 6.14: SO/STIX spectrum (black data-points with background in grey) from 2020 June 6 fitted with a single thermal model (blue line) with the performed MCMC analysis with the *maximum a posteriori* result (black) and one hundred random samples in orange. The fitted thermal parameters are displayed in blue, showing a temperature ($T1$) in units of MK and emission measure ($EM1$) of 10^{46} cm^{-3}. The green line shows the energy range the spectrum was fitted over with the residual panel showing the deviation of the fit from the data normalised by the error. Text in the residuals panel shows the best C-stat (Section 5.2.1), logarithmic value found during the MCMC analysis.

6.3.1 Simultaneous STIX and NuSTAR Spectral Fitting

As stated at the end of Section 6.2.3, we have co-temporal observations of the same microflare from both of NuSTAR's FPMs and also STIX. Figure 6.15 shows the STIX and NuSTAR FPMA & B light curves over the time of microflare SOL2020-06-06T19:47 with the flare time indicated in a magenta shaded region.

We observe that the microflare signal between 19:46:31–19:48:58 UTC is most obvious in the 5–7 keV for all three instruments (Figure 6.15, blue profiles). We perform MCMC analysis on the spectral fitting over the flare time for the three available spectra, representing all three with a single thermal model over 5–10 keV for the STIX spectrum (as in Section 6.2.3) and 5–12 keV for both NuSTAR spectra, to include the bulk of the higher energy NuSTAR counts (Figure 6.16).

From sampling the posterior we find the *maximum a posteriori* model parameters to be $10.1^{+0.2}_{-0.1}$ MK and emission measure of $8.6^{+0.6}_{-0.9}\times10^{45}$ cm^{-3}. We also include a scaling parameter C, accounting for any relative instrumental uncertainties, which is measured relative to NuSTAR's FPMA. We

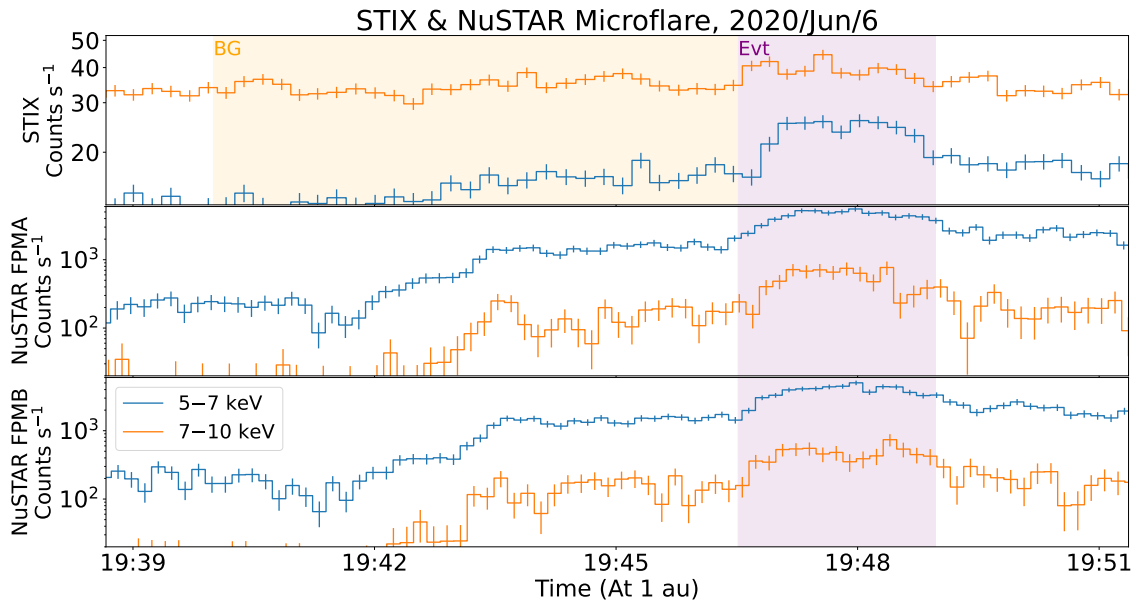

Figure 6.15: STIX (top panel, Figure 6.13) and NuSTAR FPMA & B (middle and bottom panel) time profiles for two energy ranges (5–7 keV, blue; 7–10 keV, orange) covering a flare on 2020 June 6. The background STIX spectrum is integrated over the orange region and the flare time for STIX and NuSTAR is integrated over the magenta shaded region. The STIX time axis has been corrected to 1 au.

find NuSTAR FPMB's scaling factor to be $0.95^{+0.05}_{-0.03}$ with STIX's scaling factor obtaining a value of $0.44^{+0.02}_{-0.02}$.

This is the first simultaneous fitting of NuSTAR and STIX spectra even though it is simply only utilising an isothermal model in attempting to represent the spectra. However, this demonstrates a proof-of-concept and emphasises Sunxspex's ability to perform analysis that either is difficult or not possible to perform in other spectral fitting software in their current forms. We find that, for science purposes, the spectra may require a more complicated model in order to represent the emission; for example, a multi-thermal model and/or the inclusion of a non-thermal component.

In addition, other factors will need to be considered when progressing this analysis due to this microflare being at the limits of each detector's dynamic range. The SOL2020-06-06T19:47 microflare appears as a very weak signal with STIX, only being obvious above the background in the 5–7 keV energy range in the light curve (Figure 6.13 and Figure 6.15, top panel) while being well hidden in the higher 7–10 keV energy range. The fitted spectrum (Figure 6.14) also indicates the weak nature of the 9 MK signal observed by STIX, which is why a background spectrum is used to aid its fitting. However, the microflare is incredibly bright in NuSTAR, producing livetimes down to ~0.3%. Therefore, pile-up or gain corrections must be investigated for the NuSTAR spectral fitting.

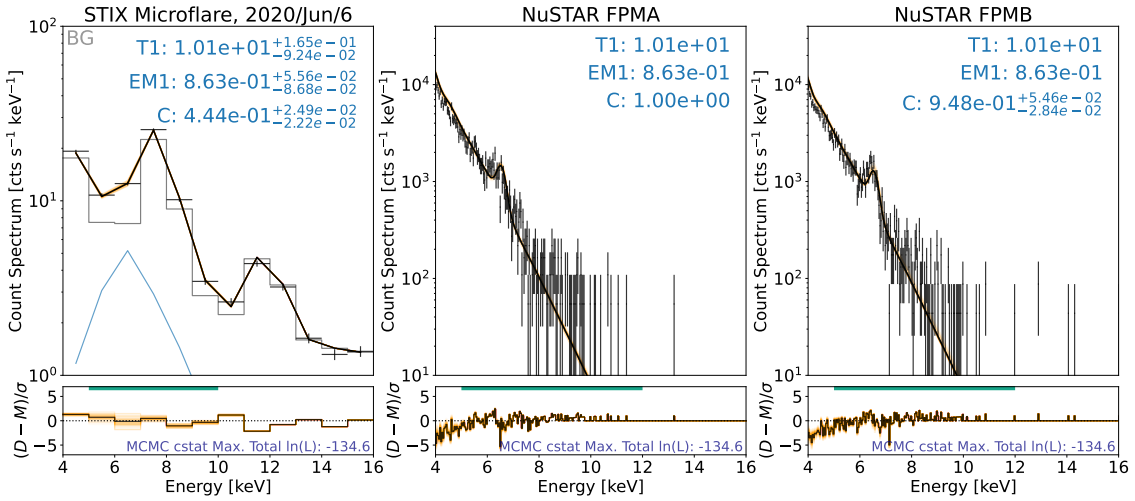

Figure 6.16: SO/STIX (black data-points with background in grey, left panel) and NuSTAR FPMA and B (black data-points, middle and right panel) spectra from 2020 June 6, simultaneously fitted with a single thermal model (black line). The performed MCMC analysis with the *maximum a posteriori* result is shown with one hundred random samples in orange. The fitted thermal parameters are displayed in blue, showing a temperature ($T1$) in units of MK and emission measure ($EM1$) of 10^{46} cm^{-3}. Note, in this example, the scalar C is measured relative to the second loaded spectrum, NuSTAR's FPMA in the middle panel.

6.4 Summary and Conclusions

Throughout this chapter, we show Sunxspex's ability to perform increasingly complex spectral analysis with excellent agreement in fitted parameter values compared to published works, and also progress that analysis where possible by being able to simultaneously fit spectra with the thick target using Poisson statistics. For example, XSPEC can simultaneously fit spectra with various fit statistics but does not have the thick target model, while OSPEX has the relevant models but assumes Gaussian statistics and cannot simultaneously fit all available data. These limitations do not exist with Sunxspex since multiple spectra can be simultaneously fitted with the thick target model for any non-thermal emission while using a variety of fit statistics or likelihoods.

We show Sunxspex's capabilities using several NuSTAR microflare spectral examples. We started by showing a fit with an isothermal model on NuSTAR's FPMA spectrum (Section 6.1.1) from microflare 3's pre-flare time taken from Section 3.4.1. We then perform the full spectral fitting process shown for microflare 10's pre-flare from Cooper et al. (2021, see Section 3.7), for both FPMs in Section 6.1.2.

We then reproduce the same fitting shown in Figure 4's left panel in Glesener et al. (2020) by fitting the spectrum with a thermal and non-thermal model (Section 6.1.3, Figure 6.4). In the same section we are able to extend the analysis beyond what was possible with OSPEX and XSPEC by fitting the

same model to both of NuSTAR's FPMs simultaneously (Figure 6.5). In addition to fitting parametric models to spectra simultaneously, we also perform the analysis in Duncan et al. (2021) for the May1618 microflare where we also fit over the gain slope response parameter.

With Sunxspex, we are able to return to the impulsive phase of microflare 3 from Section 3.4.1 and fit the non-thermal emission using the thick target model, giving electron distribution parameters instead of estimating from a correlated photon spectrum. This gave a similar fit but with a clearer non-thermal energy estimate from electron parameters (Section 1.4.5 and 4.4.2).

RHESSI data is also able to be used in Sunxspex. We show two different background handling methods (Method 3 and 5), described in Section 5.2.6, when fitting the summed 2002 October 5 RHESSI flare spectrum from Fletcher et al. (2007) (Section 6.2.1, Figure 6.10). We then show Sunxspex's ability to investigate each RHESSI detector spectrum which contributed to the summed spectrum in Section 6.2.1. This can provide an insight into each detector's relative performance to each other while also showing any abnormalities in a single spectrum which is 'smoothed out' in the summed result.

In addition to fitting NuSTAR and RHESSI data, Sunxspex also has support for SO/STIX. Therefore, we show an example of fitting a STIX spectrum in Section 6.2.3 where we find that Sunxspex is able to fit the microflare spectrum with an isothermal model to a weak signal compared to the known background (Figure 6.14). However, we note that Sunxspex could be used to fit any STIX microflare.

Finally, we also showcase Sunxspex's capabilities when fitting multiple spectra, not just those from the same observatories, seen at different locations with respect to the flare; NuSTAR being situated at Earth and STIX at ~0.5 au (Figure 6.12). In Section 6.3.1, we analyse the same microflare discussed in Section 6.2.3 but include the microflare spectra also observed by NuSTAR. Although care must be taken when performing simultaneous analysis from two completely different observatories, we show that it can easily be done in Sunxspex.

6.4.1 Future Work

Sunxspex is already a powerful tool; however, there is still further work required to improve the software. One improvement would be the ability to apply spectral model corrections when fitting. These corrections come in the form of altering the photon model being fitted either before or after being converted into count-space. For example, there is currently no simple method to incorporate pile-up corrections on the observed spectra where we account for multiple photons being detected as a single count in a detector (for examples of pile-up in relation to NuSTAR see Section 2.1 and Grefenstette et al. 2016).

Further work is needed to include specific support for other spectrometers, such as Chandrayaan-2/Solar X-ray Monitor (XSM; Vadawale et al. 2014; Shanmugam et al. 2020) which observes the Sun as a star in soft X-rays and the Fermi Gamma-Ray Burst Monitor (GBM Meegan et al. 2009) which monitors the whole sky with sensitivity to photon energies between 8 keV–40 MeV. Moreover, we aim to support any instrument that can benefit from being analysed in Sunxspex.

Sunxspex's user interface will also be improved with it being simple for any user to intuitively incorporate any Python-interacting fitting routine outside of the Sunxspex software. In addition, the default values used in plotting need to be more dynamic in order to adapt to any spectrum that is being analysed, using sensible values for all instruments and applicable from weak microflares $\ll$A-class to bright $>$X-class flares. We will also make more photon models available in Sunxspex, for example a DEM model, one which produces a multi-thermal model from a differential emission measure distribution, and include the ability for count models to be fitted directly to the data without the need to be folded through the instrument SRM. The dependence on SSW for the `CHIANTI` look-up tables used for the thermal model will also be removed and calculated entirely in Python.

It can also be difficult to locate the global maximum or minimum value throughout the parameter-space and decrease the sensitivity to the initial parameter guesses in the fitting; however, methods exist to aid in this matter. Methods like 'Differential Evolution' attempt to sample the whole of the parameter-space and uses a stochastic process to explore the landscape (Storn & Price 1997). This method spawns a number of independent vectors across the coordinate space and randomly evolves these vectors, only accepting new positions if the statistic is more desirable at the new point. The process is then terminated once a set number of new points have been attempted for every original vector. Therefore, this method can explore many dimensions quickly and is easily set-up in parallel; however, it can take many more function evaluations than other optimising methods.

In addition, some performance improvements could be made to Sunxspex to improve the speed of the fitting process. For example, many of the models started as Python copies of the OSPEX implemented IDL code. This means that the Sunxspex models are not optimised for Python and do not take advantage of many high performance libraries available. We also aim to improve performance by adding support for parallelisation throughout the computing process where independent but similar calculations are taking place, such as when performing an MCMC or other parameter exploration methods like nested sampling (Skilling 2004). Improvements made, or being made, to Sunxspex are documented and can be found on the Sunxspex GitHub issues page[9].

[9] https://github.com/sunpy/sunxspex/issues

7

Nested Sampling and Model Comparison

In this chapter, we describe the need, theory, and application of a technique called *nested sampling* (Skilling 2004, 2006) in relation to X-ray spectral fitting with NuSTAR. In Section 7.1, we discuss the need for nested sampling in solar X-ray spectral analysis with regards to model selection with the mathematical theory of the technique described in Section 7.2. In Section 7.3, we then present a NuSTAR solar microflare spectrum which could be fitted with a thermal or non-thermal model. I use the nested sampling algorithm to determine whether the thermal or non-thermal model is the most probable. The nested sampling technique allows for well evidenced model comparisons when performing spectral fitting. I have made nested sampling analysis available in Sunxspex for a given data-set and model by utilising the Python nested sampling package, Nestle[1].

7.1 The Need for Nested Sampling

Nested sampling has been used extensively in many scientific fields (Knuth et al. 2015) with many implementations (Kester & Mueller 2021; Ashton et al. 2022), but is not common in the field of solar physics. However, it is an incredibly powerful tool when performing any fitting process.

Throughout all spectral fitting examples discussed thus far we have used a fit statistic to determine the relative 'goodness-of-fit' of a model to the data. However, this approach only works when comparing different realisations of the same parametric model to the same data and only allows the determination of the best fit of parameters for that model, not necessarily if the model itself is appropriate.

Therefore, the fit statistic alone does not allow for an appropriate comparison between different models. This is, perhaps, most obvious when comparing two different parametric model fits with a different number of input parameters. It might be expected that the model with more parameters

[1] http://kylebarbary.com/nestle/

would naturally fit the data better since it has more dimensions to utilise and so it is unclear what role these extra parameters play in the 'goodness-of-fit' when comparing models.

We have used qualitative arguments in order to rule out certain model fits when observing residual artifacts indicative of ill-fitted models (e.g., microflare 3's rise time double thermal fit in Section 3.4.1) and loose quantitative arguments based on a reduced fit statistic[2] when describing the different model fits to the same data (e.g., Glesener et al. 2020).

Therefore, we utilise the nested sampling Bayesian analysis tool (Skilling 2004); a technique which integrates over the parameter posterior hyperspace in order to calculate the *evidence* term of Equation 5.13. The evidence then allows completely different parametric models over different priors to be compared, producing a quantitative value on how one model is more probable compared to another (Kass & Raftery 1995). This is invaluable and can be used to robustly determine whether, for example, non-thermal or more thermal emission is present in the X-ray spectra with reproducible values that could be used for future comparisons. Therefore, nested sampling can help establish if thermal or non-thermal processes are contributing to weak microflare spectra and if the microflares appear to be scaled versions of their larger counterparts (see Section 1.2.2). We describe the theory of nested sampling in Section 7.2.

7.2 The Theory of the Nested Sampling Algorithm

Nested sampling focusses on obtaining the evidence (the denominator of Equation 5.13) which is the probability of the data we observe occurring given a certain model, where this model is defined by its number of parameters, θ, and their priors. We adopt a new nomenclature for Equation 5.13 which avoids needless repetition of notation to give

$$P(\theta|D,M) = \frac{P(D|\theta,M)P(\theta|M)}{P(D|M)} \quad \rightarrow \quad P(\theta) = \frac{\mathcal{L}(\theta)\pi(\theta)}{Z}, \tag{7.1}$$

where everything is dependent on the given model M and the posterior, likelihood, and evidence all depend on the data D, neither of which change while fitting a data-set. The nested sampling process then uses the likelihood $\mathcal{L}(\theta)$ and the prior $\pi(\theta)$ to estimate the evidence Z and, in the process, the posterior $P(\theta)$.

[2]This divides the fit statistic by the degrees of freedom (number of data-points minus the number of parameters). However, this is a simplistic approach in order to penalise a fit that involves more parameters which becomes quickly insignificant with a large number of data-points.

7.2.1 Why Do We Want the Evidence?

In order to understand how Equation 7.1 is useful for model comparison we take a step back and investigate the probability of a hypothesis under question explaining data, D, with

$$P(H|D) = \frac{P(D|H)P(H)}{P(D)}, \tag{7.2}$$

where H is the hypothesis being tested. The hypothesis term represents the full model M in relation to Equation 7.1.

Therefore, with Equation 7.2, we would be able to compare any hypothesis with any other when it comes to explaining a given observation. However, the evidence term $P(D)$ (note this is *different* to Z in Equation 7.1)[3] is incredibly difficult to obtain, and has little meaning, since it is the probability of the observation or data occurring independent of hypothesis or model (Chapter 16 of Kurt 2019).

Fortunately, $P(D)$ does not change between hypotheses for the same data; therefore, we can compare models relatively using the posterior odds with

$$\frac{P(H_1|D)}{P(H_2|D)} = \frac{P(D|H_1)}{P(D|H_2)} \times \frac{P(H_1)}{P(H_2)}, \tag{7.3}$$

where the ratio shows which hypothesis, H_1 or H_2, is better suited to explain the data. The ratio is the product of Bayes factor (ratio of model likelihoods) and the model prior ratio with the burdensome evidence term $P(D)$ now removed (Chapter 4 of Sivia & Skilling 2006).

The model prior ratio, $P(H_1)/P(H_2)$, can be evaluated to 1 if there is no prior information as to the probability of each hypothesis; i.e., one is as likely as the other. For example, we may find it equally likely that a thermal or a non-thermal model—M_1 and M_2, respectively—could be used to fit a low count number excess in a spectral fit, having no prior preference between either hypothesis. Therefore, Equation 7.3 becomes

$$\frac{P(M_1|D)}{P(M_2|D)} = \frac{P(D|M_1)}{P(D|M_2)} = \frac{Z_1}{Z_2}, \tag{7.4}$$

where we have changed H to M in order to represent the comparison between two models and to be consistent with Equation 7.1.

This is where the power of nested sampling becomes obvious since it offers a method to calculate the model dependent evidence of Equation 7.1 ($P(D|M_i)$ or Z_i) for any model which is the likelihood term in the posterior odds ratio from Equation 7.4. Therefore, nested sampling provides the ability to

[3]The evidence Z in Equation 7.1 is equal to the likelihood $P(D|H)$ of Equation 7.2.

obtain any model dependent evidence and for the relative comparison of any model representation to another using Bayes factor (Kass & Raftery 1995).

7.2.2 Mathematically Obtaining the Evidence

Obtaining a value for the probability of the observed data—the value often omitted in MCMC analysis (e.g., shown by Equation 5.14 and 7.1)—is an incredibly important value which can then be used to compare against other model fits. We could find the evidence, Z, by integrating over all parameter dimensions for the model as

$$Z = \int_{\theta} \mathcal{L}(\theta)\pi(\theta)d\theta. \tag{7.5}$$

Therefore, using Equation 7.5, we are finding the total probability *mass* to normalise $\mathcal{L}(\theta)\pi(\theta)$ of Equation 7.1 such that

$$\int_{\theta} P(\theta)d\theta = 1, \tag{7.6}$$

which is the necessary normalisation requirement for probabilities (Chapter 9 of Sivia & Skilling 2006), this is also true for the integrated prior probability density. Again, in the case of Equation 7.1, this is all in relation to a specific model being tested against a given data-set. This integration can quickly become too complex to reliably perform, especially with increasing parameter dimensions (see Section 7.2.3).

In order to reduce the number of dimensions being integrated over, we essentially collapse all parameter dimensions down to one by calculating a quantity known as the prior mass, this is the proportion of prior with a likelihood greater than a value of λ (Sivia & Skilling 2006, pp. 182). The expression for the prior mass $> \lambda$, $\xi(\lambda)$, is given by

$$\xi(\lambda) = \int_{\mathcal{L}(\theta)>\lambda} \pi(\theta)d\theta. \tag{7.7}$$

where λ is the likelihood under question. We also note that $\lambda = 0$ at $\xi = 1$ and $\lambda = \mathcal{L}_{max}$ at $\xi = 0$ due to the probability normalisation requirements. Therefore, we find that Equation 7.5 and Equation 7.1 become

$$Z = \int_0^1 \mathcal{L}(\xi)d\xi \tag{7.8}$$

and

$$P(\xi) = \frac{\mathcal{L}(\xi)}{Z}, \tag{7.9}$$

respectively. From Equation 7.9, we observe that a random sample ξ_i taken from the area defined by Equation 7.8 provides a random sample from the posterior $P(\xi)$, equivalent to when a random sample

of θ_i provides a random sample from the posterior $P(\theta)$ shown by Equation 7.1. Therefore, we wish to calculate and integrate the distribution of likelihoods with respect to prior mass enclosed.

