Detecting non-thermal emission in a solar microflare using nested sampling

Kristopher Cooper [1,2]★ Iain G. Hannah [2] Lindsay Glesener [1] and Brian W. Grefenstette [3]

[1] *School of Physics & Astronomy, University of Minnesota Twin Cities, Minneapolis, MN 55455, USA*
[2] *School of Physics & Astronomy, University of Glasgow, University Avenue, Glasgow, G12 8QQ, UK*
[3] *Cahill Center for Astrophysics, California Institute of Technology, 1216 East California Boulevard, Pasadena, CA 91125, USA*

ABSTRACT

Microflares are energetically smaller versions of solar flares, demonstrating the same processes of plasma heating and particle acceleration. However, it remains unclear down to what energy scales this impulsive energy release continues, which has implications for how the solar atmosphere is heated. The heating and particle acceleration in microflares can be studied through their X-ray emission, finding predominantly thermal emission at lower energies; however, at higher energies it can be difficult to distinguish whether the emission is due to hotter plasma and/or accelerated electrons. We present the first application of nested sampling to solar flare X-ray spectra, an approach that provides a quantitative degree of confidence for one model over another. We analyse *Nuclear Spectroscopic Telescope Array* X-ray observations of a small active region microflare (A0.02 GOES/XRS class equivalent) that occurred on 2021 November 17, with a new PYTHON package for spectral fitting, SUNKIT-SPEX, to compute the parameter posterior distributions and the evidence of different models representing the higher energy emission as due to thermal or non-thermal sources. Calculating the Bayes factor, we show that there is significantly stronger evidence for the higher energy microflare emission to be produced by non-thermal emission from flare-accelerated electrons than by an additional hot thermal source. Qualitative confirmation of this non-thermal source is provided by the lack of hotter (10 MK) emission in Solar Dynamic Observatory's Atmospheric Imaging Assembly's extreme ultraviolet data. The nested sampling approach used in this paper has provided clear support for non-thermal emission at the level of 3×10^{24} erg s^{-1} in this tiny microflare.

Key words: methods: numerical – Sun: activity – Sun: corona – Sun: flares – Sun: X-rays, gamma-rays.

1 INTRODUCTION

Flares are explosive releases of energy that take place in active regions (ARs) thought to be produced and powered by magnetic reconnection of stressed magnetic fields in the solar atmosphere (Priest & Forbes 2002). The reconnection can release energies up to $\sim 10^{32}$ erg and readily drives mass flows, plasma heating, and particle acceleration (Fletcher et al. 2011; Benz 2017). Flares produce emission across the electromagnetic spectrum and are commonly characterized by their Geostationary Operational Environmental Satellite X-Ray Sensor (GOES/XRS) 1–8 Å soft X-ray (SXR) emission and thermal energy release. Flares with SXR emission $<10^{-6}$ W m^{-2} and energies between 10^{26} and 10^{28} erg are commonly referred to as *microflares* (Lin et al. 1984; Hannah et al. 2011).

Flares with energies of the order of $\sim 10^{24}$ erg, termed *nanoflares*, are proposed to occur everywhere on the solar disc and not just localized to ARs (Parker 1988). If microflares and nanoflares occur frequently enough, then they could provide the majority fraction of the energy heating the corona from the overall flare distribution; however, this relies on the same flaring mechanisms scaling down to the weaker events, such as non-thermal particle acceleration emission (Hudson 1991). It is often difficult to ascertain whether weak, sub-A

class microflares do indeed show positive evidence for these standard flare processes often observed in their larger counterparts.

Microflares have been extensively studied in X-rays with instruments such as the Reuven Ramaty High-Energy Solar Spectroscopic Imager (RHESSI; Lin et al. 2002), Chandrayaan-2's Solar X-ray Monitor (Vadawale et al. 2014; Shanmugam et al. 2020), and the *Nuclear Spectroscopic Telescope Array* (*NuSTAR*; Harrison et al. 2013). Statistical and individual studies have observed hard X-ray emission during microflares from non-thermal electron acceleration showing evidence that microflares down to GOES class A0.1 appear to show similar processes to their brighter counterparts (Christe et al. 2008; Hannah et al. 2008; Glesener et al. 2020; Cooper et al. 2021). However, as the microflares get weaker it becomes increasingly difficult to determine whether this flare behaviour continues to scale or has a lower limit. Sensitivity to these weak microflares is crucial in understanding whether the flare process scales across these different classifications.

NuSTAR is an astrophysical telescope capable of observing the Sun >2.5 keV (Grefenstette et al. 2016; Hannah et al. 2016). *NuSTAR* consists of two telescopes utilizing Wolter-I-type optics to focus X-rays on to two focal plane modules (FPMA and FPMB), each with a field of view (FOV) of 12 arcmin × 12 arcmin. Each FPM takes 2.5 ms to process a detected photon during which no other trigger can be recorded; the fraction of time *NuSTAR* spends open to detection during an observation is quantified by the livetime. Under quiescent

★ E-mail: kristophercooper95@gmail.com

conditions, *NuSTAR*'s livetime can reach up to 92 per cent (Paterson et al. 2023); however, the livetime drops to <16 per cent even with small microflares of A class or smaller (Cooper et al. 2021). This limited throughput, combined with flare X-ray spectra steeply falling with increasing energy, results in a low number of *NuSTAR* counts at higher energies.

NuSTAR has observed several microflares with energies from 10^{28} down to 10^{26} erg with GOES classes of B, A, and sub-A (Glesener et al. 2017; Wright et al. 2017; Hannah et al. 2019; Cooper et al. 2020, 2021; Duncan et al. 2021) and also quiescent Sun features outside ARs of energies down to 10^{26} erg (Kuhar et al. 2018; Paterson et al. 2023). Several of these studies have found consistency with the presence of non-thermal emission throughout the flare evolution, while Glesener et al. (2020) and Cooper et al. (2021) show convincing evidence of non-thermal emission in an A5.7 and an estimated A0.1 microflare, respectively.[1]

It becomes more difficult with weaker flares to confidently determine the presence of any high-energy emission and even harder to determine the mechanism behind the emission. In order to determine the nature of the higher energy emission observed by an instrument from much weaker microflares, a robust method needs to be utilized, which can investigate the difference between different model representations of the observed data. Nested sampling analysis (Skilling 2004, 2006) can be used for this purpose. Note that this analytical technique can also be applied in general cases (e.g. in larger flares) where ambiguity exists between model fits to data.

Nested sampling is a Bayesian technique that is used to map the posterior distribution and estimate a quantity called the *evidence*, which is the probability of an observation given the assumption of a model representation. The pieces of evidence of different hypotheses can then be used to compare different model fits to data, providing a level of confidence for one model over another. The nested sampling algorithm also accounts for different model parameters and different numbers of model parameters used between representations. There are many implementations of nested sampling (Kester & Mueller 2021; Ashton et al. 2022) and it is used extensively in many scientific fields (Knuth et al. 2015); however, it is not common in the solar physics community.

In this paper, we present the first use of nested sampling in the analysis of an AR *NuSTAR* X-ray microflare that took place on 2021 November 17 at ∼21:14 UT (SOL2021-11-17T21:14). In Section 2, we investigate the microflare evolution in time, space, and energy using *NuSTAR* and the Solar Dynamic Observatory's Atmospheric Imaging Assembly (SDO/AIA; Lemen et al. 2012) and Heliospheric and Magnetic Imager (SDO/HMI; Schou et al. 2012). We then describe the nested sampling algorithm and how it can be used to determine the most likely model representation to explain a given data set in Section 3. In Section 4, we then apply the nested sampling algorithm to the *NuSTAR* X-ray spectral analysis for the first time where we determine the most likely explanation for the microflare emission.

2 X-RAY MICROFLARE

A *NuSTAR* solar campaign was performed on 2021 November 17–22 where 9 h-long observations were made of ARs present on the Sun's disc. During the third observation, *NuSTAR* observed three microflares where the second one (microflare SOL2021-11-17T21:14)

---

[1]A *NuSTAR* solar observations overview is available at https://ianan.github.io/nsigh_all/.

is the subject of this paper; the other flares and observations will be investigated in future papers.