7.2.3 The Iterative Procedure

The integral in Equation 7.5 is incredibly difficult to numerically compute over all θ since the n_θ-dimensional space being integrated over increases exponentially with the number of parameters; therefore, a more cautious approach is needed. To do this, Skilling (2004) proposes to choose a number of *live-points* (n_{live}) that are spread randomly across $\pi(\theta)$. The likelihood $\mathcal{L}(\theta)$ is then calculated for each of these points and ordered from smallest to largest (e.g., $\mathcal{L}_{live\text{-}points} = (\mathcal{L}_*, ..., \mathcal{L}_{largest})$).

Figure 7.1 shows how we obtain the values for Equation 7.8 and 7.9 using an example model that has two parameters, θ_1 and θ_2. We are essentially iteratively creating contours from the lowest likelihood values ($\mathcal{L}_*$) of the live-points in the posterior (left panels) and then converting this multi-dimensional integral (Equation 7.5) to a 1-dimensional one over ξ (right panels). When the lowest live-point is identified, it is removed and a new live-point is added such that $\mathcal{L}_{new} > \mathcal{L}_*$. Three of these iterations are shown from the top to the bottom row in Figure 7.1.

We show the nested sampling process in Figure 7.1 where, in the right panels, Equation 7.7 is used to find the prior mass ξ above the smallest likelihood ($\lambda = \mathcal{L}_*$) and plot all $\mathcal{L}_*$ against their respective ξ_*. The worst likelihood *live-point* is then removed from the $\mathcal{L}_{live\text{-}points}$ list and replaced with $\mathcal{L}_{new}$ such that $\mathcal{L}_{new} > \mathcal{L}_*$. Shown in Figure 7.1's left panels, the list is ordered again and then the process is repeated until a stopping condition is met.

Therefore, this process not only provides an intuitive way to find the evidence, the most important term in Equation 7.1 for model comparisons, but also samples the posterior many times in order to obtain it. Thus, we obtain the same information that can be obtained by running MCMC analysis as a by-product.

7.2.4 Stopping Criteria

Generally, a certain number of iterations are performed until a manual condition is met and then the area under the curve in the final $\mathcal{L}$-ξ plot (Figure 7.1, right panels) is calculated. One stopping criterion is to estimate the remaining evidence between iterations and compare it to the evidence calculated from the previous iteration. If the difference is less than a certain threshold, the process can be terminated.

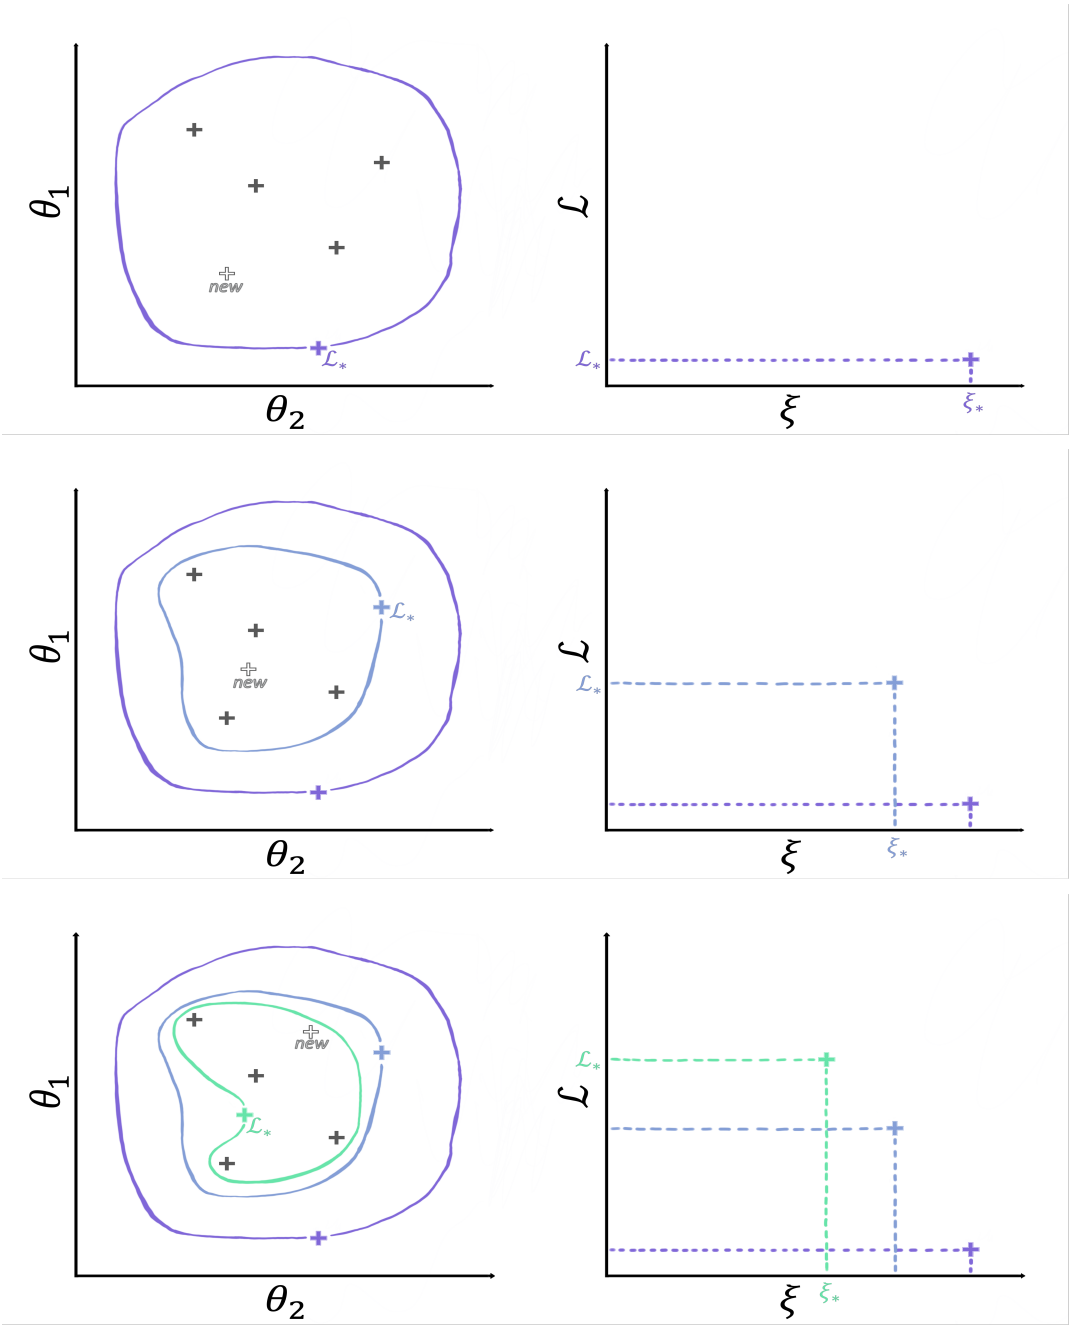

Figure 7.1: The iterative steps of the nested sampling process. The left panels show the sampling of the prior for a model with parameters θ_1 and θ_2 where the lowest likelihood value ($\mathcal{L}_*$) from the *live-points* is identified, removed, and replaced with a new live-point with $\mathcal{L}_{new} > \mathcal{L}_*$ with each row showing this iterative step. From the removed live-points we can essentially build up contours over the posterior and relate it to the prior *mass* ξ (right panels). This converts a complex multi-dimensional integral over the posterior to a simpler 1-dimensional integral of $\mathcal{L}$ over ξ when finding the value of the evidence Z (see Equation 7.8). Figure adapted from Figure 3 of Skilling (2006), Section 9.2 of Sivia & Skilling (2006), and Figure 1 of Brewer et al. (2011).

This can be done by estimating the remaining evidence with

$$Z_{remaining} = \mathcal{L}_{max,i}\xi_i. \tag{7.10}$$

where $\mathcal{L}_{max,i}$ is the maximum likelihood value of the live-points and ξ_i is the remaining prior mass at iteration i.

The prior mass at each iteration ξ_i is obtained by initially drawing samples over the original prior ($\xi_0 = 1$) and then compressing the mass upon each iteration with a value of t_i. The compression is not estimated through geometry but probability with

$$ln(t) = -\frac{1}{n_{live}}, \tag{7.11}$$

which describes an exponentially decreasing prior mass for each iteration based on the number of live-points (Ashton et al. 2022). The prior mass ξ_i for each iteration i then becomes

$$ln(\xi_i) = \sum^{i} -\frac{1}{n_{live}} = -\frac{i}{n_{live}}, \tag{7.12}$$

which describes the prior mass shrinking by 1-part-in-n_{live} for every iteration (Section 9.2.1 of Sivia & Skilling 2006).

We can then use the estimated gain of evidence and define the stopping criterion to be

$$ln(Z_{remaining} + Z_i) - ln(Z_i) < A_{threshold} \tag{7.13}$$

where $A_{threshold}$ takes a customisable value. This is an approach taken by the Python nested sampling package, Nestle[4].

7.2.5 Error in the Evidence

Nested sampling also can provide an error on the final evidence result. We do this by obtaining an expression for the *information* gained throughout the iteration steps. The information, $\mathcal{H}$, is given by

$$\mathcal{H} = \int \frac{\mathcal{L}(\xi)}{Z} ln\left(\frac{\mathcal{L}(\xi)}{Z}\right) d\xi \approx \sum_i \frac{\mathcal{L}_i(\xi_{i-1} - \xi_i)}{Z} ln\left(\frac{\mathcal{L}_i}{Z}\right) \tag{7.14}$$

[4] http://kylebarbary.com/nestle/stopping.html

where the rightmost expression is used for discrete data (Sivia & Skilling 2006, pp. 186). This expression sums over the normalised evidence multiplied by a signal-to-noise, gain-like term for each iteration. Equation 7.14 gives an estimate of how much we have 'learned' from the prior.

For example, if we obtain large steps in $\delta\xi$ and large relative increases in $\mathcal{L}$, then $\mathcal{H}$ will be large, showing that we have 'compressed' or 'shrunk' the prior-space by focussing in on significant values in the posterior. If we only obtain small $\delta\xi$ and relative $\mathcal{L}$ values then $\mathcal{H}$ will be small, showing that there are no compelling areas within the prior, learning little from the analysis.

Skilling (2004) shows that the error on the final computed evidence is related to $\mathcal{H}$ and also the number of live-points used. This uncertainty in $ln(Z)$ is given by $\sqrt{\mathcal{H}/n_{live}}$, meaning that not only can nested sampling allow models to be compared to ascertain which is the most likely to represent the data, but it also gives a confidence estimate of that result.

7.3 Application to Solar X-ray Spectra

Returning to the situation described in Section 7.2.1 of spectral fitting, we can utilise nested sampling analysis to determine whether higher energy NuSTAR counts are more likely to be produced from a thermal source or a non-thermal source. This is especially crucial when the simpler arguments discussed in Section 7.1 are not useful.

We present a microflare that falls into this exact scenario. We first discuss the initial analysis that indicates nested sampling is needed to determine whether an observable count spectrum excess above an isothermal representation is more likely to be caused by thermal or non-thermal emission. We then discuss the results of the nested sampling, concluding on the more probable fit.

7.3.1 Microflare SOL2021-11-17T21:14

We investigate a microflare that occurred on 2021 November 17 at 21:14 UTC and observed with NuSTAR (for more details on this campaign, see Section 8.2). Figure 7.2 shows the NuSTAR FPMA and B X-ray time profiles of the microflare with the microflare time (21:12:30–21:14:50 UTC) indicated with a blue shaded region and a pre-flare time (21:09:40–21:12:00 UTC) in green. The microflare time was based on the 5–10 keV emission.

From Figure 7.2, we observe that the higher energy time profile (5–10 keV) is more impulsive and peaks earlier than the 3–5 keV time profile. This is indicative of hotter plasma early in the flaring process and/or a sign of non-thermal emission (Cooper et al. 2020).

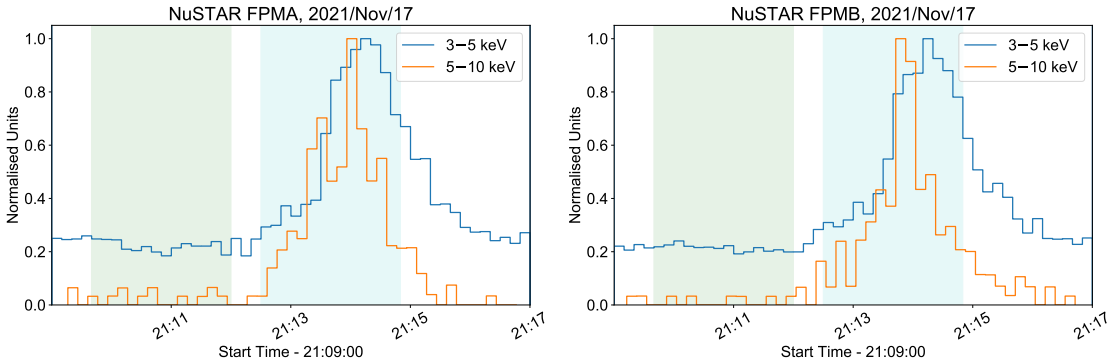

Figure 7.2: NuSTAR FPMA (left panel) and FPMB (right panel) time profiles of the 21:14 UTC microflare on 2021 November 17 over 3–5 keV (blue) and 5–10 keV (orange). The pre-flare (21:09:40–21:12:00 UTC) and microflare (21:12:30–21:14:50 UTC) time used for spectral analysis is shown by a green and blue shaded region, respectively.

Performing a spectral fit and MCMC analysis on the spectra obtained from FPMA and B over 21:09:40–21:12:00 UTC, we find a typical pre-flare temperature of 3–4 MK as shown in Section 8.2.2. We then adopt method 1 from Section 5.2.6 when incorporating the pre-flare model fit into the microflare time analysis. We fit a thermal model during the microflare time (blue shaded region in Figure 7.2), the results of which are shown in Figure 7.3.

The thermal fit provides a temperature and emission measure of 6.6 MK and 1.7×10^{44} cm^{-3}, respectively; fit parameters not too dissimilar from those obtained from other NuSTAR microflares discussed in Chapter 3. However, from Figure 7.3 we see that there is a clear excess $>$6 keV with residual artifacts from the fit $<$6 keV. This suggests an additional spectral model is required to represent the higher energy emission in the 5–10 keV range, corroborating what is eluded to by the different energy time profiles in Figure 7.2.

Therefore, we fit the spectrum with an additional model, either thermal or non-thermal, compared to what is shown in Figure 7.3. We observe that in both scenarios, the residual artifacts are removed and the excess is fitted. The non-thermal addition only just has a more preferable fit statistic but also has an additional parameter compared to the thermal addition (five parameters compared to four). It is difficult to distinguish between the models because there are so few counts in the excess, even though it is visible in Figure 7.3. This is shown and described along with a recap of this microflare when discussing the NuSTAR observation campaign it originates from in Section 8.2.2.

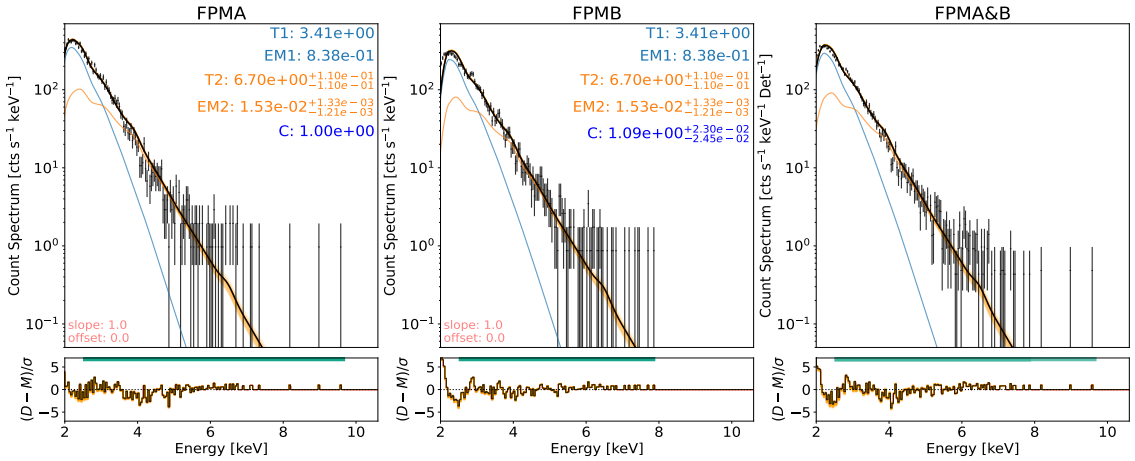

Figure 7.3: NuSTAR FPMA (left panel) and FPMB (middle panel) spectra (black data-points) of the 21:14 UTC microflare on 2021 November 17 fitted with a single thermal model in addition to a fixed pre-flare component (black line, left panel). The performed MCMC analysis with the most probable fit result (the *maximum a posteriori*, black) and one hundred random samples (orange) are plotted. The fitted thermal parameters are displayed in blue and orange, showing temperatures ($T1$ and $T2$) in units of MK and emission measures ($EM1$ and $EM2$) in 10^{46} cm^{-3}. The mean fit is shown in the right panel.

7.3.2 Thermal or Non-thermal?

We perform nested sampling on the spectral fitting of microflare SOL2021-11-17T21:14 (Figure 7.3), one with a thermal model fitting the excess counts >6 keV over the two thermal models (Figure 7.4) and one with the non-thermal model fitting the excess (Figure 7.5). The null hypothesis is also investigated, that an extra model is not needed, and so we perform the nested sampling analysis on the model fit shown in Figure 7.3. Although it appears obvious an additional model is required beyond the pre-flare and isothermal models as shown in Figure 7.3, we do not find it visually obvious whether a thermal or non-thermal model is required when comparing the fits shown in Figure 7.4 and 7.5, respectively. Therefore, we obtain the numerical result from nested sampling analysis in order to compare the different fits.

We choose 512 live-points, an $A_{threshold}$ (from Equation 7.13) value of 0.1, and use the Python nested sampling package, Nestle. The systematic C factor fitted in Figure 7.3 is fixed to the value shown when performing the nested sampling analysis and conservative ranges for the uniform prior distributions are chosen.

We find logarithmic evidence values ($ln(Z)$, with Z being the same one in Equation 7.4) of -188.5, -166.3, and -160.5 for no excess model, thermal excess model, and non-thermal excess model, respectively (Table 7.1). Therefore, the model that uses the tick-target has the largest evidence value and is most likely to represent the observation. In addition, we can determine how confident we should be in this

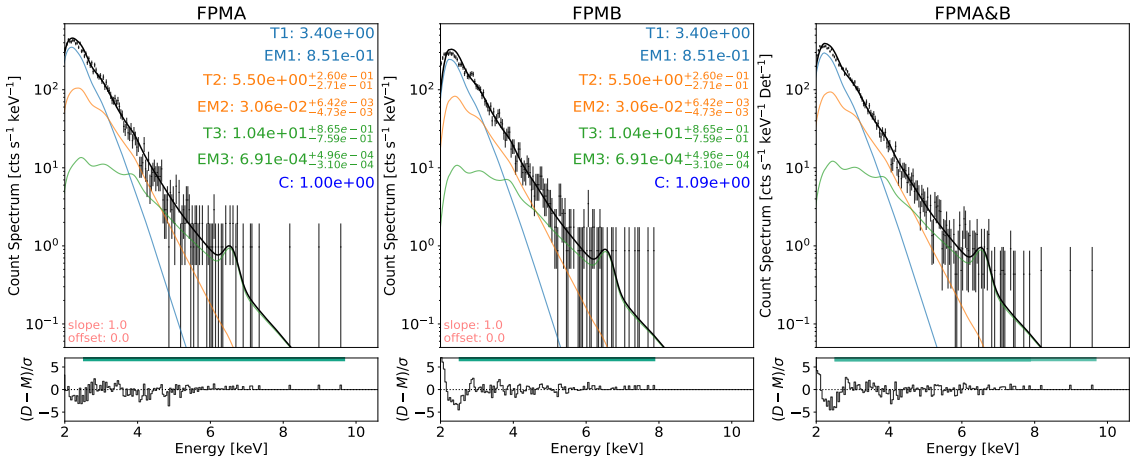

Figure 7.4: NuSTAR FPMA (left panel) and FPMB (middle panel) spectra (black data-points) of the 21:14 UTC microflare on 2021 November 17 also shown in Figure 7.3. The resulting fit is from the nested sampling result in Sunxspex (see text) where we fit the higher energy count excess >6 keV, seen in Figure 7.3, with an additional isothermal model. The parameters are shown in the same colour as their respective models with temperatures ($T1$, $T2$, and $T3$) in units of MK and emission measures ($EM1$, $EM2$, and $EM3$) in 10^{46} cm^{-3}. The mean fit is shown in the right panel.

Table 7.1: The logarithmic evidence values (Equation 7.8) calculated from nested sampling analysis for the 'pre-flare+thermal', 'pre-flare+thermal+thermal', and 'pre-flare+thermal+non-thermal' model fits to the 21:14 UTC NuSTAR microflare on 2021 November 17.

	pre-flare+isothermal ($ln(Z_{null})$)	pre-flare+2 isothermals ($ln(Z_{thermal})$)	pre-flare+isothermal+non-thermal ($ln(Z_{non-thermal})$)
$ln(Z)$	-188.5±0.1	-166.3±0.1	-160.5±0.1

result using the criteria set out in Kass & Raftery (1995) where a logarithmic Bayes factor (Equation 7.4) of 0–1 gives no confidence at all of preferring M_1 over M_2, and 1–3, 3–5, and >5 gives positive, strong, and very strong evidence, respectively.

Comparing the evidence values, the ln(Bayes factor) values—calculated from taking the logarithm of Equation 7.4—are 22.3, 28.0 for the thermal and non-thermal additions compared to the no addition case (Table 7.2). Therefore, we have a very strong reason to be confident that an additional model is needed to what is shown in Figure 7.3. We then compare the thermal and non-thermal hypotheses and find an ln(Bayes factor) value of 5.7 (Table 7.2). Therefore, there is very strong evidence that the non-thermal addition should be preferred over the thermal addition, with the total 'pre-flare+thermal+non-thermal' model being approximately 300 times more likely than the 'pre-flare+thermal+thermal' model and over 10^{12} times more likely than the initial 'pre-flare+thermal' representation shown in Figure 7.3. Figure 7.5 shows the resulting preferred model fit to the microflare spectrum resulting in a non-thermal power of $\sim 2\times10^{24}$ erg s^{-1} using Equation 1.31.

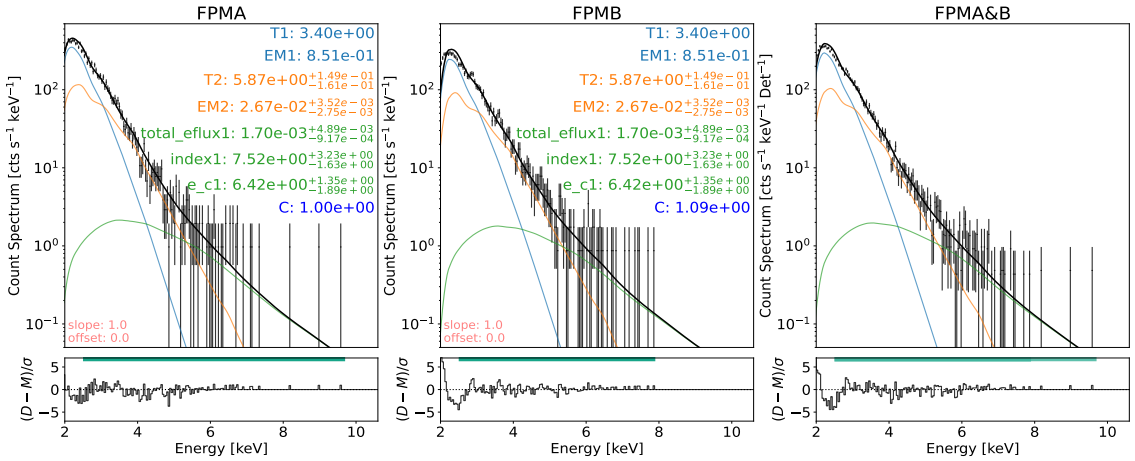

Figure 7.5: NuSTAR FPMA (left panel) and FPMB (middle panel) spectra (black data-points) of the 21:14 UTC microflare on 2021 November 17 also shown in Figure 7.3. The resulting fit is from the nested sampling result in Sunxspex (see text) which indicates that the preferred model to represent the data is one with the pre-flare (blue), an isothermal (orange), and a non-thermal (green) component. The parameters are shown in the same colour as their respective models with temperatures ($T1$ and $T2$) in units of MK and emission measures ($EM1$ and $EM2$) in 10^{46} cm^{-3}. The electron flux and the low-energy cut-off being in units of 10^{35} e^{-} s^{-1} and keV, respectively, with the electron spectral index being dimensionless. Note the residual excess >6 keV has been removed. The mean fit is shown in the right panel.

Table 7.2: The logarithmic Bayes factor values (Equation 7.4) comparing the 'pre-flare+thermal', 'pre-flare+thermal+thermal', and 'pre-flare+thermal+non-thermal' model fits to each other. These values show, in powers of e, how much more probable one model is compared to another.

	$ln(Z_{thermal}/Z_{null})$	$ln(Z_{non-thermal}/Z_{null})$	$ln(Z_{non-thermal}/Z_{thermal})$
ln(Bayes factor)	22.3±0.2	28.0±0.2	5.7±0.2

From this analysis we are able to create a corner plot as we would obtain from MCMC analysis. We calculate weights since the samples are not originally drawn from the posterior to simulate which 'steps' would have been accepted or declined throughout MCMC analysis, although some nested sampling programs return the posterior samples by default[5]. The weight for a posterior sample is proportional to the associated prior mass width multiplied by the likelihood (Skilling 2006)[6].

The corner plots for the null, thermal, and non-thermal scenarios are shown in Figure 7.6, 7.7, and 7.8, respectively. This gives a visual representation of the posterior and the more probable parameter values of the fitted model. The density distributions created from the nested sampling result are found to be consistent with those found from separate MCMC analysis.

Therefore, nested sampling is a powerful tool which enables us to quantitatively establish the most

[5] http://mattpitkin.github.io/samplers-demo/pages/samplers-samplers-everywhere/
[6] The Nestle Python package calculates the posterior sample weights as $\mathcal{L}_i \delta\xi_i / Z$, with $\delta\xi_i = \xi_{i-1} - \xi_i$.

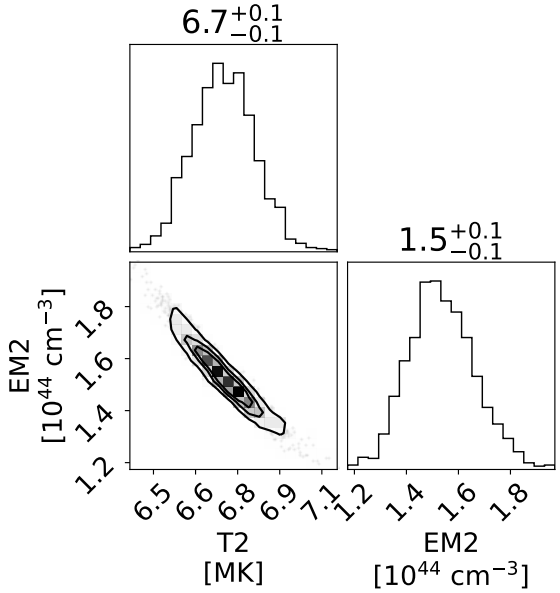

Figure 7.6: Corner plot of the nested sampling analysis performed on the null hypothesis (no extra model beyond the pre-flare and isothermal model) shown in Figure 7.3 in Sunxspex. The distributions in the off-diagonal panels show the walker density and how the emission measure ($EM2$) varies with temperature ($T2$) during the fitting, while the diagonal panels show the marginalised walker density distributions of each parameter. The most probable values, with the 68% uncertainty range, is shown above each diagonal panel.

probable model representation of any given spectra. In this section, we have shown it can help in the search for non-thermal or very hot temperature signatures in weak flares; allowing us to confidently explore and determine the physics involved at these microflare scales.

7.4 Summary and Conclusions

We discuss the need for nested sampling in any fitting process as it allows us to quantify the probability of a model in relation to data in a way that it can be compared to any other model representation. Therefore, we can gain a degree of measured confidence in concluding on a specific model fit without relying on weaker qualitative and quantitative arguments.

The mathematical theory of the nested sampling process is presented in Section 7.2, converting an n-dimensional integral that is difficult to adequately evaluate into a simpler 1-dimensional integral. A visual example of this, for a model with two parameters, is shown in Figure 7.1.

Nested sampling analysis is used with microflare SOL2021-11-17T21:14 via Sunxspex, a spectral example with a visible count excess at high energies above an isothermal model fit. We calculate that

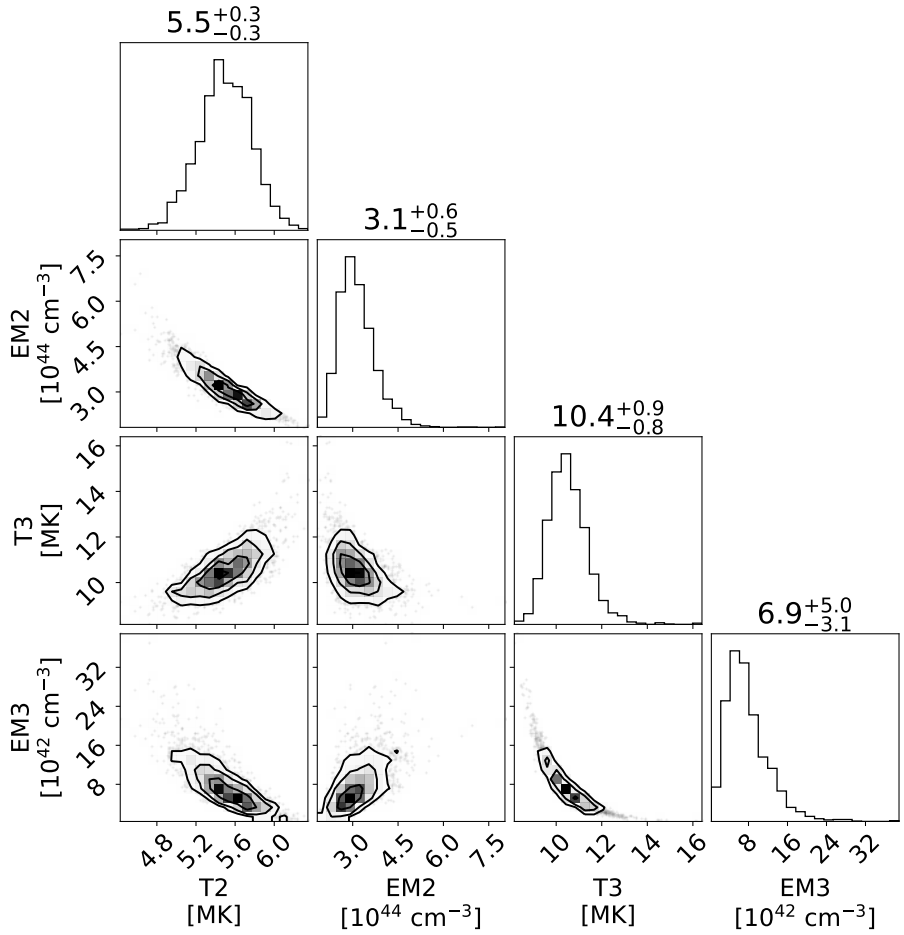

Figure 7.7: Same format as Figure 7.6 but with the spectrum being fitted with a pre-flare fixed component and two isothermal models as in Figure 7.4. We note that the walker density distribution of 'EM3' appears to be cut off at the lower bound of the prior; however, this is not the case and is due to the uniform histogram binning and the density distribution spanning multiple orders of magnitude.

there is very strong evidence an additional model beyond an isothermal and pre-flare component is needed to represent the spectrum shown in Figure 7.3. We then conclude that the most probable addition is a non-thermal model rather than a thermal model using pre-defined confidence criteria (Kass & Raftery 1995). This is the first example of nested sampling being used in relation to solar X-ray microflare spectra.

Therefore, Sunxspex and nested sampling has allowed us to confidently conclude the preferred model that best represents the microflare spectrum where previous techniques would have provided little certainty. We find that microflare SOL2021-11-17T21:14 is the weakest non-thermal AR microflare observed in hard X-rays with a non-thermal power of $\sim 5 \times 10^{24}$ erg s^{-1}. However, we note that the non-thermal model is only preferred in relation to the other models tested (the 'pre-flare+thermal',

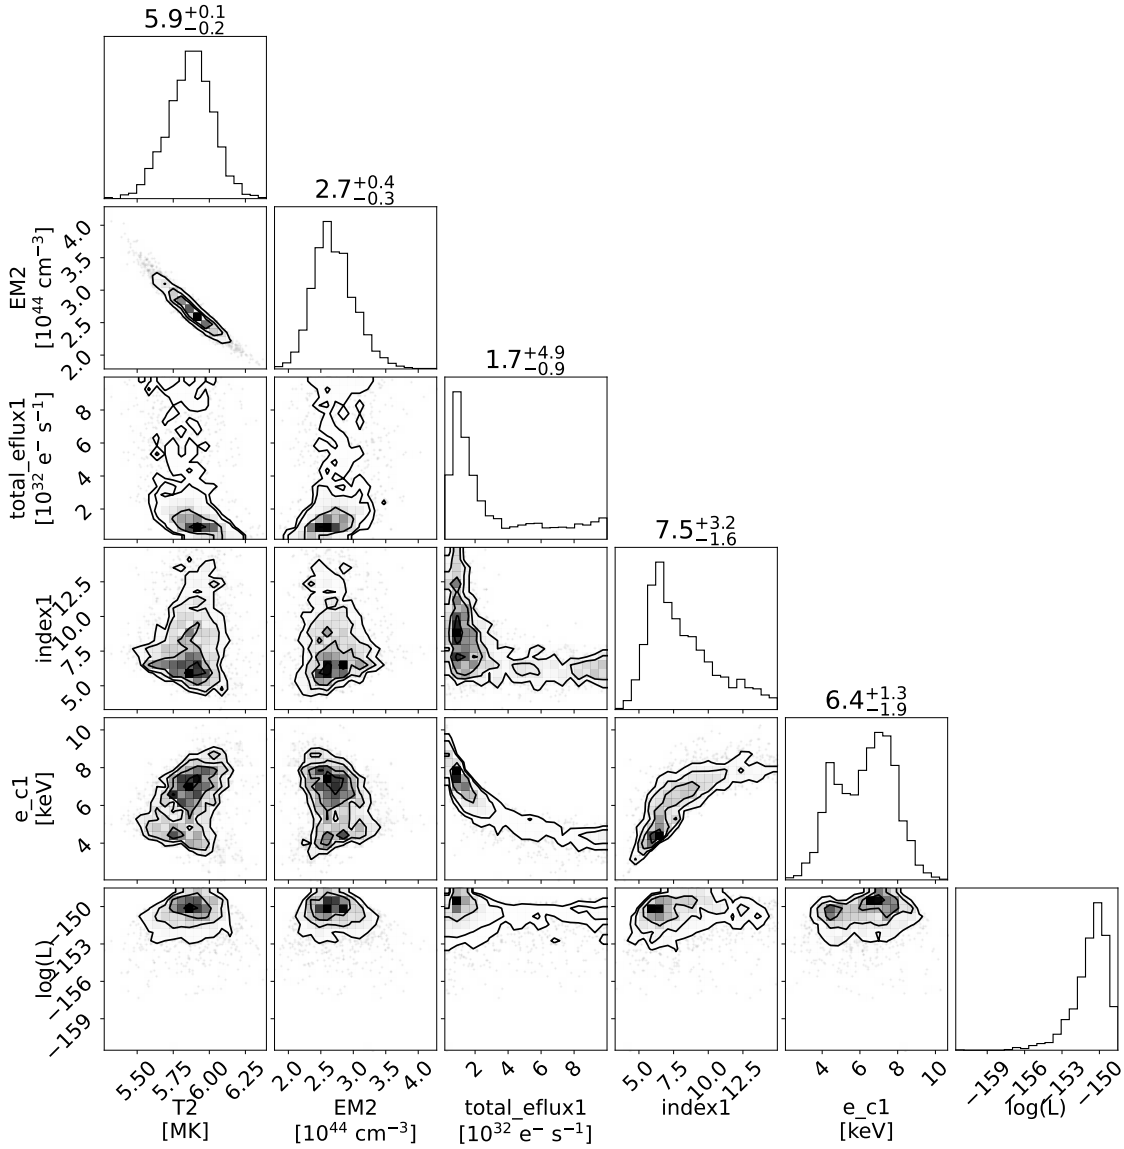

Figure 7.8: Same format as Figure 7.6 but with the spectrum being fitted with a pre-flare fixed component, one isothermal model, and a non-thermal model as in Figure 7.5. We note that the walker density distribution of 'total_eflux1' appears to be cut off at the lower bound of the prior; however, this is not the case and is due to the uniform histogram binning and the density distribution spanning multiple orders of magnitude.

'pre-flare+2 thermal' cases); therefore, another model that has not been tested could be found to be more probable to represent the spectra.

7.4.1 Future Work

Nested sampling is a tool proposed in 2004 (Skilling 2004) and we have implemented the algorithm via the Python package Nestle, but there are numerous implementations of the technique. For example, a situation may arise where it is exceedingly difficult to locate the majority volume in the posterior; therefore, *diffusive* nested sampling may be beneficial (Brewer et al. 2011). Diffusive nested sampling gives the ability for new live-points being introduced to search outside the contours being created after each iteration (contours shown in Figure 7.1, left panels). This can help stop the true most probable likelihood volume being erroneously excluded by an iteration without a chance of being located.

Sunxspex currently has the ability to set up the model parameters with a likelihood and prior, run the process, and store the nested sampling result; however, it does not have a method relating to the display of these results. Future work will include providing Sunxspex with more supporting functions for the outputs of any nested sampling approach. Further work will also involve applying the nested sampling algorithm to any event where there is any ambiguity in the appropriate model being fitted to the spectra.

8

A Further Collection of Solar Impulsive Energy Releases Observed with NuSTAR

In Section 8.1 and 8.2 we present research into two solar observation campaigns performed in 2020 January and 2021 November, respectively. Two microflares from the same general structure occurred on 2020 January 30 approximately 4.5 h apart. The first microflare is very spatially complicated, producing a complex time profile, while the second has a relatively simple physical structure and time evolution in comparison.

The final solar observation campaign we discuss, performed on 2021 November 17–22, provided a rich plethora of almost twenty separate events. From this NuSTAR observation, we easily identify twelve X-ray microflares and five other very weak events still to be identified, with potentially more events to be found through more in-depth analysis.

Both of these solar observation campaigns provide more examples of X-ray impulsive energy releases. These events can then be investigated to assess the temperatures reached and the possibility of non-thermal emission. The evolution and behaviour of these weak X-ray microflares can then be compared to what is observed in larger flares (see Section 1.2.2).

8.1 NuSTAR Solar Observation Campaign: 2020 January 30

On 2020 January 30, NuSTAR observed the Sun over many hours. During this time AR SPoCA 23783 was observed by NuSTAR and produced two weak X-ray microflares occurring on consecutive orbits where both have complex geometry. Figure 8.1 shows microflare 1 (SOL2020-01-30T13:40) and 8.2 shows microflare 2 (SOL2020-01-30T18:15). In this Section, we describe these two microflares utilising NuSTAR and SDO/AIA observations.

8.1.1 Temporal and Spatial Overview

We see in the NuSTAR contour plots that the X-ray emission traces out a corroboratory structure to that seen in Fe XVIII (Figure 8.1 and 8.2, bottom panels). The NuSTAR contours for the first microflare (Figure 8.1) also indicates that different aspects of the microflaring loops are visible in the two different energy ranges shown.

We split both microflares into several time ranges of ~2 min and ~4 min intervals for microflare 1 and 2, respectively. We do this in an attempt to isolate particular flare phases and time profile features, such as the impulsive rises and the peaks. We are also able to obtain images of NuSTAR contours overlaid on the average Fe XVIII images for the same time range. We apply the same shift to both X-ray energy ranges in each time; however, the same shift was not applied to all times due to constant CHU state shifts which produces pointing differences throughout the microflare evolution (see Section 2.1.1).

The contour image panels of Figure 8.1 and 8.2 show the co-temporal emission from NuSTAR and the SDO/AIA Fe XVIII proxy channel (Section 2.5.1) at the time of peak microflare emission. Throughout the evolution of each microflare both channels seem to respond to similarly shaped structures suggesting the same location is producing both the EUV and X-ray emission. Note that the 50%, 90%, and 99% contour levels are of the peak emission in each time step for both energy ranges to show the similarities in structures observed in by each instrument; therefore, changes in contour sizes from panel-to-panel do not represent a true change in geometry.

The first microflare (SOL2020-01-30T13:40) initially shows activity in the top right loop feature before the strong emission propagates to the more centrally located vertical loop. The second microflare (SOL2020-01-30T18:15) does not show much dynamic movement but still corroboratory EUV and X-ray emission mostly located at the loop tops of the event. Although both microflares display complex geometries, microflare 1 is more complicated spatially and temporally.

We also find that microflare 1 appears to be a complex interconnected combination of small loop structures, each ~20″ in length, that are continuously microflaring. However, microflare 2 is a relatively simpler structure with the multiple loop structures running near parallel to one another.

The two microflares appear across all coronal SDO/AIA channels (Figure 8.3 and 8.4 for microflare 1 and 2, respectively) and clearly dominate NuSTAR's FOV. These events are brighter and physically larger to all those in the 2018 September data (Chapter 3). For each microflare, the emission appears to be well positioned within the FOV meaning that pointing variations do not pose a significant issue as long as care is taken when defining regions of NuSTAR emission for different times.

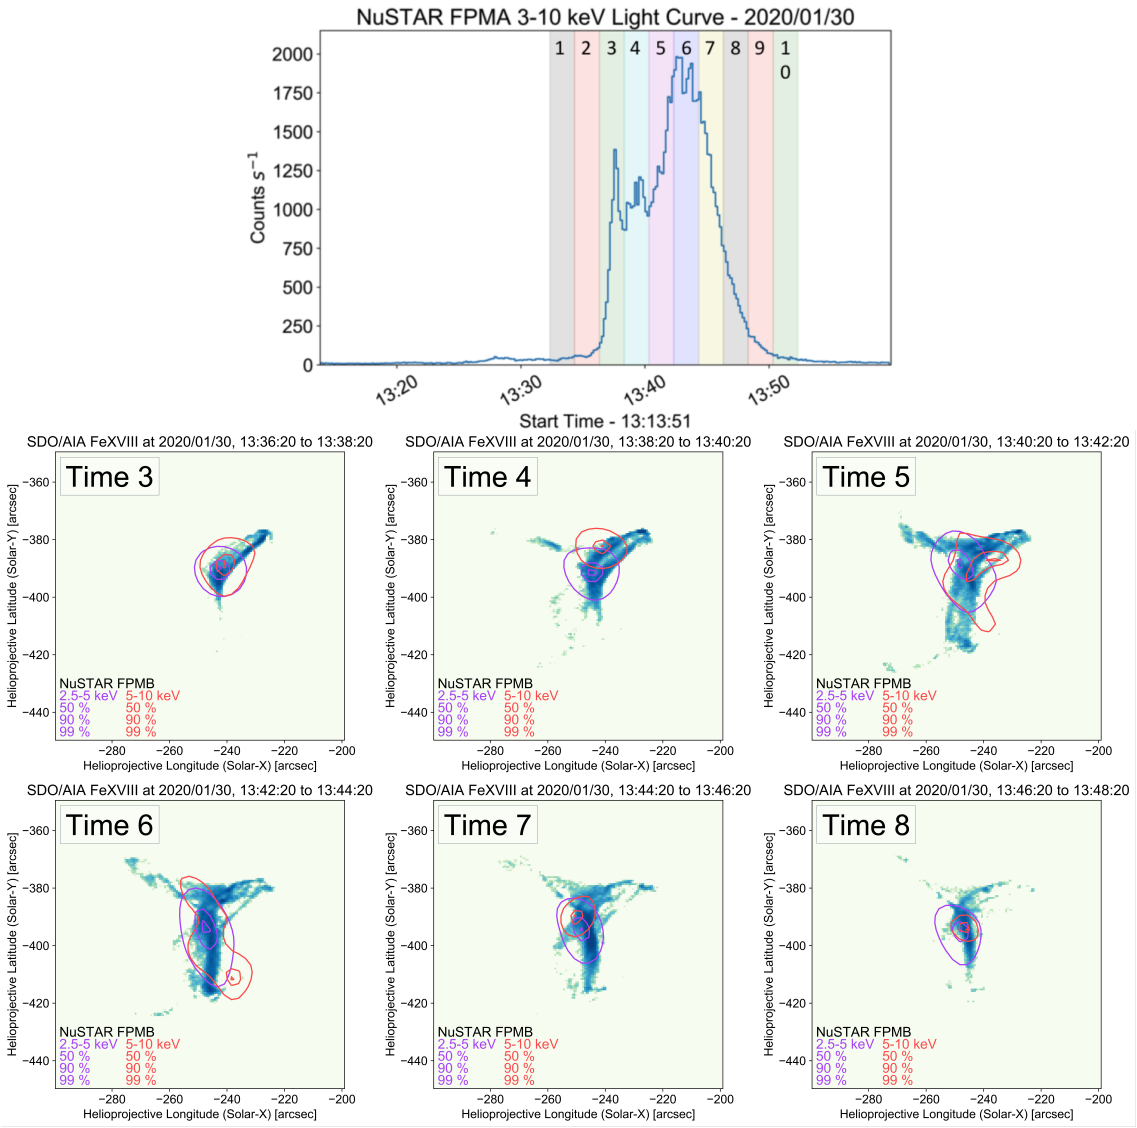

Figure 8.1: NuSTAR 3–10 keV light curve (top panel) showing the microflare 1 time profile broken into ten time ranges. SDO/AIA Fe XVIII time averaged maps of the microflare occurring on 2020 January 30 from AR SPoCA 23783 across the main microflare times (bottom panels). NuSTAR 2.5–5 keV (purple) and 5–10 keV (red) emission contours are shown for the same time ranges on top of the Fe XVIII emission. Contour levels are at 50%, 90%, and 99% of each energy range's emission at each time step. The same shift has been applied to both NuSTAR energy ranges while the images of the 2.5–5 keV and 5–10 keV emission are deconvolved with 50 and 25 iterations, respectively. All NuSTAR emission is livetime corrected with livetimes ranging between 4%–44% throughout the microflare's evolution.