2.1 Time evolution

Microflare SOL2021-11-17T21:14 is clearly observed in X-rays with *NuSTAR* as shown in Fig. 1. The extreme ultraviolet (EUV) and X-ray time profiles from SDO/AIA and *NuSTAR*, respectively, are shown (left panel) where the SDO/AIA light curves are from a region around the microflare (see Fig. 2) and the *NuSTAR* time profiles are integrated over the entire AR area indicated by a black box in Fig. 1 (right panel). The X-ray image (right panel) shows the microflare in the top of *NuSTAR*'s FOV, inside the black box, as a relatively bright source compared to the more northern AR (SPoCA 26190; Verbeeck et al. 2014) core emission.

We find that the native SDO/AIA channel light curves in Fig. 1 (top of left panel) show little similarity with the X-ray time profiles (bottom of left panel) with the exception of the 94 Å channel that shows a slight positive correlation and a possible negative correlation with the 304 Å channel. The microflare only becomes apparent in the EUV when calculating the Fe XVIII proxy channel – a linear combination of the 94, 171, and 211 Å channels – as described in Del Zanna (2013). This channel has a significant response to thermal emission from material between 4 and 10 MK, which is a crucial range for *NuSTAR* microflares (Hannah et al. 2016, 2019; Wright et al. 2017; Cooper et al. 2020, 2021; Glesener et al. 2020; Duncan et al. 2021).

The Fe XVIII light curve shows very similar progression to the 2.5–6 keV FPMA + FPMB *NuSTAR* time profile indicating that this microflare heats a significant amount of material to between 4 and 10 MK. The higher energy X-ray range, 6–10 keV, shows a more impulsive time profile that peaks earlier, indicative of hot plasma being present in the early stages of the microflare and/or non-thermal emission from an initially accelerated electron distribution.

Pre-flare and microflare times are defined from the *NuSTAR* time profiles and shown as red and blue shaded regions, respectively, in Fig. 1. The microflare time is chosen based on the more impulsive 6–10 keV *NuSTAR* light curve to better investigate the emission processes in the initial flaring stages.

The microflare, along with many events *NuSTAR* observed in the observation campaign, is obvious in the X-ray regime where it is dominant across the whole FOV; however, as indicated by the SDO/AIA light curves in Fig. 1, it is difficult to find corroborating evidence in the native EUV channels even when locating the microflare's position. Therefore, the EUV microflaring structure can only be reliably determined with information from *NuSTAR* and its comparison to Fe XVIII images.

2.2 Spatial evolution

Fig. 2 shows EUV images integrated over the microflare time. Comparable in size to the black box shown in the *NuSTAR* FOV image in Fig. 1 (right panel), the left panel of Fig. 2 shows the SDO/AIA 94 Å channel emission of the full AR with the microflare identified within a black box. The black box region identified to host the microflare is the area used to depict the other panels.

The Fe XVIII proxy channel and, to a lesser extent, the 94 Å channel show a loop structure, while the other channels do not show any corresponding activity except at the apparent loop footpoints. Footpoint activity can be seen through all EUV channels; however, the footpoints do not appear to produce the same response at all levels
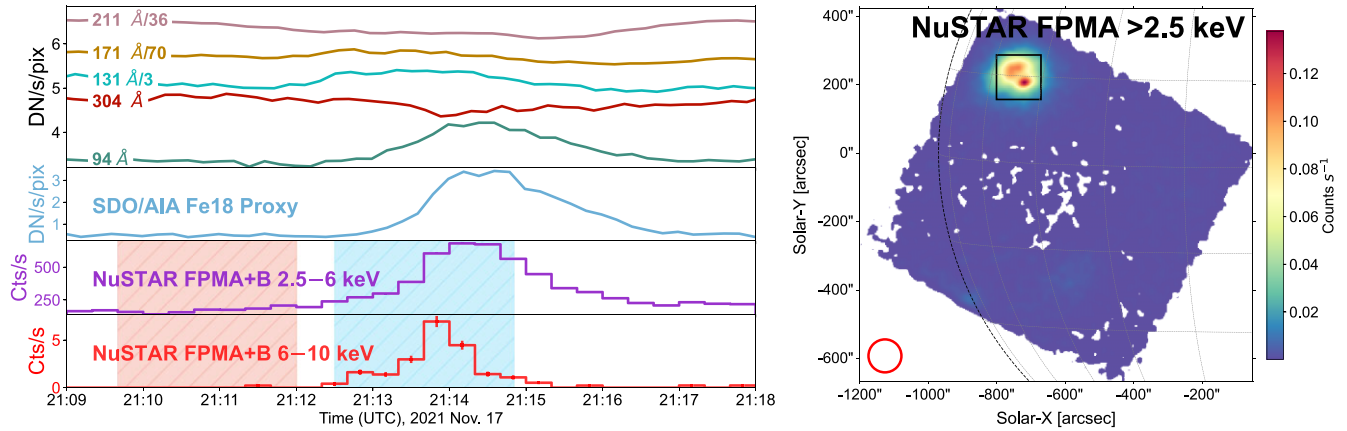
Figure 1. SDO/AIA, including Fe XVIII proxy, and *NuSTAR* livetime-corrected light curves over the time of microflare SOL2021-11-17T21:14 (left panel). The SDO/AIA time profiles are averaged over the black box region shown in Fig. 2 and a scale factor (for plotting purposes) has been applied, as shown, to several channels. The area used to create the *NuSTAR* light curves is shown as a black box across AR SPoCA 26190 in the *NuSTAR* >2.5 keV integrated, livetime-corrected FOV (right panel). Pre-flare (21:09:40–21:12:00 UT) and microflare (21:12:30–21:14:50 UT) times are indicated by the shaded red and blue areas in the *NuSTAR* light-curve panels, respectively. The red circle with a radius of 48 arcsec shown in the right panel indicates the region size used to produce spectral fit profiles for the *NuSTAR* spectra.
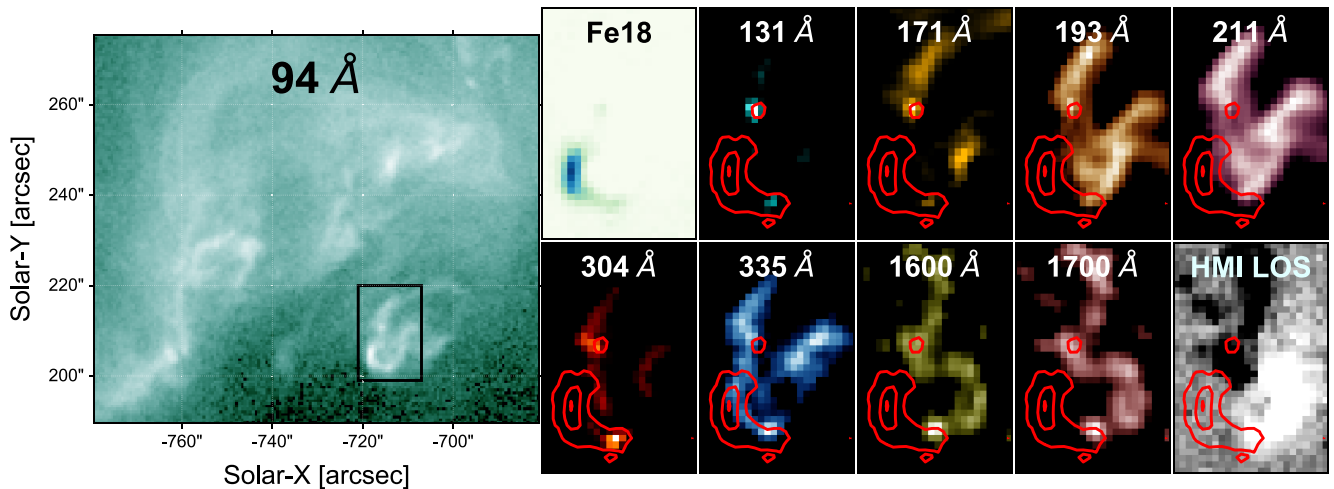
Figure 2. SDO/AIA native channels and Fe XVIII proxy are averaged over the microflare time (blue shaded region, Fig. 1), while the SDO/HMI line-of-sight magnetogram is taken from the start of the time range. The left panel shows the 94 Å channel emission over the full AR with an FOV slightly smaller than the boxed region in the *NuSTAR* image from Fig. 1. A black box indicates the location of the microflare and is used to create the zoomed-in images for the other emission channels (right panels). The Fe XVIII proxy microflaring loop emission contours at 5, 50, and 95 per cent of the maximum are overlaid on the zoomed panels and are shown in red.

in the solar atmosphere. The northern footpoint appears brighter in the upper transition region as shown by the relative brightness in the 131 and 171 Å channels, whereas the southern footpoint shows a greater response in the chromosphere and photosphere as seen in the 304, 1600, and 1700 Å channels (Lemen et al. 2012). From the SDO/HMI panel in Fig. 2, we find that the loop footpoints are anchored in opposite polarity regions.

Zoomed images with *NuSTAR* emission only showed a single circular shape to co-align with the EUV loop structure for the 2.5–6 and 6–10 keV emission. The lack of distinctive shape in X-rays may be expected due to the very few high-energy microflare counts and the fact that the angular extent of the microflare is ∼10 arcsec, which is considerably smaller than *NuSTAR*'s angular resolution with a full width at half-maximum and half-power diameter of 18 and ∼60 arcsec, respectively (Harrison et al. 2013).

From the EUV images shown in Fig. 2, we are able to obtain a microflaring loop volume by modelling it as a half torus. We estimate an upper limit for the distance between the footpoints and the diameter of the loop to be approximately 9 and 2 arcsec (6.5×10^8 and 1.5×10^8 cm), respectively. Therefore, the volume obtained from the geometry as viewed in EUV is 1.7×10^{25} cm³.

2.3 Spectral evolution

Fig. 3 shows the spectral fitting analysis (see Section 2.3.1) of the *NuSTAR* FPMA and FPMB grade 0 (single pixel) microflare emission on the pre-flare and microflare time using SUNKIT-SPEX,[2]

---

[2]Formerly, SUNXSPEX, a PYTHON spectral fitting tool: https://github.com/sunpy/sunkit-spex.
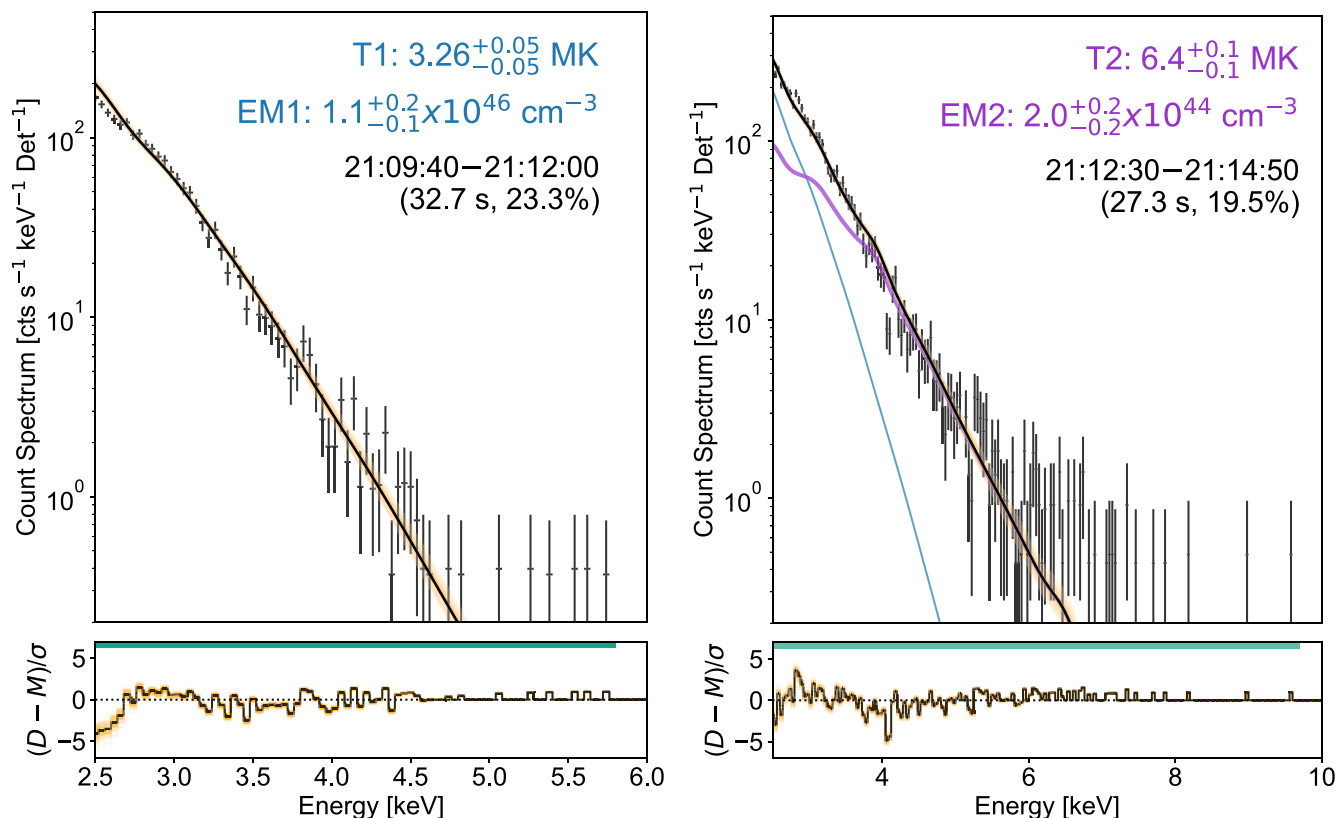
Figure 3. Simultaneous thermal model fits, using SUNKIT-SPEX, of *NuSTAR*'s FPMA and FPMB emission spectra during the pre-flare (21:09:40–21:12:00 UT, left panel) and microflare (21:12:30–21:14:50 UT, right panel) times, which are shown as shaded red and blue regions in Fig. 1, respectively. The pre-flare spectra are fitted with one thermal model (black) and then used as a fixed component for the microflare spectra (blue). The microflare spectra are fitted with an additional thermal model (purple) to account for the flaring emission with the overall model being shown in black. The spectra from both times are fitted over the energy ranges indicated by the horizontal green line in the residual panels. The temperatures and emission measures for each model fit are shown in their respective model colours where the error ranges are calculated from MCMC analysis and report the 1σ equivalent uncertainty. Additionally, 100 random model samples from the MCMC analysis are shown in orange. The relevant time ranges, effective exposures, and livetimes are also shown in each panel. The microflare fit in the right panel is also referred to as the *null scenario* in relation to model comparison.

which utilizes the EMCEE package (Foreman-Mackey et al. 2013). We find that the pre-flare time (Fig. 3, left panel) is represented well with an isothermal model of temperature 3.3 MK and emission measure 1 $\times 10^{46}$ cm^{-3} from the Markov chain Monte Carlo (MCMC) analysis, a common result from previous *NuSTAR* quiescent AR and pre-flare spectral fitting (Glesener et al. 2017; Wright et al. 2017; Hannah et al. 2019; Cooper et al. 2021). The isothermal model available in SUNKIT-SPEX is a PYTHON version of the IDL function f_vth.pro (Schwartz et al. 2002) and uses CHIANTI V7.1 (Dere et al. 1997; Landi et al. 2013) with coronal abundances (Feldman et al. 1992; Landi, Feldman & Dere 2002).

The pre-flare thermal model is then used as a fixed component when fitting the microflare time with an additional thermal model representing the flaring emission (Fig. 3, right panel). This model fit provides a temperature and emission measure in a sensible range for previous *NuSTAR* microflares, between 4 and 10 MK, as discussed in Section 2.1; however, it is clear that this model does not provide a satisfactory representation of the observed spectra. The residuals of the microflare fit in Fig. 3 indicate a poor fit between 2.5 and 6 keV and a clear count excess above the model at energies >6 keV. Therefore, we find sufficient reason to include an additional model to the microflare fit in order to accurately represent the observed spectra.

Two model candidates are tested to fit the excess: a *thermal scenario*, which includes an additional thermal model, and a *non-thermal scenario*, which involves the thick-target model (Brown 1971; Holman et al. 2011). The thick-target model assumes a power-law distribution of electrons of spectral index δ above a low-energy cut-off E_c that lose all their energy through Coulomb collisions. These hypotheses represent either hotter plasma or non-thermal emission during the impulsive phase of the microflare. In relation to these two excess additions, the microflare fit in Fig. 3 is termed the *null scenario* since the excess is not represented. Representations of the photon models utilized in each scenario are shown in Fig. A1. Note that unlike in the spectral analysis of brighter *NuSTAR* microflares, no gain correction (Duncan et al. 2021) is required due to the relatively high livetime of microflare SOL2021-11-17T21:14.