Future work will investigate the time and spatial EUV evolution in more detail. In addition, soft X-ray emission will also be studied further with GOES/XRS (Section 2.2) and Hinode's X-ray Telescope (Hinode/XRT; Kosugi et al. 2007; Golub et al. 2007).

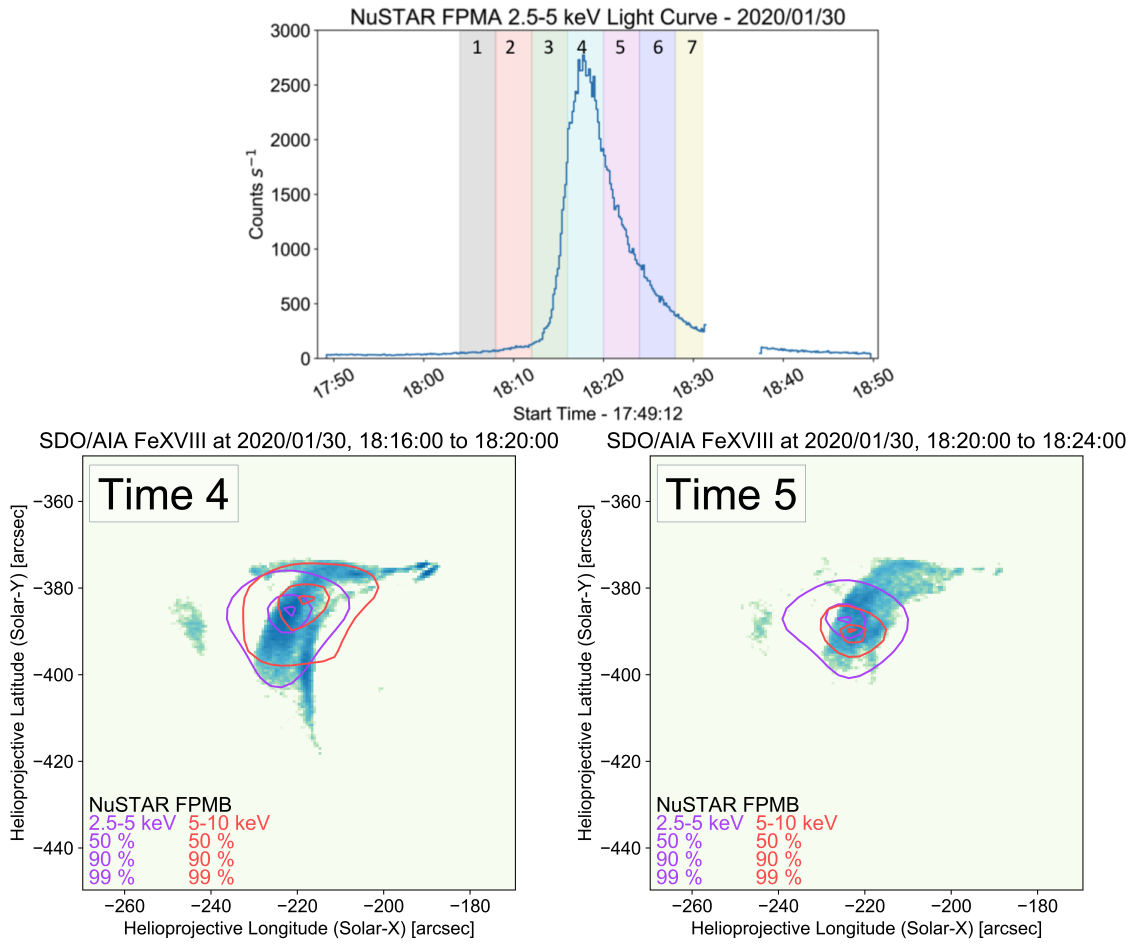

Figure 8.2: The same format as was used for microflare 1 in Figure 8.1 but for microflare 2. NuSTAR livetimes range between 6%–36% throughout the microflare's evolution.

8.1.2 Spectral Analysis and Nature of Emission

Investigating microflare 1's spectra for all ten time ranges we find strong evidence that time 2 (a pre-flare time) is multi-thermal, similar to the multi-thermal pre-flare of microflare 10 in Section 3.7, as shown by Figure 8.5 (top left panel). Time range 1 is omitted as it produces very similar result to time 2. Spectral analysis in XSPEC (Section 5.3.2) shows typical quiescent AR temperatures of ~4 MK with an excess potentially due to plasma with a temperature ~8 MK and a very small emission measure (~10^{43} cm^{-3}) or non-thermal processes. All spectra in this section are integrated over a circular region of radius ~70″ centred on the brightest emission.

Shown in Figure 8.5, the first rise (time 3, top middle panel) produces a much more complicated spectrum with a clear 6.7 keV Fe complex (Section 1.4.2). We find this line complex located at 6.4 keV

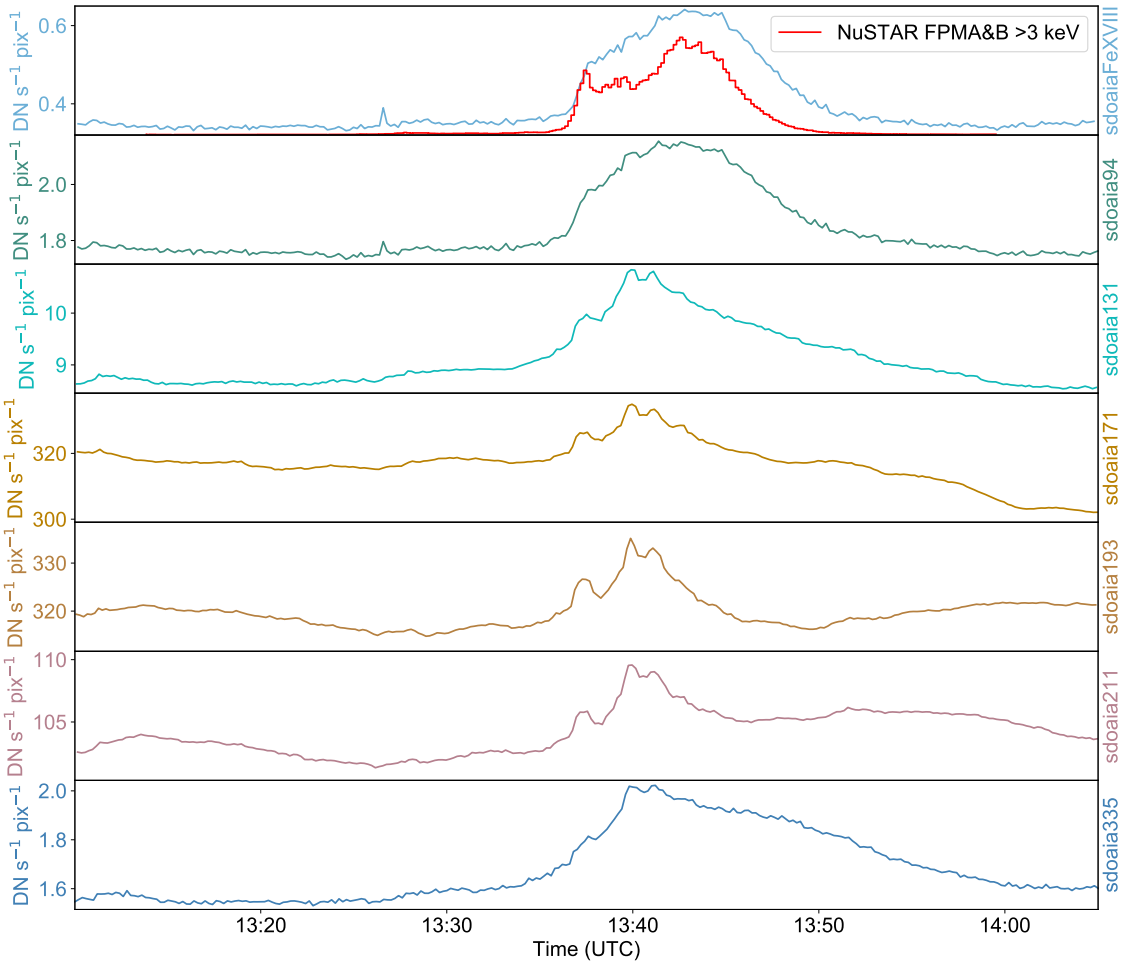

Figure 8.3: SDO/AIA AR integrated time profiles of NuSTAR's microflare 1 on 2020 January 20. We show the Fe XVIII proxy and SDO/AIA 94 Å, 131 Å, 171 Å, 193 Å, 211 Å, and 335 Å channel time profiles from the top to the bottom panel. The NuSTAR FPMA+B >3 keV light curve is shown in red in the top panel for context. The spike present at ~13:26 in the SDO/AIA Fe XVIII proxy and 94 Å time profiles appears to be due to a cosmic ray strike.

which is resolved by the approach described in Duncan et al. (2021), utilising an instrumental gain correction to NuSTAR's spectral response matrix (Section 2.1.5) due to the clear shift in the line's position. The strong presence of the line feature suggests the presence of hot plasma.

However, we do not find an acceptable representation when time 3's spectrum is fitted with two thermal models. We do, however, find a satisfactory fit (no obvious residual artifacts indicative of an ill-fit) when fitting the spectrum >5 keV with a thermal APEC model and a broken photon power-law model to represent non-thermal emission in XSPEC. The thermal model is fitted over the line emission (5–7.5 keV) and the non-thermal model represents the counts >7.5 keV. We find this fit acceptable to

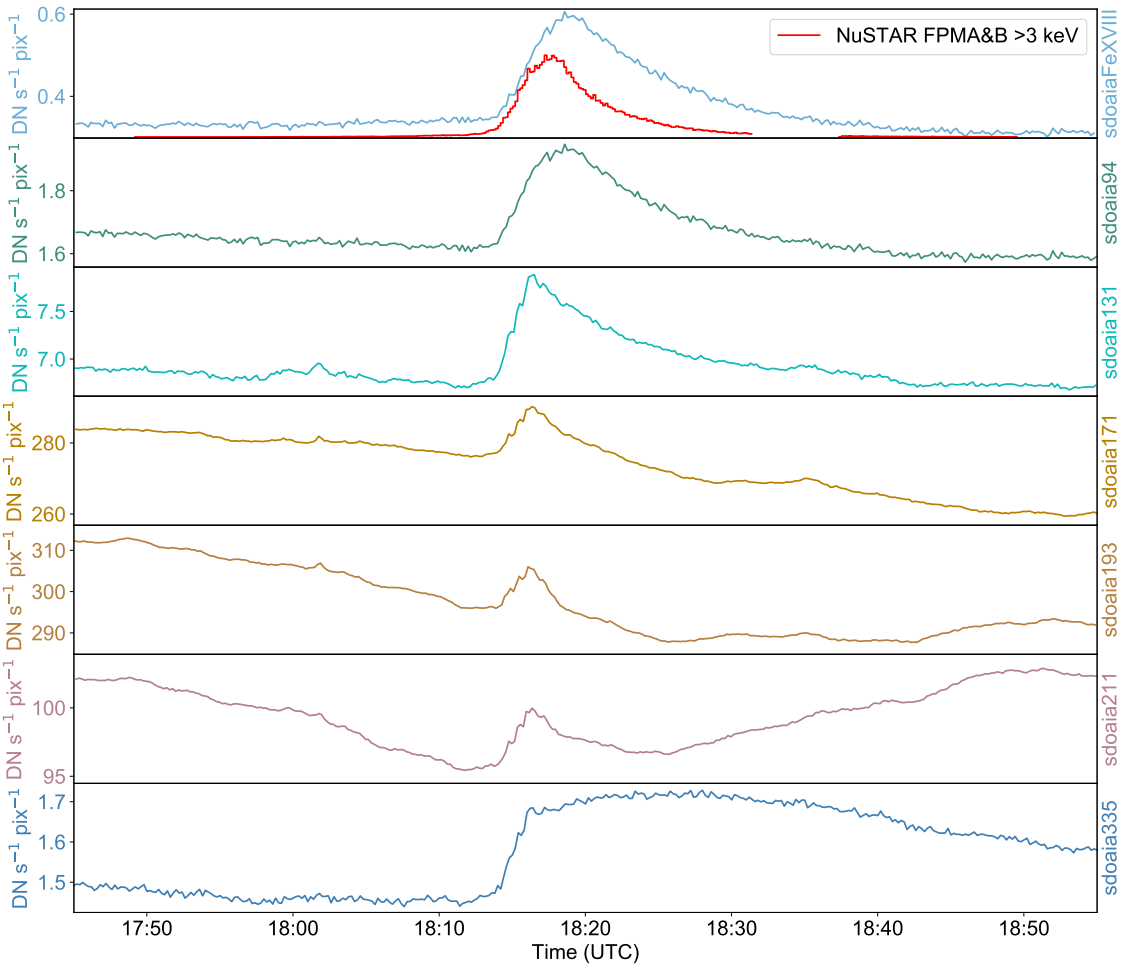

Figure 8.4: The same format as was used for microflare 1 in Figure 8.3 but for microflare 2.

energies lower than the fit range, down to 4 keV (Figure 8.5, top middle panel), with lower energy emission being assumed to be progressively dominated by enhanced pre-flare emission. The pre-flare was omitted in the fit due to its complex nature and low significance to the energy ranges of interest.

Taking the same approach for microflare 1's time 4 (Figure 8.5, top right panel), we find a slightly higher temperature and emission measure compared to time 3 (top middle panel). The non-thermal parameters show a very similar break energy and photon index with an increased normalisation constant. For the second rise at time 5 (Figure 8.5, bottom left panel) we find mostly multi-thermal emission; however, there appears to be a slight excess >7 keV which could indicate non-thermal emission. Time 6 (Figure 8.5, bottom middle panel) shows a decrease in the hotter thermal model's temperature but a large increase in emission measure from time 5.

We group times 7 to 10 together (Figure 8.5, bottom right panel) since the images in Figure 8.1 and the

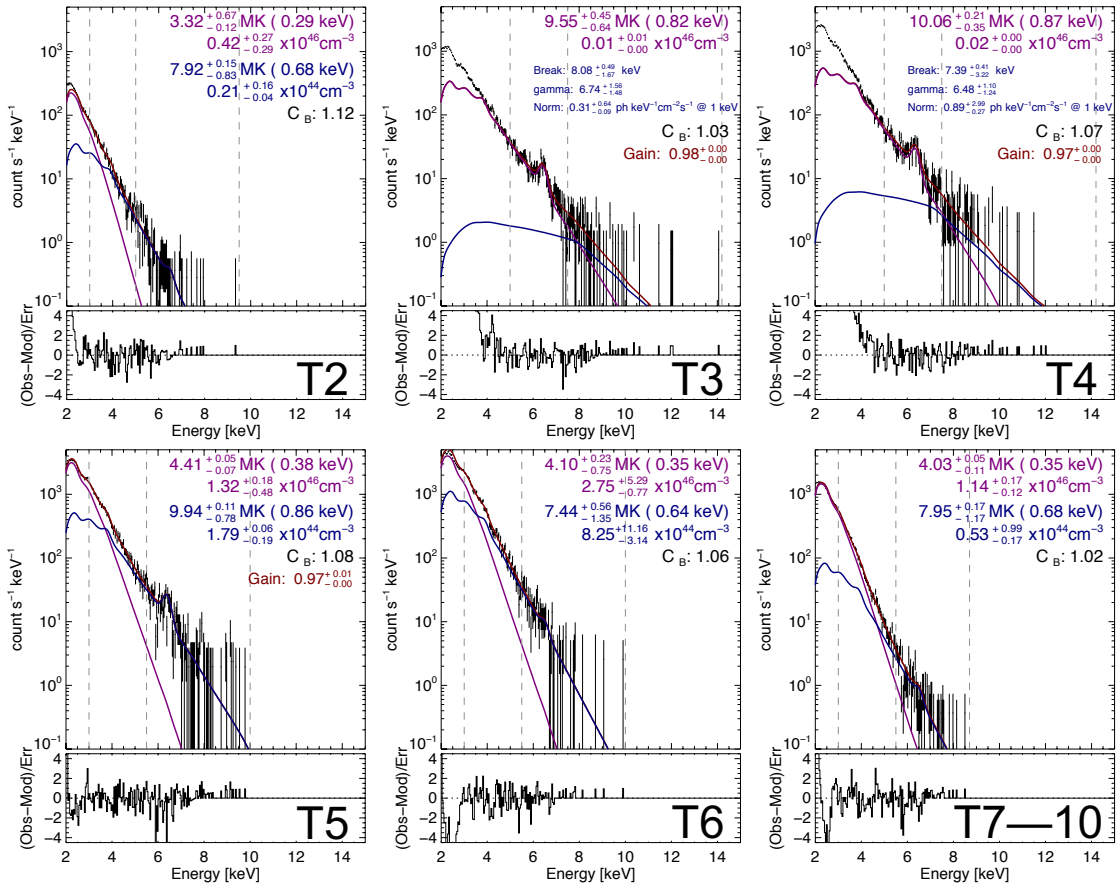

Figure 8.5: Microflare 1's spectra for time range 2, 3, 4, 5, 6 and 7–10 shown in Figure 8.1. Time range 1 is omitted as it is very similar to time range 2. The thermal or non-thermal parameter values are shown along with the FPMB scalar factor and gain correction factor if applicable. The vertical dashed lines indicate the energy fitting ranges used where one range was used to fit the lower temperature thermal model and the upper range being used to fit the hotter thermal model or non-thermal photon power-law model.

individual time range spectra show very similar behaviour and we find similar plasma temperatures to time 6 (bottom middle panel). The pre-flare temperature of ~4 MK appears to be present throughout the whole microflare as it is found when representing the emission <5 keV with a thermal model. In addition, we find consistent signatures of ~10 MK microflaring plasma as well as broken photon power-law values consistent with those found for Chapter 3's microflare 3 with a break energy of 6–7 keV and a photon index of ~8.

Unlike the first microflare, microflare 2 appears to start of as clearly isothermal with a temperature of 3–4 MK in time 1 and 2 (Figure 8.6, top left panel), where time range 1 has been omitted due to its similarity with time 2. We then fit the impulsive time with two thermal models finding the same pre-flare component temperature with larger emission measure and another component at

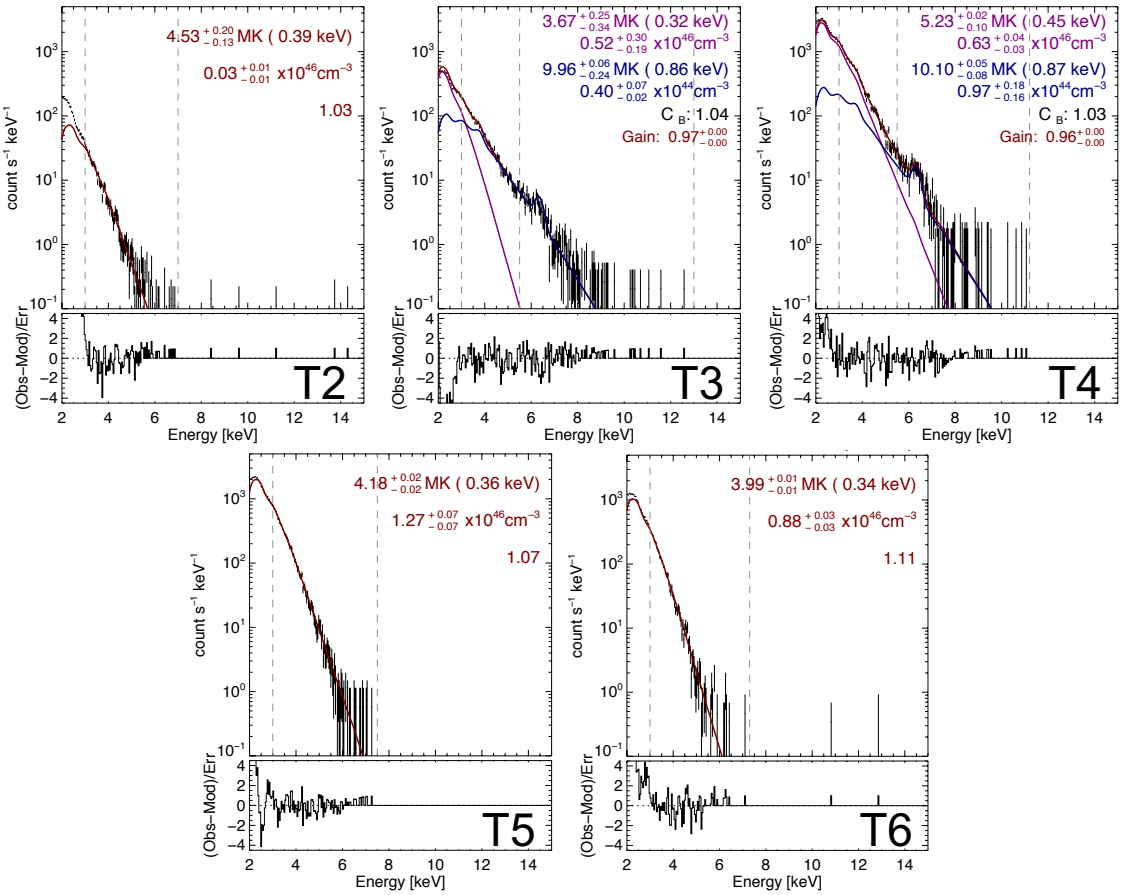

Figure 8.6: The same format as was used for microflare 1 in Figure 8.5 but for microflare 2's spectra at time 2, 3, 4, 5 and 6 indicated in Figure 8.2. Time ranges 1 and 7 were omitted due to them being very similar to time ranges 2 and 6, respectively.

~10 MK (Figure 8.6, top middle panel). The peak time (time 4, top right panel) shows an enhanced pre-flare component of 5.2 MK and factor of two increase in the emission measure for the 10 MK model. Progressing beyond the peak time we find a return to quiescent AR temperatures of 4 MK and then finally ~3 MK for time ranges 5 and 6 (Figure 8.6, bottom row). Time 7 has not been included as it gave a similar fitting result to the time 6 spectra. Despite the differences of microflare 1 and 2, they both required a gain correction and both reach similar temperatures.

Microflare 2's time 4 spectra (Figure 8.6, top right panel), and to a greater extent time 3's spectra (top middle panel), have a possible count excess >8 keV. This could indicate the presence of hotter plasma observed by NuSTAR or accelerated electrons. Therefore, an additional model may be needed to appropriately represent the observed emission.

As discussed, an instrumental response gain correction was required throughout both microflare 1

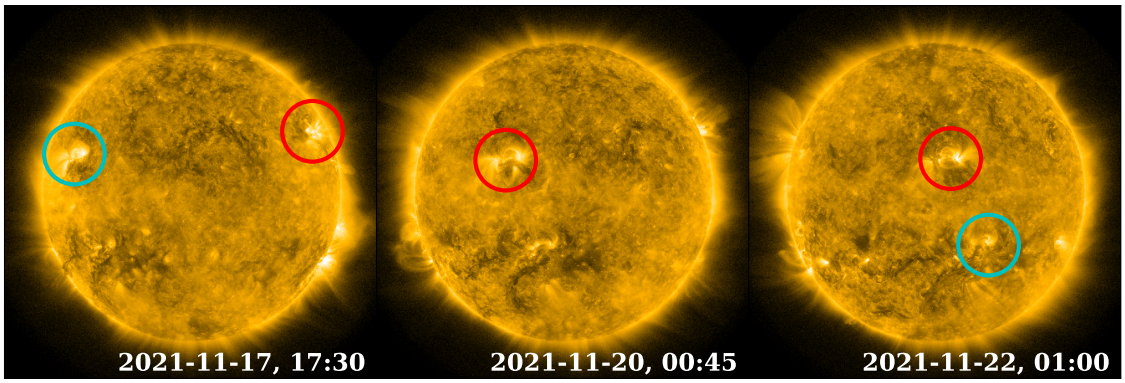

Figure 8.7: SDO/AIA 171 Å full disk images on 2021 November 17 at 17:30 UTC (left panel), November 20 at 00:45 UTC (middle panel), and November 22 at 01:00 UTC (right panel). The red circle indicates the first NuSTAR target for a set of observations while the blue, if applicable, indicates the second.

(Figure 8.5) and 2's (Figure 8.6) spectral fitting. Note that a gain correction is only applied to spectra with clear line features and maintains a consistent value between 0.96–0.98 when present. This is not to say that a spectrum without line features does not need a gain correction but that the correction is difficult to constrain without them. This is because there are no bremsstrahlung continuum features located at specific energies that cannot just be fitted by varying model parameters and so there is no strong justification to vary any instrument response parameters. In addition, the gain correction only removes features in the residuals and does not drastically affect the fitted parameters.

This analysis will be progressed by performing MCMC analysis on the fitted results to investigate the probability density distribution of the parameter-space while quantitatively obtaining the most probably total model fit to each spectra using nested sampling via Sunxspex (Chapter 7). We also intend to obtain the DEM distribution of the emitting plasma being observed to better understand the temperature evolution throughout each microflare (Section 1.4.4 and 4.5.1).

8.2 NuSTAR Solar Observation Campaign: 2021 November 17–22

A NuSTAR solar observation took place between 2021 November 17–22 where nine 1 h dwells were performed with the last one being split into a 50 min and a 10 min dwell on two separate regions. Figure 8.7 shows a full disk SDO/AIA 171 Å image on November 17 at 17:30 UTC (left panel), November 20 at 00:45 UTC (middle panel), and November 22 at 01:00 UTC (right panel). The targets for the observation are indicated with circles, red being the first target and blue being the second if applicable.

8.2.1 Campaign Overview

We find a great number of events occurring across the 2021 November 17–22 solar observation campaign. A total of seventeen events are visually identified from the field of view integrated NuSTAR time profiles with nine full microflare profiles, three partially caught microflares, and five other events. Every orbit is given a brief overview which is shown in Table 8.1.

We find that the activity of the Sun did not change drastically between the observation planning and the dwells with all targeted sources within the FOV and away from the chip-gaps as expected. The livetime remains relatively high throughout every observation (5%–35%), increasing NuSTAR's sensitivity to weak higher energy emission. We also find weak, to no ghost-ray contamination across the four-day period.

In addition, the vast majority of microflares/events occur in isolation with no other co-spatial or co-temporal activity. This means that a great number of events can be studied completely separately through time or region selection and also that a pre-activity time can be defined for many events, helping with the characterisation of background or surrounding emission, honing in on only that emission produced by the event itself. Therefore, we see that many of the events occur under very favourable conditions and have been optimally observed by NuSTAR.

We focus on two orbits from this data-set, orbit 3 on November 17 and orbit 8 on November 22. NuSTAR's orbit 3 captures three very impulsive microflares all having similar X-ray time profiles but quickly become different when considering emission in other wavelengths. In orbit 8, a jet north of the AR being targeted is also observed by NuSTAR. In the rest of this section, we investigate the differences between the three X-ray microflares by utilising SDO/AIA observations (Section 8.2.2) and study the spectral nature of the jet emission (Section 8.2.3).

8.2.2 Orbit 3: Three Impulsive Microflares

Three incredibly impulsive microflares are observed during orbit 3, each lasting approximately 3–5 min (SOL2021-11-17T21:03, SOL2021-11-17T21:14, and SOL2021-11-17T21:25). All three microflares occur spatially separate in temporal isolation and were recorded by DET-ID 0 (one of NuSTAR's four detector quadrants, see Section 2.1) on both FPMs from AR SPoCA 26190. Even though there are multiple sources across NuSTAR's FOV in EUV (seen in SDO/AIA images), we find no obvious counterpart to the clear X-ray emission. Therefore, we utilise the Fe XVIII proxy channel which allows the microflaring structures to be located in these higher spatial resolution images.

Table 8.1: NuSTAR orbit summary of the 2021 November 17–22 solar observation. The main features observed by NuSTAR are documented as well as co-observing instruments such as Hinode/XRT and EIS, the radio detecting Long Wavelength Array (LWA), and the Interface Region Imaging Spectrograph (IRIS). SDO/AIA, STEREO-A, and STIX data is available during all orbits. The three microflares observed in orbit 3 and the jet in orbit 8 are presented in Section 8.2.2 and 8.2.3, respectively. The target of each orbit is indicated with a coloured dot representing the coloured circles in Figure 8.7.

Observation Date	Orbit	Time (HH:MM)	Description	Livetime (%)	Other Instruments
Nov. 17	1•	17:28–18:30	Quiet AR emission	~10	XRT (18:03–18:08)
	2•	19:05–20:06	1 microflare Start & end of another	10–15	XRT LWA
	3•	20:42–21:33	3 microflares	15–20	XRT LWA
Nov. 19 to	4•	22:37–23:34	AR emission 2 weak microflares/events	20–30	XRT LWA (19:00–23:00)
Nov. 20	5•	00:13–01:15	3 potential microflares	15–30	XRT
	6•	01:50–02:52	3 microflares Another microflare or jet	5–30	XRT
Nov. 21 to Nov. 22	7•	22:55–23:53	Start of 1 microflare	10–30	XRT IRIS (22:50–23:15) EIS (23:20–23:40)
	8•	00:32–01:33	Quiet AR emission Small jet	20–25	XRT IRIS (01:30–end) EIS (00:50–01:00)
	9(a)•	02:09–02:59	AR emission A very weak event	20–30	XRT IRIS (start–02:30)
	9(b)•	03:00–03:10	Quiet AR emission	20–35	

Overview

During NuSTAR's third orbit on 2021 November 17, AR SPoCA 26190 is observed (Figure 8.8, top panels) where this AR was one of the planned targets for the solar campaign. We find that the AR dominates NuSTAR's throughput and is entirely located in DET-ID 0 for both focal plane modules while a weaker region is seen to the south of NuSTAR's FOV.

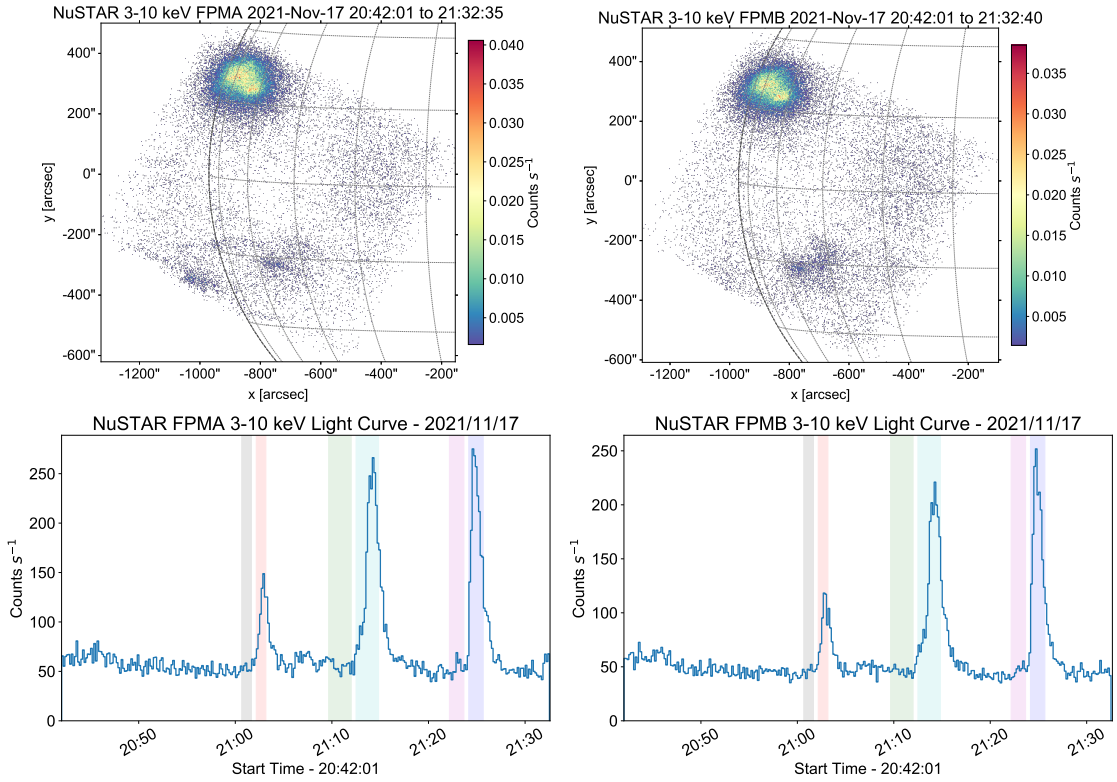

Figure 8.8: NuSTAR FPMA and B images integrated over orbit 3 (top left and right panels) where the AR corresponds to the blue circle area in Figure 8.7 (left panel). Full FOV orbit 3 NuSTAR time profiles showing the three microflares under investigation (bottom left and right panels). The orange, blue, and purple shaded regions indicate the spectral fitting time for each microflare while the grey, green, magenta shaded regions show a corresponding pre-flare time. All NuSTAR data is integrated over 3–10 keV and corrected for livetime.

Investigating the time evolution of AR SPoCA 26190, we find three very distinct microflare time profiles shown in Figure 8.8 (bottom panels). Both light curves from FPMA and B show similar behaviour and both are able to be used during spectral analysis. We aim to investigate the impulsive and peak microflare phases and so define the microflare times using orange, blue, and purple shaded regions (Figure 8.8, bottom panels). Since each microflare occurs in isolation, separated from each other, we are able to define pre-flare times to help constrain the surrounding plasma during spectral fitting.

Despite the microflares being prominent in X-rays, they are difficult to locate in EUV. Inspecting SDO/AIA observations, we find many dynamic events occurring during orbit 3's time (Figure 8.9, top row). However, producing the Fe XVIII proxy channel (Section 2.5.1), we find the locations of the microflares and are able to focus on those features across all other channels. The Fe XVIII emission shows that the first microflare (SOL2021-11-17T21:03) occurs in the northern box region with a double loop or 'M'-shaped structure and has a complicated geometry throughout its evolution

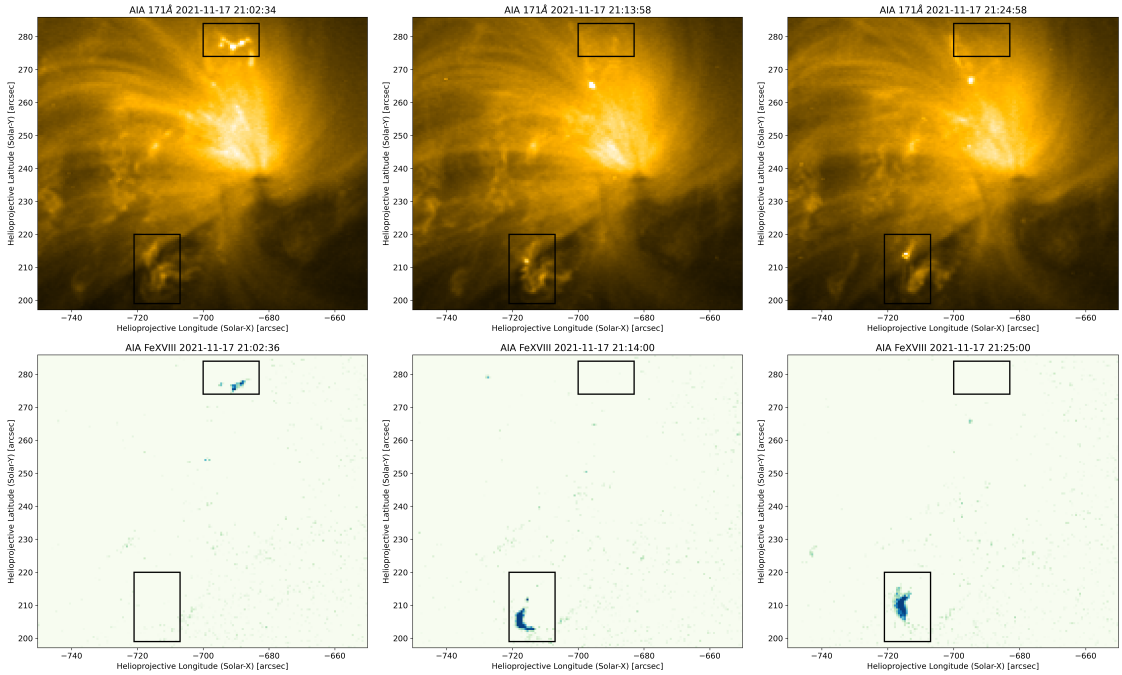

Figure 8.9: SDO/AIA 171 Å images (top row) and Fe XVIII proxy images (bottom row) of times near the peak X-ray emission of microflare 1, 2, and 3 (SOL2021-11-17T21:03, SOL2021-11-17T21:14, and SOL2021-11-17T21:25, respectively) from left to right. Black boxes indicate microflaring regions with microflare 1 (left panels) occurring in the northern box and microflare 2 and 3 (middle and right panels) occurring in the southern box. We observe that microflare 1 is shown as an 'M'-shaped structure with microflare 2 and 3 being more loop-like.

(Figure 8.9, left panels). The complex structure is clearer in the 171 Å image (top left panel). The successive two microflares, SOL2021-11-17T21:14 and SOL2021-11-17T21:25, occur in the southern region and both produce a relatively simple loop-like volume from the same general structure (middle and right panels, respectively). Therefore, despite the three microflares displaying similar behaviour in X-rays, their EUV emission hints at a different style of evolution.

Inspecting the EUV time profiles from the box regions shown in Figure 8.9 we find that the 94 Å emission shows similar behaviour to the X-ray light curve with an even stronger similarity with the Fe XVIII channel (Figure 8.10, top two panels). We find that microflare 1 creates a sizeable response across almost all SDO/AIA channels suggesting a wide range of temperatures are produced. Whereas, in microflare 2 and 3 we only find an obvious response in the hotter SDO/AIA channels, primarily Fe XVIII and 94 Å. This may be connected to the different geometries displayed in the EUV images between the first microflare and the last two (Figure 8.9).

Therefore, microflare 1 from the northern box appears to evolve very differently to the other two microflares from the southern region. This could be due to their different general microflaring

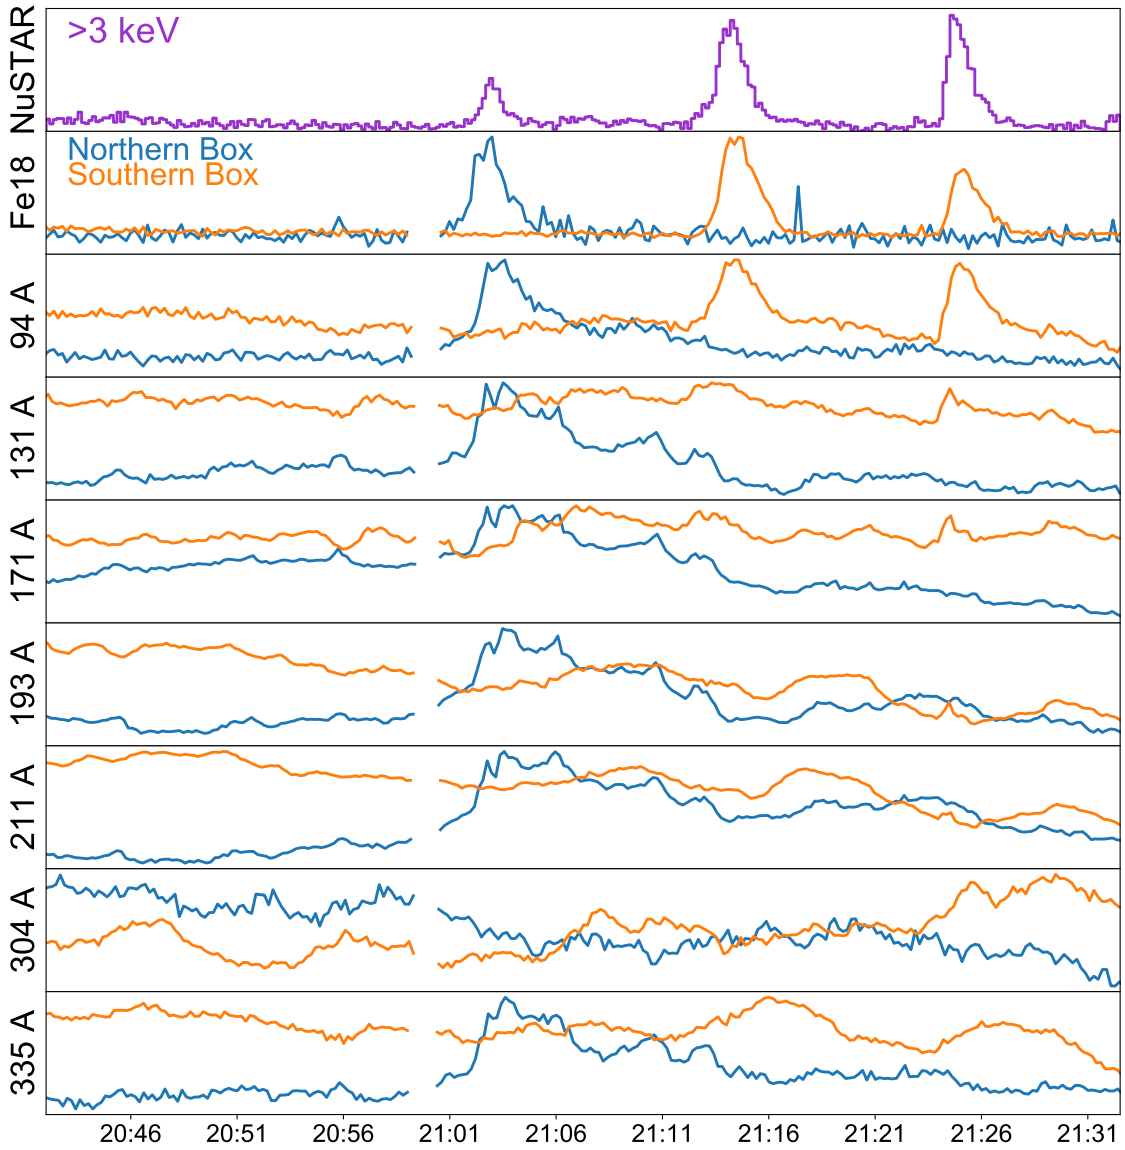

Figure 8.10: NuSTAR FPMA >3 keV X-ray (top panel) and SDO/AIA EUV (all other panels) maximum normalised time profiles. The X-ray time profile is integrated over NuSTAR's FOV and the EUV light curves show emission from the northern box (blue) and the southern box (orange) indicated in Figure 8.9. The EUV light curves shown are from the Fe XVIII, 94 Å, 131 Å, 171 Å, 193 Å, 211 Å, 304 Å, and 335 Å SDO/AIA channels. There is an unknown spike in the blue Fe XVIII time profile at ~21:17; however, it should be noted that this light curve is much noisier than the one produced from the southern box (orange) and, therefore, it might be an artifact due to noise in the Fe XVIII proxy.

structure and not just because microflare 1 appears the weakest in X-rays (Figure 8.8). We investigate the difference between the three microflares further by utilising NuSTAR's spectroscopic abilities.

This will provide a more detailed insight into the microflaring plasma's nature, allowing us to probe the physics at play.

Spectral Analysis

We perform the spectral fitting utilising Sunxspex described in Section 5.4 and Chapter 6; therefore, we perform spectral analysis on microflare 1's FPMA & B spectra (Figure 8.11). We first fit the pre-flare time and find that those spectra are well described by an isothermal model of temperature 3.6 MK and emission measure 6.2×10^{45} cm^{-3}, this is kept as a fixed component during the microflare fitting. The pre-flare isothermal model is shown in Figure 8.11 (blue, all panels). Upon simultaneously fitting FPMA & B with an additional thermal model or non-thermal model (thick target model, see Section 1.4.5) it quickly becomes apparent that two models are required above the pre-flare component; this is determined from residual artifacts and relative fit statistic values.