Fig. 4 shows, in red, the thermal scenario fit (left panel) and the non-thermal scenario fit (right panel) to the microflare spectra. We find that both scenarios resolve the poor fit to the 2.5–6 keV range and represent the count excess well >6 keV.

The thermally fitted excess representation proposes the presence of a relatively weak emitting source with a temperature of ~10 MK, whereas the non-thermal case suggests thick-target emission with a δ of ~8 and an E_c of ~6 keV. We note that the non-
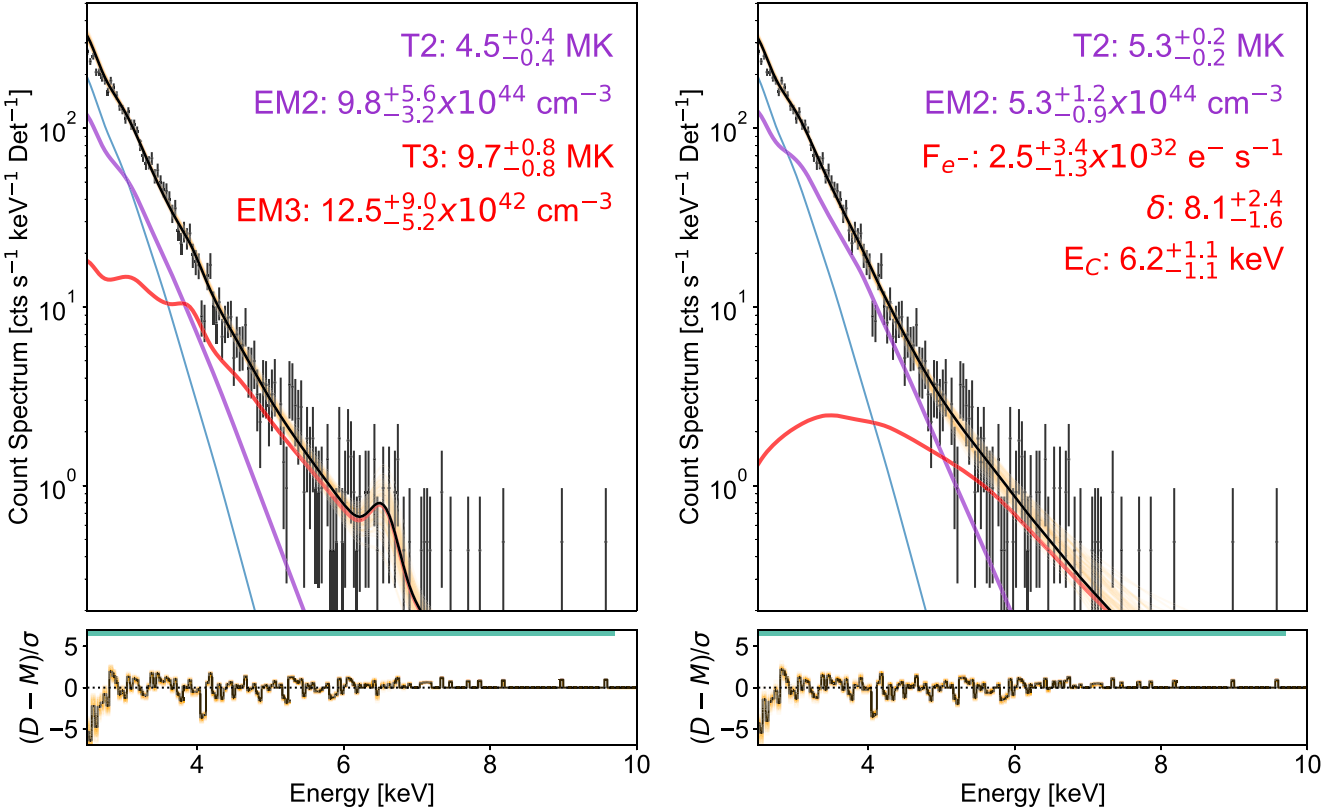
Figure 4. Two different spectral fits of the microflare spectra from the right panel of Fig. 3, one with the count model excess fitted with a thermal model (the *thermal scenario*, left panel) and the other fitted with a non-thermal model (*non-thermal scenario*, right panel) shown in red. The model parameters are shown in their respective plots and colours with their error values showing 1σ equivalent uncertainties resulting from MCMC analysis with 100 random model samples shown in orange.

thermal parameters are well within those previously reported for other microflares (Hannah et al. 2008; Glesener et al. 2020; Cooper et al. 2021). The lower energy thermal models (purple) show a decrease in temperature and increase in emission measure compared to the values obtained in Fig. 3 (right panel).

Using the temperatures and emission measures obtained from all microflare fits, we find the microflare to be an estimated GOES class between A0.02 and A0.03. The GOES class is roughly estimated using the `goes_flux49.pro` IDL routine as the microflare was too weak to provide a direct GOES/XRS measurement.

From the flare volume upper limit determined in Section 2.2 and the microflare thermal parameters, we find thermal energies of $1.6^{+0.1}_{-0.1} \times 10^{26}$, $3.0^{+0.4}_{-0.4} \times 10^{26}$, and $2.1^{+0.1}_{-0.1} \times 10^{26}$ erg for the null, thermal, and non-thermal scenario fits, respectively, using equations (3) and (A4) from Hannah et al. (2008) and Aschwanden et al. (2015). In addition, the non-thermal power associated with the thick-target component in Fig. 4 (right panel) is calculated to be $2.7^{+2.7}_{-1.2} \times 10^{24}$ erg s^{-1} (equation 5 from Wright et al. 2017), which results in $3.7^{+3.8}_{-1.6} \times 10^{26}$ erg over the course of the impulsive phase. All thermal energies are in the expected range for a weak microflare and we note that in the non-thermal case both energies are consistent with each other.

Therefore, both hypotheses shown in Fig. 4 present an interesting and unique conclusion, either showing a microflare of this scale producing such hot temperatures in the impulsive phase or providing positive evidence for the presence of visible non-thermal particle acceleration via a clear count excess in the spectral fits.

2.3.1 Best spectral fit

The fits shown in Figs 3 and 4 are determined using the Poissonian likelihood, $\mathcal{L}(\theta)$, given by

$$\mathcal{L}(\theta) = \prod_{i=0}^{N} \frac{M_i(\theta)^{D_i} \times e^{-M_i(\theta)}}{D_i!}, \tag{1}$$

where $M_i(\theta)$ represents the predicted model counts with parameters θ and D_i is the observed number of counts in a given energy bin, i. The number of data bins included in the optimization is represented by N.

We find the optimum fit to the microflare spectra with the thermal, double-thermal, and thermal and non-thermal models give a Poissonian log-likelihood value, $\ln(\mathcal{L}_{opt})$, of -520, -484, and -480, respectively. This loosely suggests that the non-thermal scenario fit (Fig. 4, right panel) best represents the observed spectra out of the tested models. However, it may be expected that this is the case since the thick-target model has more free parameters; therefore, it is difficult to ascertain whether the non-thermal scenario should be trusted more with such a small difference in log-likelihood values given the different number of parameters available to each model.

It is possible that there is a discrepancy between the thermal excess scenario (Fig. 4, left panel) and the EUV time profiles displayed in Fig. 1. The double-thermal fit suggests the presence of material at ~10 MK during the microflare's impulsive phase. However, the SDO/AIA 131 Å channel has a significant response to material of this temperature (Lemen et al. 2012) but the light curve does not show

a noticeable increase in emission around this time. Folding the 10 MK thermal component through the EUV channel response, we find that the hot component is very faint and may be too weak to produce a significant increase in 131 Å channel meaning that the emission mechanism of the higher X-ray energy excess remains ambiguous.

A more robust approach is needed to assess the thermal and non-thermal representations of the microflare spectra. The nested sampling approach described in Skilling (2004) provides a reproducible, intuitive, and reliable method to ascertain how much more likely one hypothesis is compared to another. Using this approach, we can investigate the most likely representation of the microflare spectra.