From the spectral fit we find that the two fitted thermal models could represent an enhanced pre-flare component of 4.8 MK and 5.6×10^{44} cm^{-3} (Figure 8.11), similar to that found for Chapter 3's microflare 10; however, we find evidence of much hotter and fainter material in this weak, sub-A class microflare with a temperature and emission measure of ∼15 MK and $\sim10^{42}$ cm^{-3}, respectively. It is unlikely that these higher energy counts could be fitted with a non-thermal model instead of the 15 MK component as this would suggest that only cool, 3.5–4.5 MK, plasma is present. However, no matter the model used it is important to acknowledge the few counts constraining this component and, therefore, detailed comparisons with SDO/AIA emission and DEM analysis must be performed in order to investigate this further.

Performing spectral fitting with microflare 2 and 3 in a similar nature we also find that two photon models are required beyond the fixed pre-flare component (Figure 8.12 and 8.13, respectively). However, unlike the analysis for microflare 1, we find similar fits and fit statistics (shown in residual panels for each fit) when either using a thermal (top panels) and non-thermal (bottom panels) model for the higher energy counts. Moreover, both microflare 2 and 3's spectral fits tend to slightly favour the results that utilise the thick target. Therefore, we either have more incredibly weak sub-A microflares, which is on the order of 10^{25-26} erg (Equation 1.15), reaching very hot temperatures of ∼15 MK or direct evidence of non-thermal emission that provides $\sim10^{26}$ erg of energy over the course of the microflare time (Equation 1.31).

Therefore, we find that microflare 1's spectral analysis is consistent to what is found in EUV where microflare 1 was found to likely be strongly multi-thermal from the SDO/AIA time profiles. In addition, we find that the microflare 2 and 3 non-thermal fits have a more preferable fit statistic value to fit the higher energy counts compared to a thermal model. However, there are few counts

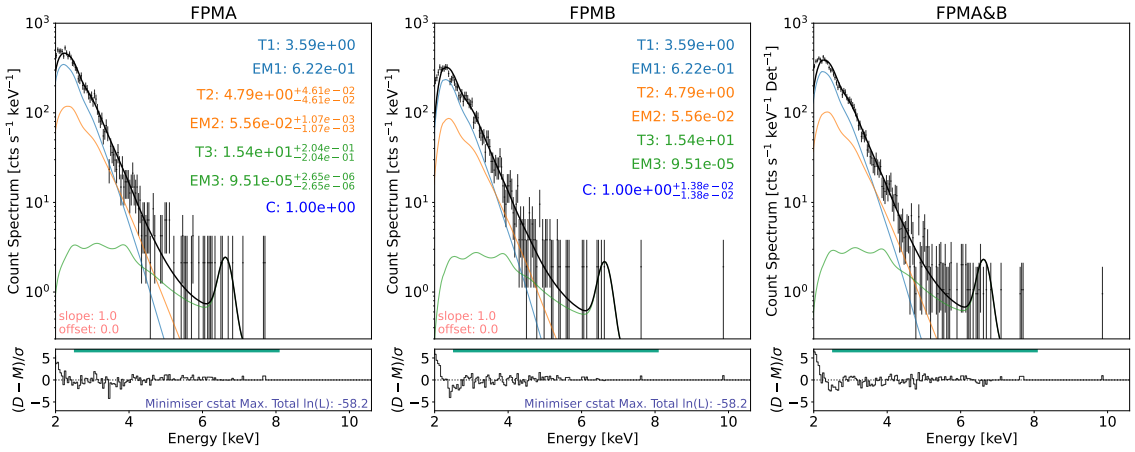

Figure 8.11: NuSTAR FPMA (left panel) and FPMB (middle panel) spectra for microflare 1 for the time range 02:10–03:10 (Figure 8.8, bottom row) fitted with three thermal models, one representing pre-flare emission and two for the microflare excess. The right panel depicts the mean result from both FPMs. The pre-flare spectra (00:40–01:40) were found to be well fitted by an isothermal model of 3.6 MK and 6.2×10^{45} cm^{-3} and so was fixed for the microflare time fitting. The horizontal green line shows the energy fitting range and the 'T' & 'EM' parameters are in MK & 10^{46} cm^{-3}, respectively. The 'C' parameter represent a scalar factor to account for systematic differences between the instruments.

constraining this space; therefore, we used a more powerful analytical tool, *nested sampling*, for microflare 2's model selection in Chapter 7. In Section 7.3.2, we find that there is very strong evidence that microflare 2's fit using the non-thermal component for the higher energy range describes the data better than the fit with the higher temperature thermal model. Nested sampling is yet to be used to determine the most probable model fit to microflare 3's spectra.

A potential explanation for the X-ray and EUV behaviour is that the flare accelerated particles are hindered in their propagation in some way by the complex geometry of microflare 1, causing energy to be distributed in the corona and somehow heating material through the channel response temperatures. Whereas, the relative simple, loop-like geometry of the last two microflares could allow the flare accelerated particles an unobstructed path to lower levels of the atmosphere, heating material up quickly and not producing a clear response in the cooler SDO/AIA channels.

We inspect the SDO/AIA UV channels, 1600 Å and 1700Å, to probe potential microflare corresponding signatures further down in the solar atmosphere. Therefore, we can investigate if microflare 1 has any, or a weaker, impact on the solar chromosphere compared to microflare 2 and 3.

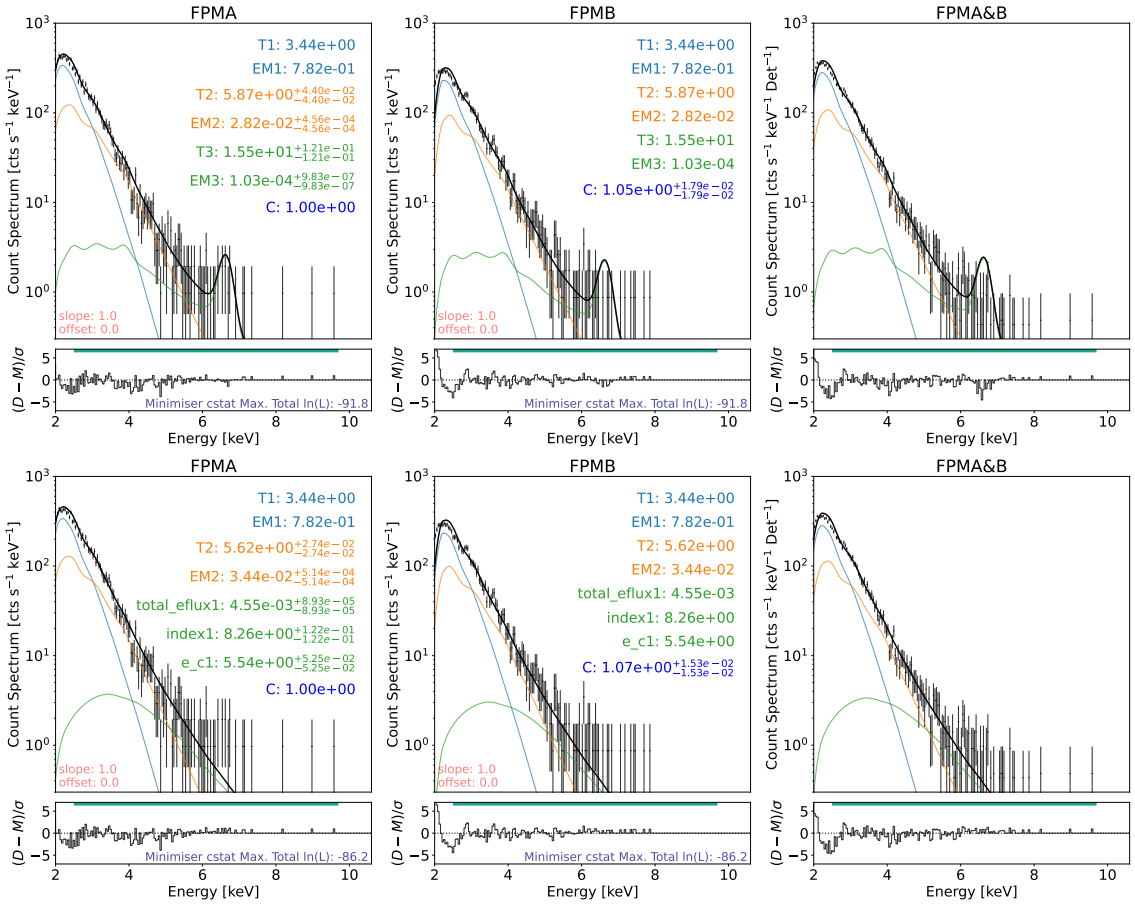

Figure 8.12: The top row is in the same format of that shown in Figure 8.11 but for microflare 2. The bottom row is similar but the higher energy emission is fitted using the thick target model (Section 1.4.5). The pre-flare time range is 09:40–12:00 with a microflare time of 12:30–14:50 (Figure 8.8, bottom row). The units for the 'total_eflux' and 'e_c' parameters is 10^{35} e^{-} s^{-1} and keV, respectively, with 'index' being dimensionless.

SDO/AIA UV Channels

Figure 8.14 shows the time profile of the Fe XVIII proxy and the 1600 Å and 1700 Å SDO/AIA UV channels from the northern (blue) and southern (orange) boxes shown in Figure 8.9. We find that the UV channels show oscillating behaviour throughout the NuSTAR orbit.

The northern region appears to produce the same undisturbed wave pattern for the full time displayed for both UV channels. This may suggest limited or no interaction between microflare 1 and the lower solar atmosphere. However, when considering the southern region, we find that the oscillation appears to increase in amplitude during microflare 2's onset and then is perturbed further at the onset of microflare 3. This may indicate an interaction between microflare 2 and 3 and the chromosphere.

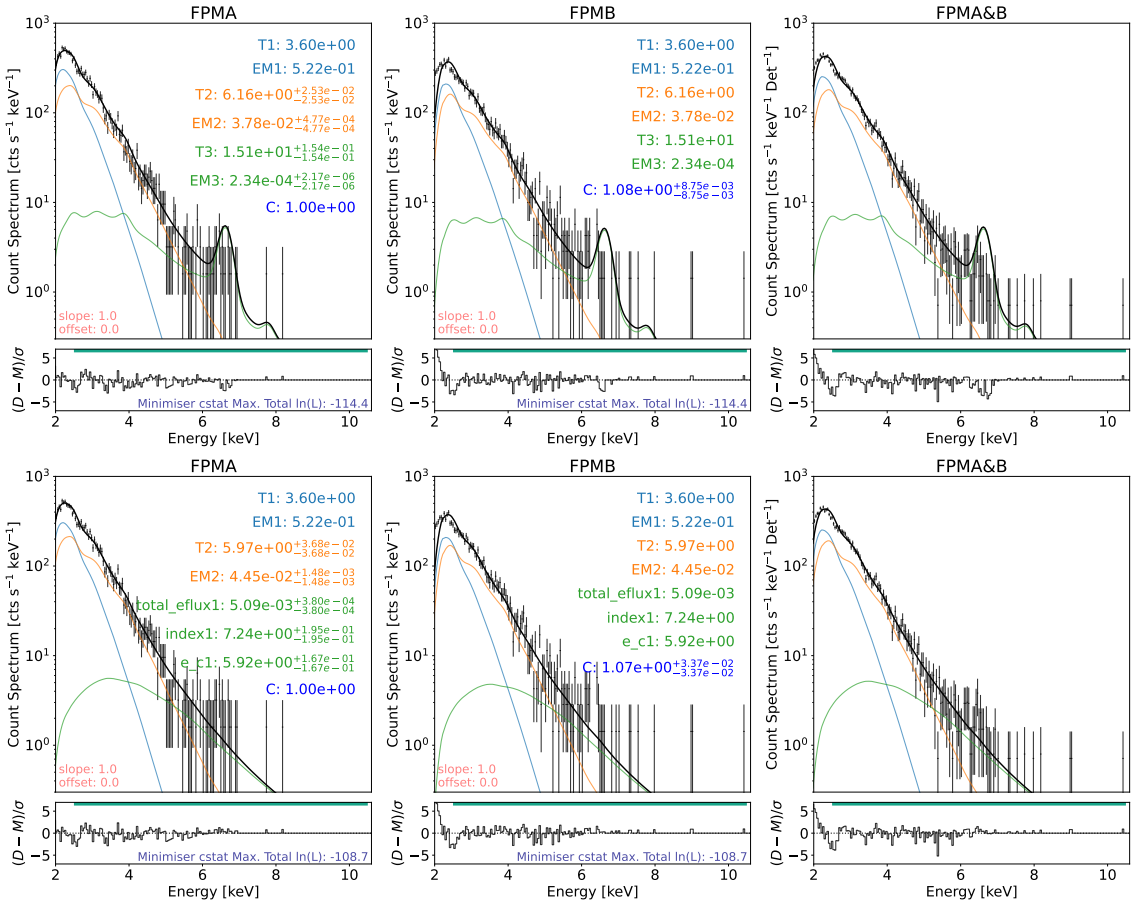

Figure 8.13: Same format as Figure 8.12 but for microflare 3. The pre-flare time range is 22:10–23:40 with a microflare time of 24:10–25:40 (Figure 8.8, bottom row).

Future work on the spatial, temporal, and spectral analysis of these three microflares will include investigating their similarities and differences further with all instruments that observed them, such as Hinode's soft X-ray Telescope and the radio Long Wavelength Array, shown in Table 8.1. This will help in the understanding of how the apparent different geometries between microflare 1, 2, and 3 affect their different responses across different wavelengths. In addition, we can investigate the perturbations of the wave pattern observed in Figure 8.14 with wavelet analysis (Torrence & Compo 1998). This would allow a measure of the oscillating signal from both regions and determine whether any other waveform is introduced, or significant deviation is caused, by the microflares.

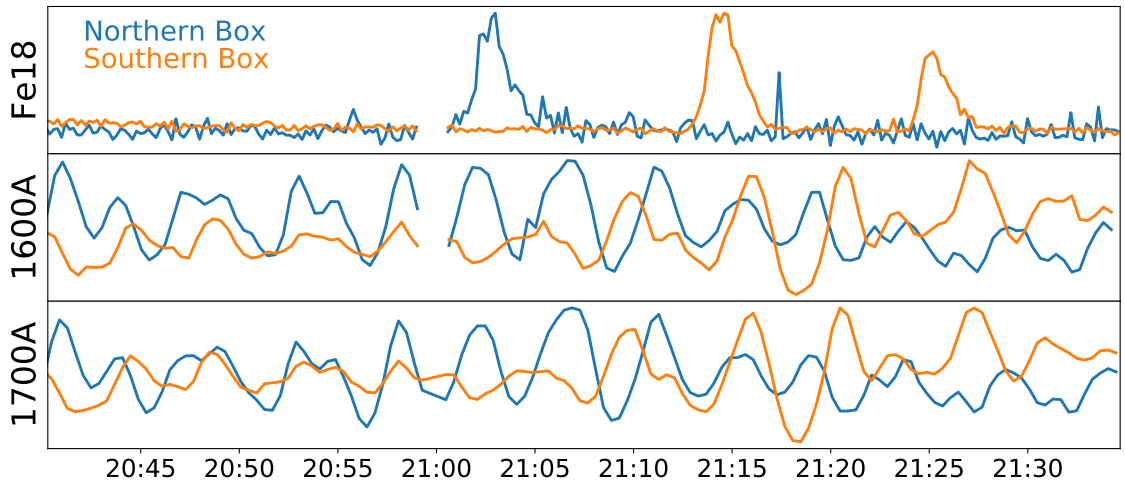

Figure 8.14: Fe XVIII and UV maximum normalised time profiles of the northern box (blue) and the southern box (orange) shown in Figure 8.9. The light curves shown are from the Fe XVIII, 1600 Å, and 1700 Å SDO/AIA channels.

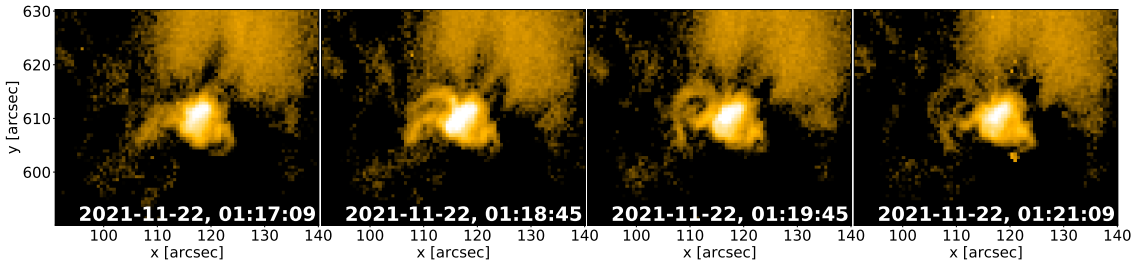

Figure 8.15: SDO/AIA 171 Å images zoomed of the jet event that occurred on 2021 November 22. The jet takes place between ~01:18–01:20 where 171 Å images are shown before (01:17:09, left panel), during (01:18:45 and 01:19:45, middle two panels), and after (01:21:09, right panel) the time range.

8.2.3 Orbit 8: A Jet

During orbit 8 of NuSTAR's 2021 November 17–22 solar observation campaign we target AR SPoCA 26196 which appears to dominate NuSTAR's FOV (Figure 8.16, top panels). However, a jet like event appears to occur between 01:18–01:20 to the left of a bright structure (Figure 8.15) which is detected by NuSTAR to the north of the AR which is corroborated via visual inspection of co-temporal SDO/AIA images. The jet is observed in X-rays as a cluster of counts indicated by a red arrow in Figure 8.16 (top panels).

From Figure 8.17 (blue curves), we observe that the AR dominates 87%–89% of the emission seen throughout orbit 8 and is solely viewed by DET-ID 0. However, the jet is observed by NuSTAR even though its signal is only ~30 counts above the surrounding emission. The high livetime during

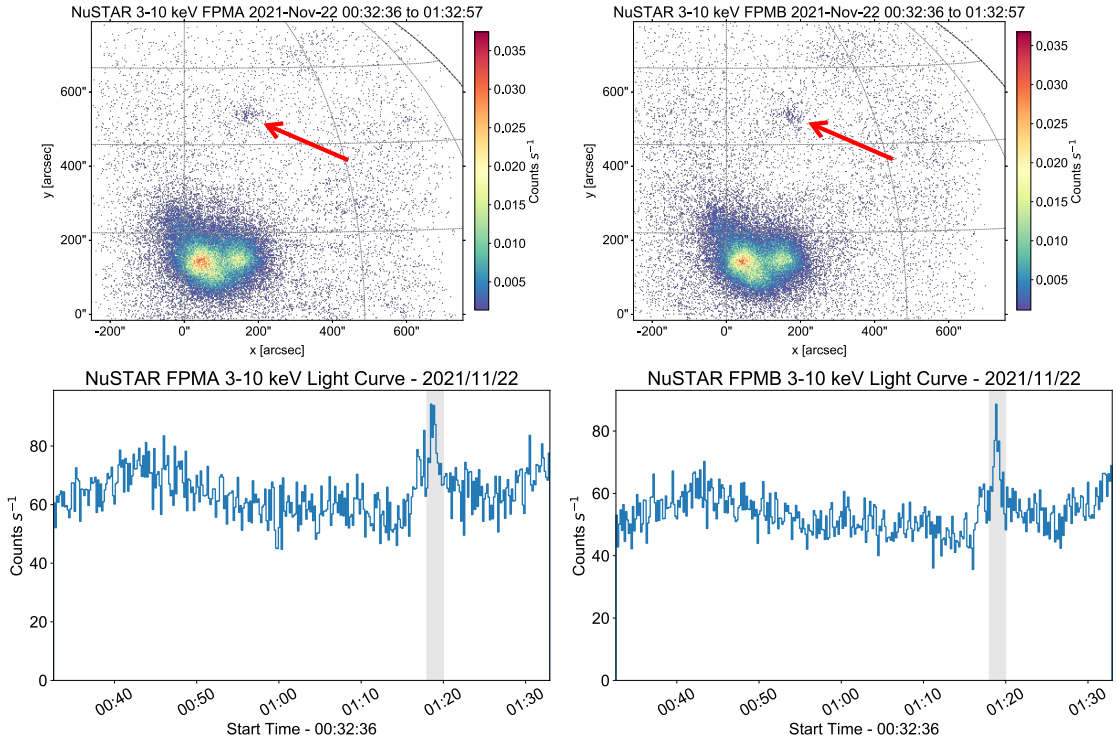

Figure 8.16: NuSTAR FPMA and B images integrated over orbit 8 (top left and right panels) with the jet location indicated by a red arrow. The AR corresponds to the red circle area in Figure 8.7 (right panel). Full FOV orbit 8 NuSTAR time profiles showing the jet under investigation (bottom left and right panels). The grey shaded region indicates the spectral fitting time for the jet (01:18–01:20). All NuSTAR data is integrated over 3–10 keV and corrected for livetime.

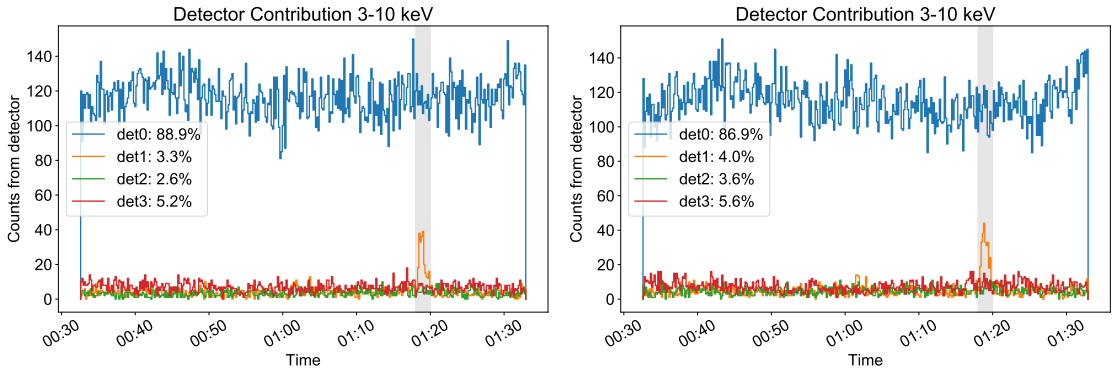

Figure 8.17: NuSTAR time profiles for orbit 8 produced from each of the four detectors on both FPMA (left) and B (right). The DET-ID 0, 1, 2, and 3 (blue, orange, green, and red) time profiles correspond to the bottom left, top left, top right, and bottom right of the NuSTAR FOV, respectively, in Figure 8.16 (top panels). The grey shaded region indicates the spectral fitting time for the jet and the percentage of total emission each detector observes is displayed. The data is integrated over 3–10 keV and in count-space; therefore, they are not livetime corrected.

this time, about 20%–25%, undoubtedly aides the detection of the jet. In addition, we find that the jet is captured on a single detector for both FPMs (Figure 8.17). This allows both FPMs to be used for quantitative analysis of the event which is very beneficial, especially with such a low number of counts.

We continue the analysis of the jet by performing spectral fitting using XSPEC. The defined 2 min time range for analysis is shown in Figure 8.16 and 8.17 with a grey shaded region. Figure 8.18 shows the result of the spectral fitting where we find the spectrum to be predominantly isothermal with a temperature of 4.3 MK and emission measure 2.4×10^{44} cm^{-3}. Therefore, this jet appears to have similar thermal properties as those studied in Kuhar et al. (2018) which also occurred outside an AR.

We investigate the effect of including a background time fixed spectral component in the fitting process. We find that the jet temperature and emission measure stay well within uncertainty ranges when this model is included and, therefore, it is omitted. A visual inspection of SDO/HMI (Figure 8.19) reveals two small oppositely polarised regions colliding in the line-of-sight magnetic field. This could, therefore, be another example of photospheric magnetic flux cancellation causing the onset of solar events as discussed in Section 3.9.

Future work regarding the jet event will include investigating observations made by other instruments, such as Hinode's X-ray Telescope, as indicated by Table 8.1. Utilising the soft X-ray observations will allow us to obtain a more complete picture of the jet's temporal and spatial evolution. We may also be able to investigate the jet region shown in Figure 8.15 before and after the event with the Interface

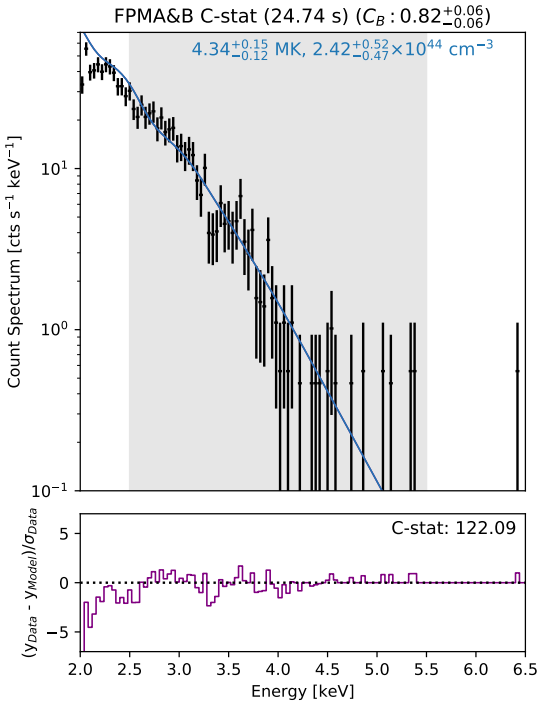

Figure 8.18: NuSTAR FPMA&B spectral fit of the jet observed at 01:18–01:20. The isothermal fitted model temperature and emission measure are shown along with the effective exposure and instrument scaling factor for FPMB relative to FPMA. The energy fitting range is shown by the grey shaded region.

Region Imaging Spectrograph (IRIS; De Pontieu et al. 2014) and Hinode's EUV Imaging Spectrometer (EIS; Culhane et al. 2007), respectively.

8.3 Summary and Conclusions

NuSTAR observed two microflares on 2020 January 30, both of which show strong evidence that ~10 MK plasma is present despite the different geometries. In addition, the microflare with the more complicated evolution shows convincing non-thermal signatures in the initial microflaring stages.

The 2021 November 17–22 NuSTAR solar observation campaign was incredibly successful and we observe a plethora of weak solar events, both from the planned targets and other sources. In addition to a large number of events we also have a number of co-observations by other instruments throughout many of the orbits (Table 8.1). We then analyse the events in two orbits, orbit 3 and orbit 8.

Similar to the microflares discussed in Chapter 3, it appears that the brighter sub-A microflares observed by NuSTAR from orbit 3, presented in Section 8.2.2, are represented well with two thermal

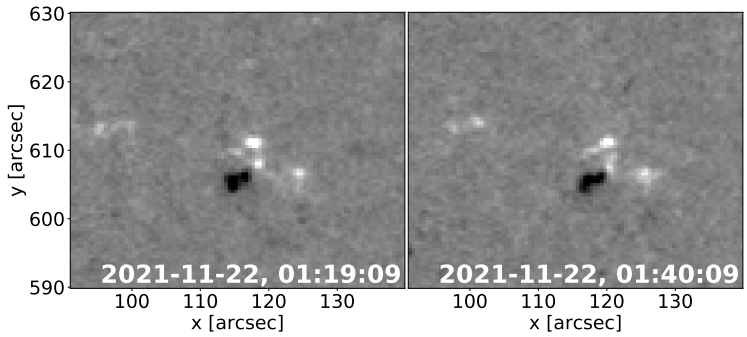

Figure 8.19: The SDO/HMI photospheric line-of-sight magnetic field co-spatial with the jet event in Figure 8.15 at the times indicated. The left panel shows the magnetic flux during the jet (01:19) while the right panel shows the same region ~20 min later (01:40). White and black represents positive and negative polarity, respectively.

models. One thermal model represents the pre-flare/quiescent AR emission and the other represents flare heated plasma. However, an addition model is required to fit a residual count excess above these two thermal models. This can either be another thermal model indicating even hotter temperatures or a non-thermal model to represent non-thermal electron emission.

We find that the first microflare of orbit 3 may reach temperatures $>$15 MK. Microflare 2 and 3 are also fitted with thermal models of temperatures $>$15 MK but they are also fitted equally well with a non-thermal model instead. The parameters of the non-thermal model would provide an energy of a similar magnitude needed for the thermal component (Section 8.2.2). To solve this issue we performed nested sampling (performed in Chapter 7 and discussed in Section 7.3.2 and 8.2.2) to conclude that microflare 2 residual excess above the two thermal models was more likely to be represented with the non-thermal model. We also find a temperature of ~4 MK for a jet, a typical hot AR core temperature, far removed from an AR.

Much of the spectral analysis in this chapter has been perform using XSPEC (Section 8.1.2 and 8.2.3). However, we also utilise Sunxspex for Section 8.2.2, performing crucial analysis that would be difficult to execute in other spectral fitting tools.

8.3.1 Future Work

Although much NuSTAR spectral analysis has been performed for the 2020 January 30 microflares discussed in Section 8.1, more work is to be done in studying the time and spatial evolution of each microflare by incorporating observations from other observatories like GOES/XRS and Hinode/XRT. Using the additional instrumental data, we can perform DEM analysis to better describe the temperature evolution throughout the microflares. More work should be performed to investigate the

possibility of separately resolved X-ray footpoints in these microflares since they are observed to have loop lengths greater than NuSTAR's spatial resolution. This could not have been done for any other presented microflare as they have similar or smaller lengths compared to NuSTAR's resolution.

Moreover, the parameter-spaces of the spectral model fits for the 2020 January 30 microflares should also be explored through an MCMC process to obtain a better understanding of the parameter posterior distribution and more robust uncertainty estimates. Performing nested sampling on the spectral fits would achieve this while also giving a confidence value for which model fits should be preferred. This, in turn, provides greater confidence in the energetic values we must obtain from the spectral fits.

The three 2021 November 17 microflares discussed in Section 8.2.2 must be studied in more detail spatially, temporally, spectrally, and energetically in order to identify what makes microflare 1 so different to microflare 2 and 3. Incorporating data from GOES/XRS, the Long Wavelength Array (LWA; Ellingson et al. 2009), and Hinode/XRT (shown in Table 8.1) will help with this, providing a more comprehensive understanding of all three microflares and why they appear differently to each other across wavelengths.

In addition, nested sampling will be performed on microflare 3 in order to determine whether a non-thermal model is most probable to fit the spectra, like what was done for microflare 2 in Section 7.3.2. We also aim to perform wavelet analysis (Torrence & Compo 1998) on the oscillatory signal shown in Figure 8.14 for both chromospheric channels (bottom two panels). The wavelet analysis will indicate if new periods related to the microflares are introduced into the chromospheric oscillatory signal. This analysis will help asses if or how each microflare interacts with the lower solar atmosphere.

The jet on 2021 November 22 (Section 8.2.3) has been identified in the X-ray data and was found to be located north of the AR being targeted using NuSTAR and SDO; however, Hinode/XRT also observed the jet, as well as IRIS and Hinode/EIS observations before and after the event, respectively. Using these other observations, we will be able investigate this very weak event more completely.

We also aim to progress this research by investigating all events in the 2021 November NuSTAR solar campaign shown in Table 8.1. Similar analysis to what is already shown will be performed on each event as well as the analysis described above. This research will greatly increase the number of studied NuSTAR microflares and other events, allowing a more statistical view of very weak X-ray microflares to be obtained.

9

Conclusions and Future Work

In this thesis, we initially introduce the relevant background for microflares, the phenomenon predominantly studied in the following chapters, and the mathematical modelling used in their analysis (Chapter 1). In the following chapter, Chapter2, we provide a brief overview of all instruments used throughout this thesis with particular emphasis on the Nuclear Spectroscopic Telescope ARray (NuSTAR; Harrison et al. 2013).

In Chapter 3 and 4, we present the research from Cooper et al. (2021) and Cooper et al. (2020), respectively. These two chapters give a detailed insight into X-ray microflare analysis, showing how NuSTAR data can be used in combination with data from other observatories and showing the use of analytical tools, such as the spectral fitting programs OSPEX (Section 5.3.1) and XSPEC (Section 5.3.2).

Chapter 5 details the statistical framework of these spectral analytical tools, indicating their individual advantages and limitations, while presenting the need for a new spectral fitting program in the form of Sunxspex (Section 5.4). Chapter 6 then showcases spectral fitting examples using Sunxspex, both re-analysing published works and also performing analysis that was not possible within OSPEX and XSPEC. For example performing spectral fitting simultaneously on multiple spectra (Section 5.2.5) with the thick target model (Section 1.4.5) using Poisson statistics (Section 5.2.1) or performing nested sampling analysis for model comparison as shown in Chapter 7 (Section 7.3). In Chapter 8, we present more NuSTAR solar events (five microflares and a jet) being analysed by the tools already described in previous chapters. These events further emphasise the need for Sunxspex.

The following section, Section 9.1, brings together the conclusions presented in previous chapters, collating and summarising all work presented. Potential further work is described in Section 9.2, indicating the next steps for some research presented in this thesis, particularly from Chapter 6, 7, and 8.

9.1 Conclusions

Chapter 3 presents the Cooper et al. (2021) study of ten NuSTAR X-ray microflares from AR12721. The ten microflares occur over a two-day period (2018 September 9–10) where we search for evidence of high temperatures and non-thermal emission emitted throughout their evolution. We find that the higher HXR NuSTAR energy ranges consistently showed more impulsive and earlier peaking time profiles when compared to the lower energy emission, indicative of either hot plasma in the initial flare stages or non-thermal emission. We find typical non-flaring AR temperatures of 3–4 MK consistent with other studies (Glesener et al. 2017; Wright et al. 2017; Hannah et al. 2019; Mitra-Kraev & Del Zanna 2019) as well as corroboration of a higher ~10 MK component during microflare evolution (Duncan et al. 2021). The energies of the microflares are in the range 10^{26}–10^{28} erg placing these towards the lower boundary as to what is still considered a microflare (Table 1.1).

One of the weakest non-thermal X-ray microflares (microflare 3, Section 3.4.1) is also presented in Chapter 3 with a non-thermal energy of $1.3^{+0.7}_{-0.4}\times10^{27}$ erg, the same order of magnitude as the thermal energy upper limit ($6.49^{+0.04}_{-0.03}\times10^{27}$ erg), and a GOES class of A0.1. In addition, we find evidence of mixed polarity photospheric magnetic fields at, or close to, the footpoints in eight out of the ten microflares (Section 3.9). The polarities at the footpoints of the two microflares with the most complex physical structures are able to be investigated more closely finding similar magnetic flux cancellation rates to those found in previous studies (Chitta et al. 2020).

In Chapter 4, we discuss the weakest microflare of the 2018 September 9–10 NuSTAR solar campaign, microflare 4, which is the subject of Cooper et al. (2020). We find that this weak microflare, which is barely visible in EUV, reaches temperatures of 6.7 MK with a thermal energy of $1.1^{+0.2}_{-0.2}\times10^{26}$ erg. This is the weakest X-ray microflare in literature and is of a comparable energy to small EUV brightenings of magnetically braided loops discussed in Cirtain et al. (2013). Although there is no direct indication of non-thermal emission, we find that the higher energy time profile (4–7 keV) peaks before the lower energy Fe XVIII and 2.5–4 keV emission. We investigate the possibility that non-thermal emission could be present but consistent with a null detection. A section of the non-thermal parameter space is consistent with this scenario (Section 4.4.2) which would produce a non-thermal energy comparable with the calculated thermal energy. Therefore, non-thermal processes could be present during this very weak and miniscule microflare.

The fundamental properties and statistical nature of X-ray spectral fitting is discussed in Chapter 5. Several likelihoods and fit statistics are discussed as well as common parameter exploration techniques (Section 5.2). These concepts are then related to the abilities of spectral fitting software widely in use, such as OSPEX and XSPEC, where the advantages and disadvantages of both are discussed

(Section 5.3). Finally, the Python X-ray spectral fitting package Sunxspex is introduced (Section 5.4) which aims to combine many of the features present in OSPEX and XSPEC into one program.

Chapter 6 showcases many X-ray spectral fitting examples with Sunxspex. The analysis from several publications was re-performed using Sunxspex, finding good agreement to the published spectral fitted values. We re-analyse the NuSTAR spectral data for microflare 10's pre-flare time in Section 6.1.2 (Cooper et al. 2021, see Section 3.7), a non-thermal fit to NuSTAR's FPMB presented in Glesener et al. (2020) (Section 6.1.3), and a multi-thermal fit to NuSTAR's FPMA and B spectra from Duncan et al. (2021) (Section 6.1.4). In addition to NuSTAR spectra, Sunxspex is designed to handle other instruments such as RHESSI and STIX. Therefore, we are also able to reproduce analysis from Fletcher et al. (2007) for the 2002 October 5 flare (Section 6.2.1).

Moreover, in Chapter 6 using Sunxspex, we are able to perform analysis that either has been difficult or not possible to do in the past with OSPEX and XSPEC. Instead of being limited to just FPMB's spectrum for the Glesener et al. (2020) microflare, we are able to fit the spectra from both FPMs with the thick target model. We are also able to return to microflare 3's rise time spectrum (from Section 3.4.1) in Section 6.1.5 and perform fitting using the thick target model instead of the broken power-law approximation explained in Section 1.4.5. This gives a direct estimate of the non-thermal electron distribution and provides an energy closer to the thermal energy upper limit. In Section 6.2.2, we simultaneously fit separate RHESSI detectors using Sunxspex which can help identify any systematic differences between them. We also show Sunxspex's ability to fit STIX microflare spectra with a microflare observation on 2020 June 6 (Section 6.2.3). Although this spectrum is at the lower sensitivity range of STIX, with a visibly high background, the same microflare was observed with NuSTAR (see Figure 1.4). Therefore, using Sunxspex, we perform the first STIX-NuSTAR simultaneous spectral fitting (Section 6.3.1).

Chapter 7 describes the nested sampling algorithm (Skilling 2004) which provides a method to obtain the relative probability of one model fit to data compared to another and has been made available in Sunxspex through the Python package Nestle[1]. In Section 7.3, we investigate three model fits to a microflare spectrum (SOL2021-11-17T21:14) observed on 2021 November 17. One model is composed of two thermal sub-models, one of three thermal sub-models, and the final one is made of two thermal sub-models with the thick target. The first model (two thermal sub-models) is taken to be the most basic form to fit the microflare spectrum as one thermal sub-model is derived from the pre-flare with the other fitting the microflare emission. The latter two models include the same two thermal sub-models but with an extra sub-model (one thermal, one non-thermal) in order to fit an observed count residual excess present in the two thermal sub-model fit case. We find the model

[1] http://kylebarbary.com/nestle/

that includes the thick target is almost 10^{12} times more probable to fit the data than the two thermal scenario and ~300 times more likely than the three thermal case.

Chapter 8 presents two microflares observed on 2020 January 30 from AR SPoCA 23783, three microflares on 2021 November 17 from AR SPoCA 26190, and one jet on 2021 November 22 north of AR SPoCA 26196. From the 2020 January microflares (Section 8.1), we obtain evidence of temperatures up to 10 MK being reached from the complex microflaring structures with potential direct evidence for non-thermal emission. The microflares from 2021 November (Section 8.2.2) also reach high temperatures; however, despite their similar appearance in their X-ray time profiles, the first microflare appears to not have non-thermal signatures and reaches hotter temperatures than the other two. The first microflare also shows a different response in the SDO/AIA EUV channels compared to the following two microflares. Finally, Section 8.2.3 presents a jet which reaches a temperature comparable to those in quiescent, non-flaring ARs even though it is far removed from the closest active region.

Throughout this thesis, we have observed that weak X-ray microflares do exhibit behaviour commonly seen in larger flares. Chapter 3, 4, and 8 show that microflares of energy ~10^{26}–10^{28} erg heat material to temperatures up to 10 MK, much hotter than the surrounding AR, with the higher energy emission peaking earlier and being more impulsive than the lower energy emission.

Furthermore, we find that some of the microflares show very strong evidence for the presence of non-thermal emission (thought to be produced from flare loop-top accelerated electrons interacting with the dense chromosphere at the loop footpoints) during their impulsive phase. The non-thermal fits provide an energy that is comparable to the required thermal energy for the heated material (e.g., microflare 3 from Chapter 3 and the microflare studied in Chapter 7), an observation well documented in larger flares (Hannah et al. 2008, 2011; Warmuth & Mann 2016). Non-thermal emission may still be present in the other investigated microflares but be below the detection threshold either due to insufficient livetime and/or complex surrounding emission; however, when upper non-thermal limits have been investigated for a selected microflare we find non-thermal electron distributions consistent with a null detection which are still able to provide the required thermal energy (Section 4.4).

In addition, where it is possible to separate a microflare into different phases to analyse the evolving loop densities and temperatures, we find values consistent with chromospheric evaporation (Section 3.4.1) which is another phenomenon observed in larger flares caused by the heated material at the loop footpoints expanding into the loop (Fletcher et al. 2011). These findings are compatible with the description of the 'standard' flare as presented in Chapter 1 (Section 1.2), providing evidence the X-ray microflares observed at these weak scales are of the same population as their larger counterparts.

9.1.1 Flare Trends and Future Modelling

As discussed in Section 1.2, flares span over orders of magnitudes in size and emission; however, they all appear to undergo the same general behaviour. The 'standard' solar flare is described as having three phases; a pre-flare phase, an impulsive phase, and a gradual phase. During the pre-flare phase the UV/EUV and SXR emission starts to increase before the impulsive phase occurs which is characterised by a sharp increase in emission across all wavelengths. This phase hosts the impulsive time profile of the high energy non-thermal emission (HXRs and microwaves) which is caused by the flare accelerated electrons interacting with the denser chromosphere. Throughout the impulsive phase the SXR and UV emission rise to a peak until the gradual phase (Priest & Forbes 2002; Fletcher et al. 2011; Benz 2017).

From this simple scenario we find that a charactieristic of flare evolution is that the high energy HXR emission should peak earlier and be more impulsive than the lower energy EUV and SXR emission. Although the X-ray microflares presented in Chapter 3, 4, 7, and 8 are very weak we still find that the higher energy emission peaks earlier and appears more impulsive than the lower energy emission. However, this could also be produced by very hot material being present earlier in the flaring process (Caspi & Lin 2010). Therefore, using NuSTAR's spectroscopic capabilities, we find the higher energy emission produced during the earlier microflare stages is likely due to non-thermal processes where this was able to be investigated.

It is difficult to conclude on the location of the HXR emission, which may be expected to be found at the loop footpoints, due to the physical size of the flaring loops and NuSTAR's spatial resolution. However, we do find the microflares that release the most energy tend to have more complex structures in EUV (e.g., microflare 3 and 10 in Chapter 3).

We are able to determine the energies associated with each microflare from spectral fitting NuSTAR data (Chapter 5, 6). We obtain the thermal energy release using Equation 1.15 and, if there is strong evidence of non-thermal emission, we can also calculate the power input of the accelerated electron distribution using Equation 1.31 or Equation 1.33. Commonly, the non-thermal energy is found to exceed the required flare thermal energy (Aschwanden et al. 2017); however, when investigated, we find the non-thermal energies to be comparable with the thermal energy obtained. This could be due the thermal energies being conservative upper limits since the volume estimates were also upper limits.

In addition to non-thermal powers being calculated from the spectral fitting results, the fitted non-thermal parameters (e.g., the total electron flux $F(E)$, low-energy cut-off E_C, and electron spectral index δ of the thick target model described in Section 1.4.5) can be used in elaborate flare simulations. Flare simulations have been performed where material is heated with non-thermal beams of electrons

in order to investigate the observational affects of such a scenario (Reep et al. 2020; Kerr et al. 2021). Moreover, data from multiple instruments can be forward-modelled by simulations in order to obtain parameters that are likely to produce the solar flare being observed (Kerr et al. 2020). Many simulations have been performed with focus on larger flares; however, with the non-thermal parameters obtained and yet to be obtained from the microflares presented in this thesis, more attention can be given to simulating X-ray microflares at the weakest scales to help investigate similarities and differences with their larger counterparts.

9.2 Future Work

The Sunxspex spectral fitting package has many capabilities already; however, there are still improvements to be made. From Chapter 6, we note that there are model corrections that can be made during the spectral fitting that are not yet implemented in Sunxspex, for example no pile-up correction exists. We also note the ongoing effort to include support for more instruments to which high-energy spectral fitting applies, such as XSM (Vadawale et al. 2014; Shanmugam et al. 2020) and GMB (Meegan et al. 2009).

The user interface for Sunxspex is to be updated for it to be more intuitive while default values, like default model parameter values, should be made more dynamic and adapt to the loaded spectra. More models will be made available in Sunxspex, for example a multi-thermal model which builds upon the already existent isothermal model. Additionally, Sunxspex still has a dependency on IDL SSW for the isothermal model calculation; this will be removed with all model calculations being performed purely in the Sunxspex package.