3 COMPARING HYPOTHESES

Bayes formula (Jaynes 2003) is used extensively to test how well a certain hypothesis, H, explains given data, D, via the posterior, $P(H|D)$, and is given by

$$P(H|D) = \frac{P(D|H)P(H)}{P(D)}, \tag{2}$$

where $P(D|H)$ is termed the likelihood that describes how likely the data are observed if the hypothesis is true and $P(H)$ represents the prior information on the hypothesis (Sivia & Skilling 2006). The term $P(D)$ is the evidence that describes how likely the data are to occur independent of hypothesis that is difficult to obtain and, in some cases, has very little meaning.

Different hypotheses for the given data can then be compared relative to each other using equation (2) by calculating the posterior odds. To compute this between two different hypotheses – say H_1 and H_2 – for the same data, the ratio is taken between the respective posteriors as

$$\frac{P(H_1|D)}{P(H_2|D)} = \frac{P(D|H_1)}{P(D|H_2)} \frac{P(H_1)}{P(H_2)}, \tag{3}$$

where we note that the evidence term, $P(D)$, from equation (2) cancels (Ashton et al. 2022). Equation (3) therefore provides a relative metric as to which hypothesis between H_1 and H_2 is more likely. It is common that the *prior odds*, ratio of $P(H_1)$ and $P(H_2)$, is evaluated to 1 if there is no prior justifiable preference for one hypothesis over another and so the posterior odds in this case is fully equivalent to the ratio of the likelihoods, termed *Bayes factor* (BF).

Using a nested sampling algorithm (Skilling 2004, 2006), we can estimate the likelihoods by investigating the parameter posterior under the assumption that the hypothesis, now model M, is true. Nested sampling is a Bayesian tool used to compare parametric model fits to data where the models have parameters, θ, and priors, $P(\theta|M)$. This is done by calculating the parameter posterior, $P(\theta|D, M)$, by

$$P(\theta|D, M) = \frac{P(D|\theta, M)P(\theta|M)}{P(D|M)}, \tag{4}$$

where we note that all terms depend on the model, M, and the data, D, being fitted. This explicit form is typically presented in Bayesian analysis, such as in MCMC analysis (Ireland et al. 2013) where the $P(D|M)$ term is often excluded in practice.

However, to avoid the repetition of nomenclature, we adopt another equivalent representation:

$$P(\theta) = \frac{\mathcal{L}(\theta)\pi(\theta)}{\mathcal{Z}}, \tag{5}$$

where the nested sampling process uses the likelihood $\mathcal{L}(\theta)$ (e.g. equation 1) and the prior $\pi(\theta)$ to estimate both the evidence $\mathcal{Z}$ and the posterior $P(\theta)$.

Note that $\mathcal{Z}$ is equivalent to the likelihood term $P(D|H)$, in equations (2) and (3), and is *different* to the previous evidence term $P(D)$ from equation (2).

The evidence, $\mathcal{Z}$, could be calculated as

$$\mathcal{Z} = \int_\theta \mathcal{L}(\theta)\pi(\theta)\mathrm{d}\theta, \tag{6}$$

ensuring that the integral of $P(\theta)$ over the parameter space is equal to 1 (Sivia & Skilling 2006). However, this integral can quickly become complex with an increasing number of parameter dimensions. The nested sampling process takes an approach to simplify this calculation.

Estimating the pieces of evidence, $\mathcal{Z}$ [likelihood $P(D|H)$ in equation 3] allows for a robust comparison between models and how well they explain the data. Criteria from Kass & Raftery (1995), for example, can then be used to assign a relative confidence of one hypothesis compared to another. Importantly, this is also true even if the hypotheses rely on a different number of parameters or parameter priors. The next section details how nested sampling is performed and how $\mathcal{Z}$ is obtained mathematically.

3.1 Calculating the evidence $\mathcal{Z}$ with nested sampling

The evidence $\mathcal{Z}$ is calculated by sampling a predefined and physically motivated constrained prior, $\pi(\theta)$, for the model parameters while calculating the likelihood, $\mathcal{L}(\theta)$, at every corresponding location. From this sampling, a quantity called the *prior mass* can be defined (Skilling 2006). The prior mass, ξ, is the amount of prior enclosed by some likelihood λ, which is given by

$$\xi(\lambda) = \int_{\mathcal{L}(\theta)>\lambda} \pi(\theta)\mathrm{d}\theta, \tag{7}$$

and contains all sampled points with $\mathcal{L}(\theta) > \lambda$. From equation (7), we note that $\xi = 1$ when $\lambda = 0$ (all sampled points are enclosed) and so $\xi = 0$ when $\lambda = \mathcal{L}_{max}$.

The definition of the prior mass can then simplify equations (5) and (6) to

$$P(\xi) = \frac{\mathcal{L}(\xi)}{\mathcal{Z}} \tag{8}$$

and

$$\mathcal{Z} = \int_0^1 \mathcal{L}(\xi)\mathrm{d}\xi, \tag{9}$$

respectively, where the evidence integral is now only performed over one dimension, the prior mass. Equation (8) shows that taking a random sample of ξ_i from 0 to 1, the range defined in equation (9), provides a random sample from the posterior $P(\xi)$. This is equivalent to obtaining a random sample from the posterior $P(\theta)$ with a random parameter sample of θ_i as shown by equation (5). A nested sampling algorithm will then iterate through the random samples spread over the constrained prior to calculate the evidence and posterior distribution.

3.1.1 The information $\mathcal{H}$

A useful value obtained throughout the nested sampling process is a quantity known as the information, $\mathcal{H}$. This quantity describes where the bulk of the posterior lies, with the posterior occupying a fraction of $\mathrm{e}^{-\mathcal{H}}$ within the constrained prior (Skilling 2006) and can also be used to estimate the likely number of iterations to compress the majority of the prior (Ashton et al. 2022). For example, if the

majority of the posterior takes up a small region of the prior then $\mathcal{H}$ is large and if most of the posterior occupies a large area of the prior then $\mathcal{H}$ is small.

3.1.2 Iterating through the samples

The PYTHON nested sampling package, NESTLE,[3] used throughout the analysis presented in Section 4 initially starts by sampling the constrained prior n_{live} times where the corresponding likelihood value is calculated at each sample. These *live-points* are then ordered from smallest likelihood, $\mathcal{L}_*$, to largest. The prior mass is then calculated for $\mathcal{L}_*$ before the corresponding live-point is removed and replaced with another sample with a likelihood value $> \mathcal{L}_*$ to conserve the number of samples across iterations. The prior mass and the likelihood of the removed live-point is recorded and used to build the integral described by equation (9). Once a new live-point is chosen, the process begins again and continues until the stopping criteria have been met (Skilling 2004).

The stopping criterion used in NESTLE is one that estimates the remaining prior mass to be sampled. The iterative process can be stopped when the remaining estimated mass falls below a defined threshold, A_{thresh}, using

$$\ln(\mathcal{L}_{\text{max}, i}\xi_i + \mathcal{Z}_i) - \ln(\mathcal{Z}_i) < A_{\text{thresh}}, \qquad (10)$$

where $\mathcal{L}_{\text{max}, i}$ is the maximum likelihood of the live-points, ξ_i is the remaining prior mass, and $\mathcal{Z}_i$ is evidence value at iteration i.

In addition to simplifying the evidence calculations, nested sampling also obtains many random samples of the posterior through the iterative procedure. Therefore, almost as a by-product, the same information that MCMC analysis provides is also obtained in the process.

A nested sampling algorithm can be applied to many different problems involving model comparison. In Section 4, we apply the nested sampling algorithm used by NESTLE to the AR X-ray microflare presented in Section 2 to determine the emission mechanisms present during its impulsive phase by testing different hypotheses for the observed data: the null, thermal, and non-thermal scenarios.

4 THERMAL OR NON-THERMAL EMISSION

The analysis presented in Section 2 seems to indicate that non-thermal emission is present during this estimated A0.02 microflare's impulsive phase; however, the arguments are qualitative or weakly quantitative and far from decisive. Section 2.1 shows evidence of an earlier peaking and more impulsive time profile in the higher, 6–10 keV X-ray energies observed by *NuSTAR* compared to the lower energy microflare emission (Fig. 1, left panels). This supports the presence of an additional process beyond the one producing the EUV microflare signal, but it is not clear whether this is due to a thermal or non-thermal source.