More powerful parameter search algorithms will also be made available in Sunxspex to ensure greater confidence that global, and not local, optimal fit values are being found. In relation to the parameter exploration, we also aim to improve the performance of Sunxspex with speed increases in runtime using optimised code and parallel processes.

Chapter 7 presents the nested sampling algorithm which is made available in Sunxspex and allows comparison between different model fits to data. Further work on Sunxspex will include adding more support for the nested sampling result, such as methods to display and customise corner plots. Other nested sampling based algorithms will become available in Sunxspex such as *diffusive* nested sampling Brewer et al. (2011) which employs techniques to avoid excluding the most probable likelihood region during the normal nested sampling iterations. In addition, nested sampling will be applied to any event where there is ambiguity in the model that best represents the data, for example both microflares from Section 8.1 and microflare SOL2021-11-17T21:25 (the last microflare from Section 8.2.2).

In Chapter 8, we discuss a total of six solar events, five microflares and one jet. Future work on these events will involve calculating each DEM to better describe the temperature distribution in the emitting sources and performing MCMC analysis on the spectral fitting results to obtain better uncertainty estimates. From the results of this analysis, the energetics of the events can then be obtained with greater confidence in their values. Wavelet analysis will be performed on the chromospheric emission at the footpoints of the three microflares observed on 2021 November 17 (Section 8.2.2) to investigate any microflare introduced oscillations. In addition, the microflares and jet observed during the 2021 November NuSTAR solar campaign were also co-observed by other instruments (see Table 8.1). Therefore, all observations of each event will be studied in order to obtain a more comprehensive understanding.

Moreover, the role photospheric magnetic flux activity may play in microflare produciton, as well as any other potential triggering mechanism, will be investigated further for all events. Chromospheric emission (e.g., the SDO/AIA 304 Å channel) from the microflaring regions will also be explored to identify possible corresponding features between the photosphere and corona. In addition, observations of microflares similar to the ones discussed with instruments like IRIS will provide insight into any heating events in the chromosphere that relate to microflares, again helping to connect features viewed throughout the solar atmosphere.

There are possibly over seventeen events from the 2021 November 17–22 NuSTAR solar campaign, with at least twelve microflares. As such, all of these events will be identified and studied across all available wavelength ranges. This will greatly increase the number of NuSTAR analysed solar phenomena, allowing a more statistical view of these events to be obtained.

Bibliography

Antonucci E., Gabriel A. H., Dennis B. R., 1984, The Astrophysical Journal, 287, 917 36

Arnaud K. A., 1996, Astronomical Data Analysis Software and Systems V, A.S.P. Conference Series, 1996, George H. Jacoby and Jeannette Barnes, Eds., p. 17.. Astronomical Society of the Pacific Conference Series Vol. 101, ASP 49, 85, 114

Aschwanden M. J., 2004, The Astrophysical Journal, 608, 554 x, 4

Aschwanden M. J., Nightingale R. W., Tarbell T. D., Wolfson C. J., 2000, The Astrophysical Journal, 535, 1027 7

Aschwanden M. J., Boerner P., Ryan D., Caspi A., McTiernan J. M., Warren H. P., 2015, The Astrophysical Journal, 802, 53 19

Aschwanden M. J., et al., 2017, The Astrophysical Journal, 836, 17 190

Asgari-Targhi M., van Ballegooijen A. A., Davey A. R., 2019, The Astrophysical Journal, 881, 107 73

Ashton G., et al., 2022, Nature Reviews Methods Primers, 2, 39 146, 152

Athiray P. S., et al., 2020, The Astrophysical Journal, 891, 78 95

Battaglia A. F., et al., 2021, Astronomy & Astrophysics, 656, A4 x, 8, 9, 140

Benz A. O., 2017, Living Reviews in Solar Physics, 14, 2 3, 5, 190

Berger T. E., De Pontieu B., Fletcher L., Schrijver C. J., Tarbell T. D., Title A. M., 1999, Solar Physics, 190, 409 6

Bertero M., Mol C. D., Pike E. R., 1985, Inverse Problems, 1, 301 100

Bhalerao V., 2012, PhD thesis, California Institute of Technology 28, 31, 32, 34

Blackburn J. K., 1995, Astronomical Data Analysis Software and Systems IV, FTOOLS: A FITS Data Processing and Analysis Software Package. Astronomical Society of the Pacific Conference Series Vol. 77 115

Boerner P., et al., 2012, Solar Physics, 275, 41 20, 41

Brewer B. J., Pártay L. B., Csányi G., 2011, Statistics and Computing, 21, 649 151, 161, 191

Brown J. C., 1971, Solar Physics, 18, 489 22, 23, 24, 25, 56

Brown J. C., Emslie A. G., Holman G. D., Johns-Krull C. M., Kontar E. P., Lin R. P., Massone A. M., Piana M., 2006, The Astrophysical Journal, 643, 523 100

Carrington R. C., 1859, Monthly Notices of the Royal Astronomical Society, 20, 13 3

Carroll B. W., Ostlie D. A., 2014, An Introduction to Modern Astrophysics, pearson new international, 2nd edn. Pearson Education Limited, Essex 1, 2, 14, 17

Cash W., 1979, The Astrophysical Journal, 228, 939 49, 104

Caspi A., Lin R. P., 2010, The Astrophysical Journal, 725, L161 5, 113, 190

Charbonneau P., 2010, Living Reviews in Solar Physics, 7, 91 6

Chen B., Battaglia M., Krucker S., Reeves K. K., Glesener L., 2021, The Astrophysical Journal, 908, L55 118

Chitta L. P., et al., 2017a, The Astrophysical Journal Supplement Series, 229, 4 73, 74, 76, 78

Chitta L. P., Peter H., Young P. R., Huang Y.-M., 2017b, Astronomy & Astrophysics, 605, A49 73

Chitta L. P., Peter H., Solanki S. K., 2018, Astronomy & Astrophysics, 615, L9 73, 76, 78

Chitta L. P., Sukarmadji A. R. C., Rouppe van der Voort L., Peter H., 2019, Astronomy & Astrophysics, 623, A176 73

Chitta L. P., Peter H., Priest E. R., Solanki S. K., 2020, Astronomy & Astrophysics, 644, A130 73, 74, 76, 78, 79, 187

Chitta L. P., Peter H., Young P. R., 2021, Astronomy & Astrophysics, 647, A159 7

Christe S., Hannah I. G., Krucker S., McTiernan J., Lin R. P., 2008, The Astrophysical Journal, 677, 1385 8

Cirtain J. W., et al., 2013, Nature, 493, 501 87, 96, 187

Cooper K., Hannah I. G., Grefenstette B. W., Glesener L., Krucker S., Hudson H. S., White S. M., Smith D. M., 2020, The Astrophysical Journal Letters, 893, L40 9, 38, 44, 52, 61, 62, 79, 82, 83, 84, 86, 94, 95, 115, 153, 186, 187

Cooper K., et al., 2021, Monthly Notices of the Royal Astronomical Society, 507, 3936 10, 44, 46, 51, 55, 58, 62, 64, 65, 69, 70, 74, 76, 80, 81, 125, 126, 132, 143, 186, 187, 188

Craig I. J. D., Brown J. C., 1976, Astronomy and Astrophysics, 49, 239 18

Craig I. J. D., Brown J. E., 1977, Nature, 267, 2 19, 92

Crosby N. B., Aschwanden M. J., Dennis B. R., 1993, Solar Physics, 143, 275 6, 7, 97

Culhane J. L., et al., 2007, Solar Physics, 243, 19 42, 183

De Pontieu B., et al., 2014, Solar Physics, 289, 2733 183

Del Zanna G., 2013, Astronomy & Astrophysics, 558, A73 42

Del Zanna G., Dere K. P., Young P. R., Landi E., 2021, The Astrophysical Journal, 909, 38 12

Dere K. P., Landi E., Mason H. E., Monsignori Fossi B. C., Young P. R., 1997, Astronomy and Astrophysics Supplement Series, 125, 149 12, 38, 53

Doe S., et al., 2007, Astronomical Data Analysis Software and Systems XVI, Developing Sherpa with Python. Astronomical Society of the Pacific Conference Series Vol. 376 116

Duncan J., et al., 2021, The Astrophysical Journal, 908, 29 ix, xi, 8, 9, 35, 36, 42, 49, 77, 110, 115, 128, 129, 130, 131, 144, 166, 187, 188

Ellingson S., Clarke T., Cohen A., Craig J., Kassim N., Pihlstrom Y., Rickard L., Taylor G., 2009, Proceedings of the IEEE, 97, 1421 185

Emslie A. G., Sturrock P. A., 1982, Solar Physics, 80, 99 4

Emslie G. A., Phillips K. J. H., Dennis B. R., 1986, Solar Physics, 103, 89 36

Feldman U., Mandelbaum P., Seely J. F., Doschek G. A., Gursky H., 1992, The Astrophysical Journal Supplement Series, 81, 387 16, 49

Fletcher L., Hudson H. S., 2008, The Astrophysical Journal, 675, 1645 4

Fletcher L., Hannah I. G., Hudson H. S., Metcalf T. R., 2007, The Astrophysical Journal, 656, 1187 xi, 39, 134, 135, 136, 144, 188

Fletcher L., et al., 2011, Space Science Reviews, 159, 19 4, 5, 7, 56, 189, 190

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, Publications of the Astronomical Society of the Pacific, 125, 306 106, 119

Gallagher P. T., Mathioudakis M., Keenan F. P., Phillips K. J. H., Tsinganos K., 1999, The Astrophysical Journal, 524, L133 53

Gao F., Han L., 2012, Computational Optimization and Applications, 51, 259 105

Gburek S., et al., 2011, Solar System Research, 45, 182 7

Gehrels N., 1986, The Astrophysical Journal, 303, 336 90

Glesener L., Krucker S., Hannah I. G., Hudson H., Grefenstette B. W., White S. M., Smith D. M., Marsh A. J., 2017, The Astrophysical Journal, 845, 122 9, 52, 77, 85, 86, 96, 110, 187

Glesener L., et al., 2020, The Astrophysical Journal, 891, L34 ix, xi, 8, 9, 42, 57, 72, 78, 91, 115, 127, 128, 129, 143, 147, 188

Golub L., et al., 2007, Solar Physics, 243, 63 164

Goodman J., Weare J., 2010, Communications in Applied Mathematics and Computational Science, 5, 65 119, 124

Grefenstette B. W., et al., 2016, The Astrophysical Journal, 826, 20 9, 27, 29, 31, 32, 34, 44, 45, 84, 144

Hannah I. G., Kontar E. P., 2012, Astronomy & Astrophysics, 539, A146 19, 92, 93, 94

Hannah I. G., Kontar E. P., 2013, Astronomy & Astrophysics, 553, A10 92

Hannah I. G., Christe S., Krucker S., Hurford G. J., Hudson H. S., Lin R. P., 2008, The Astrophysical Journal, 677, 704 8, 9, 19, 25, 53, 85, 91, 113, 189

Hannah I. G., Hudson H. S., Battaglia M., Christe S., Kašparová J., Krucker S., Kundu M. R., Veronig A., 2011, Space Science Reviews, 159, 263 ix, 7, 8, 25, 49, 97, 189

Hannah I. G., et al., 2016, The Astrophysical Journal, 820, L14 9, 27, 42, 44, 96

Hannah I. G., Kleint L., Krucker S., Grefenstette B. W., Glesener L., Hudson H. S., White S. M., Smith D. M., 2019, The Astrophysical Journal, 881, 109 9, 42, 52, 77, 85, 86, 87, 94, 96, 110, 187

Hansen P. C., 1992, Inverse Problems, 8, 849 93

Harrison F. A., et al., 2013, The Astrophysical Journal, 770, 103 x, 1, 8, 27, 28, 29, 30, 32, 186

Hastings W. K., 1970, Biometrika, 57, 13 106

Haug E., 1997, Astronomy and Astrophysics, 326, 417 17

Holman G. D., et al., 2011, Space Science Reviews, 159, 107 17

Hudson H., 1991, Solar Physics, 133, 357 6, 7, 97

Humphrey P. J., Liu W., Buote D. A., 2009, The Astrophysical Journal, 693, 822 103, 104, 114

Hurford G., Curtis D., 2002, Solar Physics, 210, 101 39

Hurford G., et al., 2002, Solar Physics, 210, 61 39

Ireland J., Tolbert A. K., Schwartz R. A., Holman G. D., Dennis B. R., 2013, The Astrophysical Journal, 769, 89 108, 109, 115

Ishikawa S.-n., Glesener L., Krucker S., Christe S., Buitrago-Casas J. C., Narukage N., Vievering J., 2017, Nature Astronomy, 1, 771 95

Jefferies J. T., Orrall F. Q., Zirker J. B., 1972, Solar Physics, 22, 10 19, 92

Judge P., 2020, The Sun, A Very Short Introduction, 1st edn. Oxford University Press, Oxford 1

Judge P. G., Hubeny V., Brown J. C., 1997, The Astrophysical Journal, 475, 275 92

Kaastra J. S., Mewe R., Nieuwenhuijzen H., 1996, SPEX: A New Code for Spectral Analysis of X \& UV Spectra. UV and X-ray Spectroscopy of Astrophysical and Laboratory Plasmas Vol. 101 116

Kass R. E., Raftery A. E., 1995, Journal of the American Statistical Association, 90, 773 147, 149, 156, 159

Kerr G. S., Allred J. C., Polito V., 2020, The Astrophysical Journal, 900, 18 191

Kerr G. S., Xu Y., Allred J. C., Polito V., Sadykov V. M., Huang N., Wang H., 2021, The Astrophysical Journal, 912, 153 191

Kester D., Mueller M., 2021, Astronomy and Computing, 37, 100503 146

Knuth K. H., Habeck M., Malakar N. K., Mubeen A. M., Placek B., 2015, Digital Signal Processing, 47, 50 146

Kontar E. P., Piana M., Massone A. M., Emslie A. G., Brown J. C., 2004, Solar Physics, 225, 293 100

Kosugi T., et al., 2007, Solar Physics, 243, 3 42, 164

Krucker S., Christe S., Lin R. P., Hurford G. J., Schwartz R. A., 2002, Solar Physics, 210, 445 49

Krucker S., et al., 2020, Astronomy & Astrophysics, 642, A15 39

Kuhar M., et al., 2017, The Astrophysical Journal, 835, 6 6, 9, 27, 113

Kuhar M., Krucker S., Glesener L., Hannah I. G., Grefenstette B. W., Smith D. M., Hudson H. S., White S. M., 2018, The Astrophysical Journal, 856, L32 9, 182

Kurt W., 2019, Bayesian Statistics the Fun Way, 1st edn. No Starch Press, Inc., San Francisco 102, 148

Landi E., Young P. R., Dere K. P., Del Zanna G., Mason H. E., 2013, The Astrophysical Journal, 763, 86 38, 53

Lemen J. R., et al., 2012, Solar Physics, 275, 17 2, 12, 40, 42, 47

Lin R. P., 1974, Space Science Reviews, 16, 189 25

Lin R. P., Schwartz R. A., Kane S. R., Pelling R. M., Hurley K. C., 1984, The Astrophysical Journal, 283, 421 7

Lin R. P., et al., 2002, Sol Phys, 210, 30 8, 39

Longair S. M., 2018, High Energy Astrophysics, 3rd edn. Cambridge University Press 15, 23

Loumou K., Hannah I. G., Hudson H. S., 2018, Astronomy & Astrophysics, 618, A9 x, 7

MacCombie W. J., Rust D. M., 1979, Solar Physics, 61, 69 5

Madsen K. K., et al., 2015, The Astrophysical Journal Supplement Series, 220, 8 30, 32, 34, 36, 48

Meegan C., et al., 2009, The Astrophysical Journal, 702, 791 145, 191

Milligan R. O., Gallagher P. T., Mathioudakis M., Bloomfield D. S., Keenan F. P., Schwartz R. A., 2006, The Astrophysical Journal, 638, L117 5

Mitra-Kraev U., Del Zanna G., 2019, Astronomy & Astrophysics, 628, A134 77, 78, 187

Mondal B., et al., 2021, The Astrophysical Journal, 920, 4 115, 116

Moreno Cárdenas F., Cristancho Sánchez S., Vargas Domínguez S., 2016, Advances in Space Research, 57, 257 3

Morozov V. A., 1967, Dokladi Akademii Nauk SSSR, 175, 1225 93

Mossman J. E., 1989, Quarterly J. R. Astr. Soc., 30, 59 2

Mulay S. M., Matthews S., Hasegawa T., Del Zanna G., Mason H., Shimizu T., 2018, Solar Physics, 293, 160 113

Müller D., et al., 2020, Astronomy & Astrophysics, 642, A1 39

Musset S., et al., 2019, in Siegmund O. H., ed., UV, X-Ray, and Gamma-Ray Space Instrumentation for Astronomy XXI Vol. 11118, Ghost-Ray Reduction and Early Results from the Third FOXSI Sounding Rocket Flight. SPIE, San Diego, United States, pp 321–336, doi:10.1117/12.2530029 44

Neupert W. M., 1968, The Astrophysical Journal, 153, L59 5

Norton A. A., et al., 2006, Solar Physics, 239, 69 43

O'Dwyer B., Del Zanna G., Mason H. E., Weber M. A., Tripathi D., 2010, Astronomy and Astrophysics, 521, A21 41, 42, 47

Parker E. N., 1988, The Astrophysical Journal, 330, 474 7

Parnell C. E., Jupp P. E., 2000, The Astrophysical Journal, 529, 554 7

Perri M., et al., 2017, The NuSTAR Data Analysis Software Guide 31, 32, 33

Pesnell W. D., Thompson B. J., Chamberlin P. C., 2012, Solar Physics, 275, 3 40, 75

Phillips K. J. H., Feldman U., Landi E., 2012, Ultraviolet and X-ray Spectroscopy of the Solar Atmosphere. No. 44 in Cambridge Astrophysics Series, Cambridge University Press, New York 3, 6, 10, 13, 17, 18, 37, 92

Piana M., 1994, Astronomy and Astrophysics, 288, 949 100

Piana M., Brown J. C., 1998, Astronomy and Astrophysics Supplement Series, 132, 291 100

Press W. H., ed. 1996, FORTRAN Numerical Recipes, second edn. Cambridge University Press, Cambridge [England] ; New York 108

Priest E., Forbes T., 2002, The Astronomy and Astrophysics Review, 10, 313 4, 5, 190

Priest E. R., Chitta L. P., Syntelis P., 2018, The Astrophysical Journal, 862, L24 73

Reale F., McTiernan J. M., Testa P., 2009, The Astrophysical Journal, 704, L58 95

Reep J. W., Antolin P., Bradshaw S. J., 2020, The Astrophysical Journal, 890, 100 191

Richardson W. H., 1972, Journal of the Optical Society of America, 62, 55 11, 48, 83

Saqri J., et al., 2022, Astronomy & Astrophysics, 659, A52 113

Schmelz J. T., et al., 2009, The Astrophysical Journal, 704, 863 95

Schou J., et al., 2012, Solar Physics, 275, 229 12, 40, 43

Schwartz R. A., Csillaghy A., Tolbert A. K., Hurford G. J., McTIERNAN J., Zarro D., 2002, Solar Physics, 210, 165 113

Shanmugam M., et al., 2020, Current Science, 118, 45 7, 145, 191

Simões P. J. A., Hudson H. S., Fletcher L., 2015, Solar Physics, 290, 3625 38

Simões P. J. A., Reid H. A. S., Milligan R. O., Fletcher L., 2019, The Astrophysical Journal, 870, 114 40

Sivia D. S., Skilling J., 2006, Data Analysis: A Bayesian Tutorial, second edn. Oxford Science Publications, Oxford University Press, Oxford ; New York 100, 101, 102, 103, 104, 107, 148, 149, 151, 152, 153

Skilling J., 2004, in AIP Conference Proceedings. AIP, Garching (Germany), pp 395–405, doi:10.1063/1.1835238 145, 146, 147, 150, 153, 161, 188

Skilling J., 2006, Bayesian Analysis, 1 146, 151, 157

Smith H. J., Smith E. v. P., 1963, Solar Flares, 1st edn. The MacMillan Company, New York 2

Stewart B., 1861, in Proceedings of Royal Society of London. London, pp 423–430, doi:10.1098/rstl.1861.0023 3

Storn R., Price K., 1997, Journal of Global Optimization, 11, 341 145

Sylwester J., et al., 2012, The Astrophysical Journal, 751, 111 7

Syntelis P., Priest E. R., Chitta L. P., 2019, The Astrophysical Journal, 872, 32 73

Tandberg-Hanssen E., Emslie A. G., 2009, High-Energy Aspects of Solar Flares, 1st edn. Cambridge University Press, New York 3, 14, 15, 17, 18, 20, 21, 22, 23

Testa P., Reale F., 2020, The Astrophysical Journal, 902, 31 75, 78

Testa P., et al., 2014, Science, 346, 26 91

Testa P., Polito V., Pontieu B. D., 2020, The Astrophysical Journal, 889, 124 78

Tikhonov A. N., 1963, Dokladi Akademii Nauk SSSR, 151, 501 93

Torrence C., Compo G. P., 1998, Bulletin of the American Meteorological Society, 79, 61 179, 185

Vadawale S., et al., 2014, Advances in Space Research, 54, 2021 7, 145, 191

Vadawale S. V., et al., 2021a, The Astrophysical Journal Letters, 912, L12 7

Vadawale S. V., et al., 2021b, The Astrophysical Journal Letters, 912, L13 7

Vievering J. T., et al., 2021, The Astrophysical Journal, 913, 15 9

Warmuth A., Mann G., 2016, Astronomy & Astrophysics, 588, A115 8, 9, 189

Warren H. P., Winebarger A. R., Brooks D. H., 2012, The Astrophysical Journal, 759, 141 42, 54

Weisskopf M. C., Tananbaum H. D., Van Speybroeck L. P., O'Dell S. L., 2000, in Truemper J. E., Aschenbach B., eds, Astronomical Telescopes and Instrumentation. Munich, Germany, pp 2–16, doi:10.1117/12.391545 116

Woods T. N., et al., 2012, Solar Physics, 275, 115 40

Wright P. J., et al., 2017, The Astrophysical Journal, 844, 132 9, 24, 38, 42, 52, 77, 85, 88, 91, 96, 110, 115, 187

Zarro D. M., Lemen J. R., 1988, The Astrophysical Journal, 329, 456 5

Zirin H., 1966, The Solar Atmosphere, 1st edn. Blaisdale Publishing Company 3

Zouganelis I., et al., 2020, Astronomy & Astrophysics, 642, A3 39