In Section 2.2, we discussed the observed presence of a clear microflare loop structure in the Fe XVIII emission (Fig. 2), suggesting the presence of material between 4 and 10 MK, which appears to be corroborated by all microflare spectral fits (Fig. 3, right panel, and Fig. 4). We also see loop footpoints across the SDO/AIA EUV channels that are anchored in two opposite magnetically polarized regions as viewed from SDO/HMI. Fig. 2 indicates the presence of bright footpoints, which could potentially be a result from heated

[3]https://github.com/kbarbary/nestle

Table 1. The lower and upper limits to define the uniform constrained prior for each model parameter used for nested sampling. The same notation for each parameter used in the spectral fitting is given in brackets as well as the parameter units. The parameters $T2$ and $EM2$ are present in all microflare fits (Figs 3 and 4), while $T3$ and $EM3$ are only in the thermal scenario fit and F_{e^-}, δ, and E_c are only in the non-thermal scenario fit shown in Fig. 4 to fit the count model excess.

Model parameter	Lower	Upper
Microflare temperature, $T2$ (MK)	2.5	10
Microflare emission measure, $EM2$ (cm^{-3})	1×10^{42}	1×10^{46}
Excess temperature, $T3$ (MK)	8	15
Excess emission measure, $EM3$ (cm^{-3})	1×10^{40}	1×10^{44}
Excess electron flux, F_{e^-} (e$^-$ s^{-1})	1×10^{30}	1×10^{33}
Excess electron spectral index, δ	4	15
Excess low-energy cut-off, E_c (keV)	3	10

material through microflare energized electrons decelerating in the lower solar atmosphere. This interpretation would align with the presence of non-thermal emission.

Performing spectral fitting analysis on the X-ray microflare spectra (Section 2.3) shows more evidence that there is different behaviour below and above 6 keV. We find that an additional model is required to reasonably fit the count excess >6 keV beyond an isothermal fit and to resolve the poor fit <6 keV shown in Fig. 3 (right panel). However, it is still not possible to confidently determine whether an additional thermal or non-thermal model is required (Fig. 4).

In order to obtain a robust measure of confidence over the nature of the high-energy 6–10 keV emission observed from the microflare, we make use of the nested sampling process (see Section 3). The nested sampling result will then provide a vigorous, intuitive, and repeatable level of confidence on the emission mechanism most likely responsible for the higher energy X-ray emission.

4.1 Nested sampling application to SOL2021-11-17T21:14

Since nested sampling relies on a constrained prior (Section 3.1), effort must be taken to define appropriate bounds for each parameter. We make use of a uniform prior for all parameters where the bounds of each are determined by previous microflare studies with consideration to the thermal and non-thermal energetics involved.

Table 1 shows the bounds for the uniform priors used in the nested sampling analysis. Studies of microflares comparable in scale to the one studied here, estimated GOES class ~A0.02, have suggested that those fitted with a single-thermal model tend to have a temperature no higher than ~10 MK with emission measure $\lesssim 1 \times 10^{46}$ cm^{-3} (Duncan et al. 2021; Vadawale et al. 2021). Additionally, some enhanced pre-flare emission may be present while some isothermal *NuSTAR* microflares have emission measures as low as ~1×10^{43} cm^{-3} (Cooper et al. 2021). Therefore, sensible prior bounds for the microflare temperature and emission measure ($T2$ and $EM2$ from Figs 3 and 4) are taken to be 2.5–10 MK and 1×10^{42} to 1×10^{46} cm^{-3}, respectively.

From the physically motivated parameter prior ranges for $T2$ and $EM2$, and using the volume estimate from Section 2.2, the possible thermal energies resulting from models within these bounds fall between approximately 4×10^{24} and 2×10^{27} erg. Considering past *NuSTAR* sub-A class microflares, this is a conservative thermal energy range and helps support the choice of the defined parameter prior bounds.
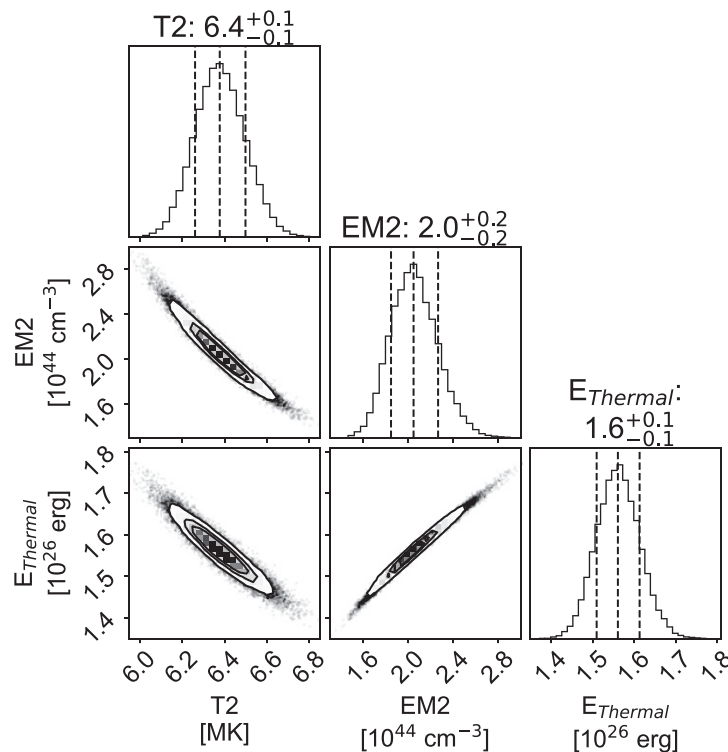
Figure 5. Posterior distributions for the isothermal, *null scenario*, representation of the microflare spectra shown in Fig. 3 (right panel) as determined from nested sampling analysis. The thermal energy distribution is calculated from the temperature ($T2$) and emission measure ($EM2$) samples. The median and 1σ confidence intervals are displayed above the diagonal panels for each marginalized parameter distribution and shown with dashed lines. Contours of 50, 68, and 90 per cent levels are shown in the 2D marginalized posterior distribution.

Next, priors must be defined for the thermal ($T3$ and $EM3$) and non-thermal (F_{e^-}, δ, and E_c) excess model parameters. When other microflares have been found to be best fitted with two isothermal models, we find that the emission measure for the higher temperature component is ~ 2 orders of magnitude less than the lower temperature model (Cooper et al. 2021; Duncan et al. 2021) and so we choose a prior emission measure range of 1×10^{40} to 1×10^{44} cm^{-3}. In order to investigate a potential hotter temperature component producing the count model excess in Fig. 3 (right panel), we set a prior temperature range of 8–15 MK such that it overlaps with, and goes beyond, the range being investigated for $T2$.

The priors for the non-thermal excess models are determined from previous non-thermal X-ray microflare studies. Hannah et al. (2008) find that 90 per cent of the RHESSI microflares studied have a low-energy cut-off between 9 and 16 keV with photon indices between 4 and 10 (corresponding to electron indices between 5 and 11 for thick-target emission; Brown 1971). However, the majority of these non-thermal X-ray microflares are of a greater GOES class and brighter than ones viewed by *NuSTAR* (Glesener et al. 2020; Cooper et al. 2021), which do not have emission extending to energies as high. Therefore, the conservative uniform prior bounds for the low-energy cut-off and the electron index are defined to be 3–10 keV and 4–15, respectively.

The microflare presented here is much weaker in magnitude than the non-thermal microflares investigated previously by *NuSTAR* and so the electron flux is expected to be less than previously obtained. Therefore, the prior range is estimated to be 10^{30} to 10^{33} electrons s^{-1} as this microflare is ~ 3 orders of magnitude weaker in GOES classification than the non-thermal microflare

studied in Glesener et al. (2020), which is found to have an electron flux of $\sim 2 \times 10^{35}$ electrons s^{-1}. The prior constraints for the non-thermal model provide an energy range $\sim 10^{23}$ to 10^{27} erg, which is complementary to the thermal energy range being explored for the microflare providing additional physical justification for the thick-target parameter constrained prior bounds.
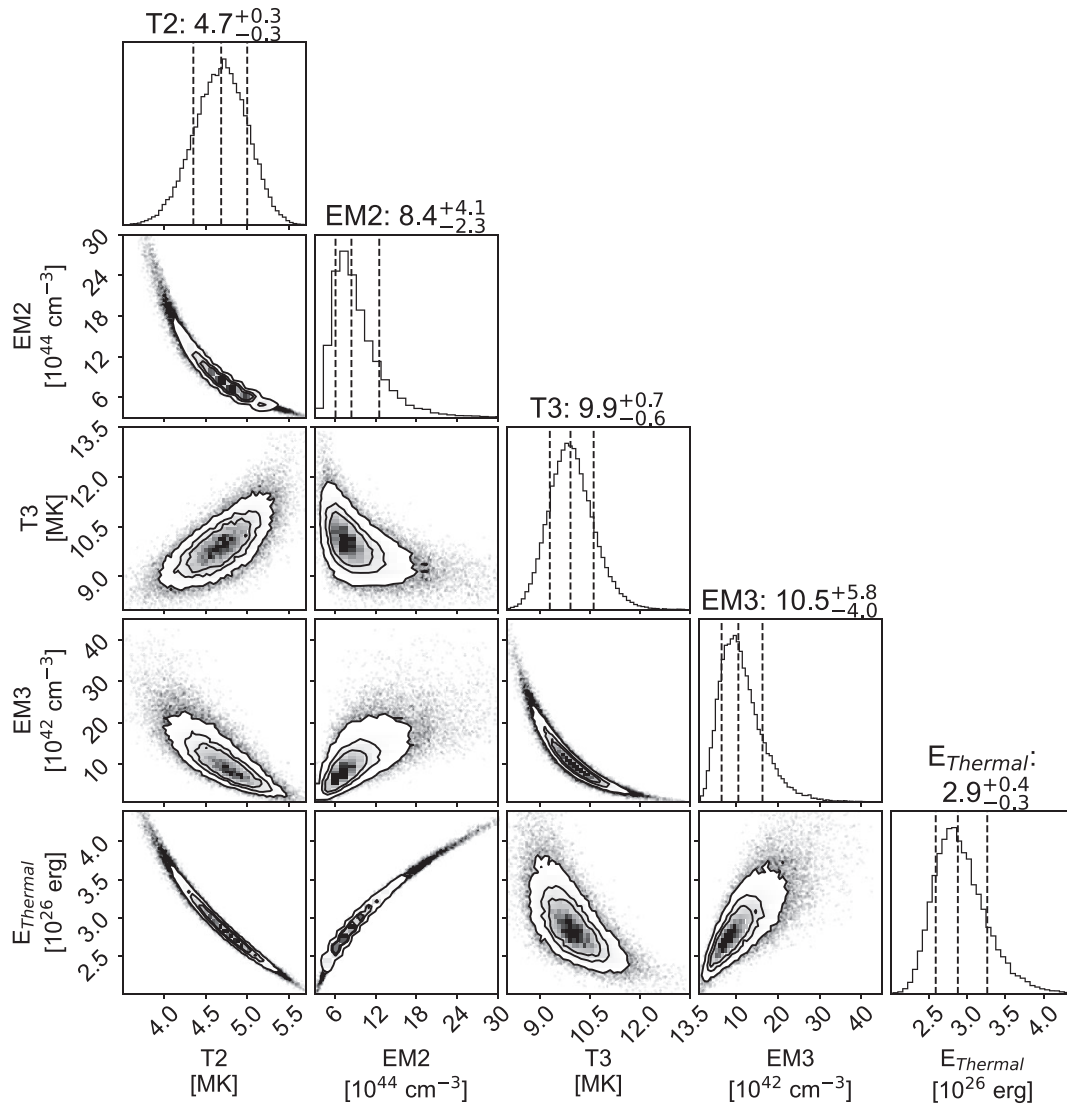
Throughout the nested sampling analysis, all three scenarios are run with 10,000 live-points and an A_{thresh} value of 0.1. These, again, are conservative values to ensure that the majority of the prior is sampled adequately. The nested sampling result does not change significantly with fewer live-points or a larger A_{thresh}, only the uncertainties become larger.

4.2 Nested sampling results

Nested sampling analysis is performed on all three model representations of the microflare spectra where the process took of the order of hours to complete for each model scenario using an Apple M1 Max MacBook Pro. The thermal and non-thermal representations of the excess from Fig. 4 (left and right panels) give $\ln(\mathcal{Z})$ values of -497.31 ± 0.03 and -493.63 ± 0.03, respectively, while the null scenario, where the excess is not represented (Fig. 3, right panel), is estimated to have a value of -530.14 ± 0.03.

The $\mathcal{H}$ values (as described in Section 3.1.1) reported for the analysis in each case are given to be ~ 9 for the null scenario and ~ 11 for both the thermal and non-thermal cases. The similar values indicate that the bulk of the posterior lies within a similar fraction of the constrained prior for each tested hypothesis.

Figure 6. Similar to Fig. 5 but the nested sampling analysis result for the *thermal scenario* shown in the left panel of Fig. 4. The thermal energy distribution is calculated from the temperature (*T*2 and *T*3) and emission measure (EM2 and EM3) samples.

The posterior distribution corner plots that accompany the pieces of evidence are shown in Figs 5–7. We find that the constructed corner plots from the nested sampling analysis agree closely with those produced from the previous MCMC analysis performed in Section 2.3 with all parameter values agreeing well within the 1σ equivalent uncertainties. This provides confidence that the posterior distribution for each model parameter, bound by the defined priors, is sampled well in the MCMC and that there are no other relevant solutions inside the physically motivated prior bounds.

With the confirmation from the corner plots produced from the nested sampling analysis, and corroboration with the MCMC analysis that the pieces of evidence have been calculated robustly, we are able to compare each model representation to the others using BF from equation (3). This is reasonable as we have no evidence or reason to prefer one model representation over the other; therefore, we assume a prior odds ratio of 1 making the BF equal to the posterior odds.

The BF for each comparison is shown in Table 2. We find that the BFs comparing the thermal and non-thermal representations of the

count excess are far more probable than not representing the excess at all in the null scenario. Therefore, nested sampling has given a quantitative measure of how much an additional model is required beyond the pre-flare and isothermal fit with the thermal and non-thermal scenarios being $>10^{14}$ times more probable than the null hypothesis.

Additionally, the nested sampling also shows that the non-thermal representation is $\sim$40 times more likely to be the cause of the count excess compared to the thermal model, corresponding to strong evidence for this conclusion. Therefore, the nested sampling analysis would suggest that this event is the weakest X-ray microflare to show direct evidence for non-thermal emission.

To check the sensitivity of the nested sampling result on the choice of physically motivated prior bounds from Table 1, we adjust the bounds of each parameter prior by up to 20 per cent. This still results in the non-thermal scenario being more preferable to the thermal representation with ln(BF) values between 3.5 and 3.8.
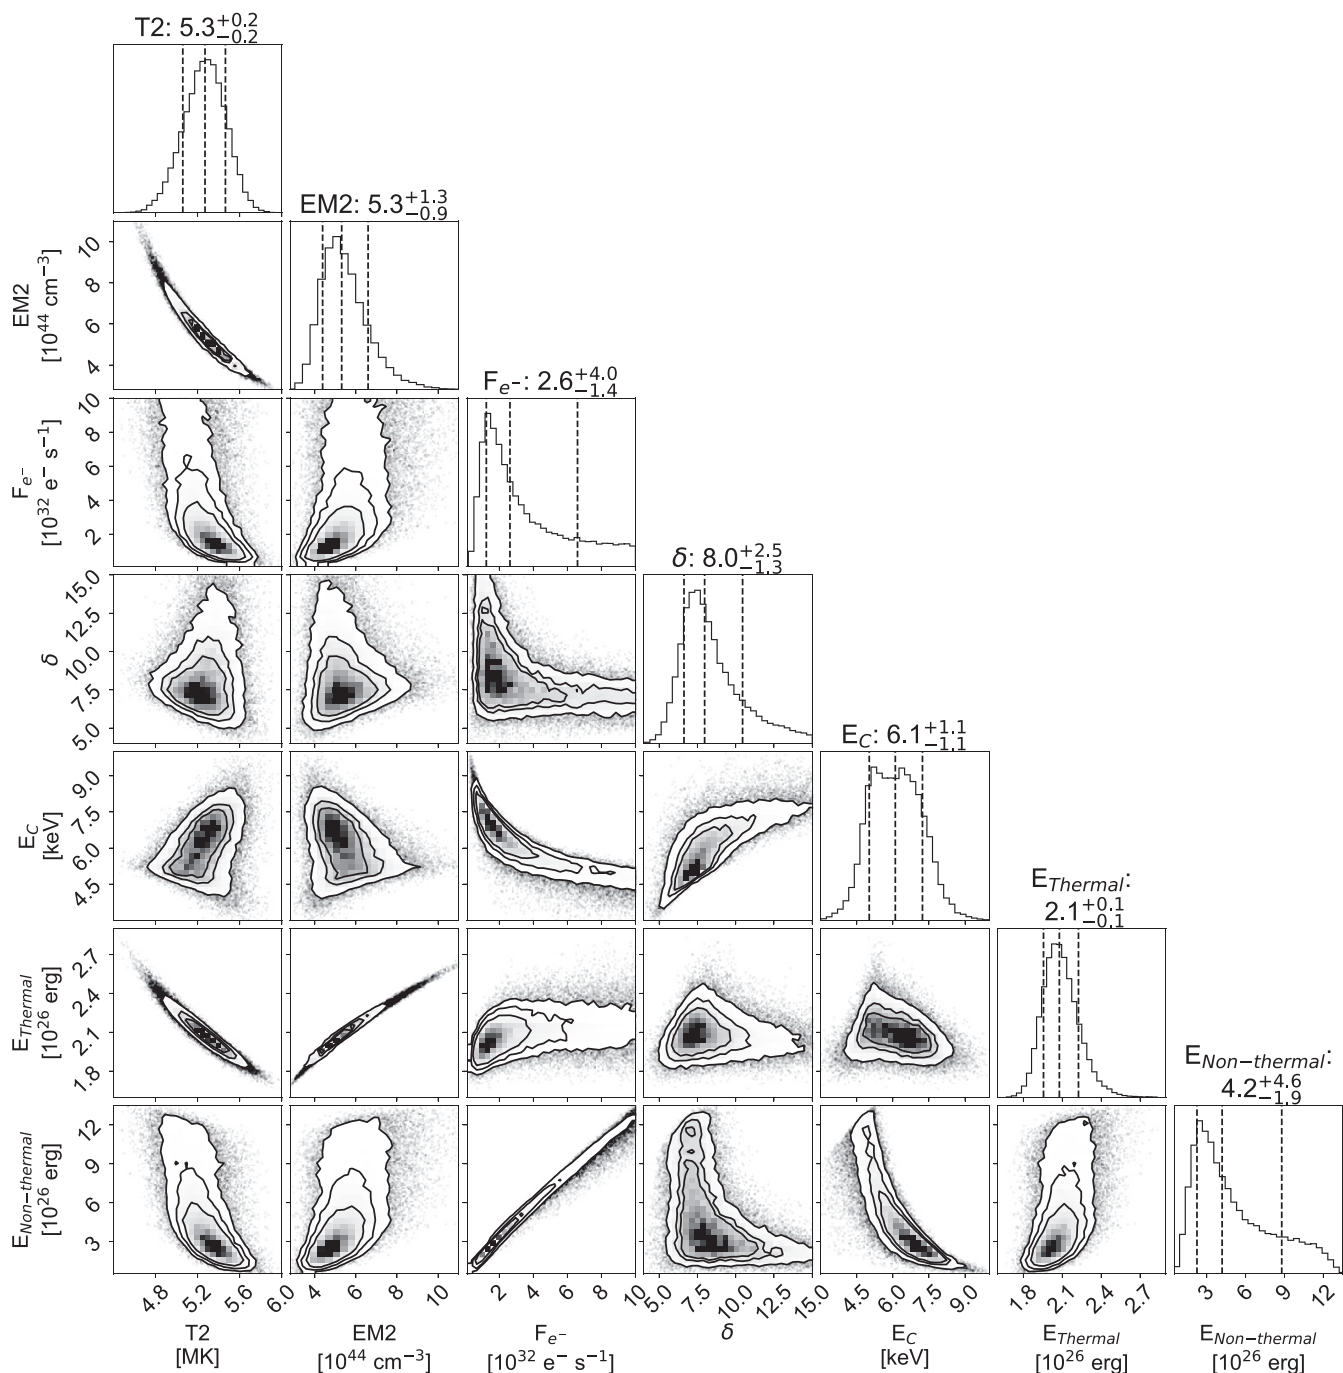
Figure 7. Similar to Fig. 5 but the nested sampling analysis result for the *non-thermal scenario* shown in the right panel of Fig. 4. The thermal energy distribution is calculated from the temperature ($T2$) and emission measure (EM2) samples and the non-thermal energy distribution is calculated from the electron flux (F_{e^-}), electron spectral index (δ), and the low-energy cut-off (E_c) samples.

Table 2. The logarithm of BF, BF, and the corresponding level of confidence from Kass & Raftery (1995) are shown for each model scenario comparison for the microflare spectra. The error on ln(BF) is ± 0.04 for all values.

Comparison	ln(BF)	BF	Confidence
$\mathcal{Z}_{\text{thermal}}$ to $\mathcal{Z}_{\text{null}}$	32.8	2×10^{14}	Very strong
$\mathcal{Z}_{\text{non-thermal}}$ to $\mathcal{Z}_{\text{null}}$	36.5	7×10^{15}	Very strong
$\mathcal{Z}_{\text{non-thermal}}$ to $\mathcal{Z}_{\text{thermal}}$	3.7	40	Strong

Furthermore, we narrow the priors around the bulk of the $T3$ and EM3 distributions shown in Fig. 6 while leaving the non-thermal parameter priors unchanged in order to be more favourable to the thermal scenario. A ln(BF) value of ~ 2 is obtained from this test, which is still classified as positive evidence for the non-thermal representation of the microflare (Kass & Raftery 1995).

We note that the best-fitting log-likelihood values, $\ln(\mathcal{L}_{\text{opt}})$, stated in Section 2.3.1 show similar differences to the log-evidence values, $\ln(\mathcal{Z})$, for the different scenarios; however, the nested sampling result

is more informative and robust for model comparison. The pieces of evidence can sufficiently compare models with different parameters and priors, whereas a clear conclusion cannot be drawn from the best-fitting log-likelihood values as discussed in Sections 2.3.1 and 3.

5 SUMMARY AND CONCLUSIONS

In this paper, we present the first analysis of X-ray flare spectra using a nested sampling algorithm. This microflare is estimated to be approximately an A0.02 GOES equivalent class flare that shows signs of high temperatures and/or non-thermal particle acceleration; however, determining the extent of this emission is difficult using arguments and methods performed in the past.

From the temporal, spatial, and spectral information discussed in Section 2, we find evidence that the X-ray microflare emission is not well represented with an isothermal model as a count excess is present at higher energies in the *NuSTAR* spectra (Fig. 3). Upon investigating two emission mechanisms potentially responsible for this, we find that both a thermal and a non-thermal explanation for the excess appear to be equally valid with no decisive reason to pick one over the other.

Qualitative evidence for a non-thermal representation is found in the form of the lack of higher temperature (∼10 MK) emission response in the 131 Å SDO/AIA channel (Section 2.1) and potentially in the presence of bright EUV loop footpoints visible throughout the solar atmosphere (Section 2.2). When performing spectral fitting, the model with the non-thermal component produces a slightly more preferable log-likelihood value (Fig. 4); however, this metric may be expected to be better due to the non-thermal component having an extra free parameter to use during the fitting process compared to the thermal model. Therefore, a conclusion cannot be made based on the best-fitting log-likelihood values.

We use a more robust model comparison method called nested sampling (Section 3), which is used to determine an estimate for the evidence attributed to each model representation of the data, therefore allowing BF to be computed for model comparison. In the process, we also obtain the posterior distributions for each model as shown in Figs 5–7. We find that the model containing the non-thermal component is most likely to represent the observed microflare spectra with BFs of 7×10^{15} and 40 compared to the null and thermal representations, respectively.

Therefore, we report the weakest X-ray microflare to have direct observation of non-thermal particle emission suggesting that X-ray flares of the estimated scale A0.02 do indeed show similar traits to their larger counterparts. From the nested sampling results, this microflare produces a thermal energy of $2.1^{+0.1}_{-0.1} \times 10^{26}$ erg, which is rational compared to the non-thermal energy of $4.2^{+4.6}_{-1.9} \times 10^{26}$ erg released over the course of its impulsive phase at $3.0^{+3.3}_{-1.4} \times 10^{24}$ erg s^{-1}. The corresponding energy posterior distributions of each model scenario are shown in Figs 5–7. This shows that microflares this weak, and likely weaker, continue to undergo the same physical processes with the same/similar emission mechanisms as brighter flares.

The nested sampling algorithm is in no way specific to X-ray microflare analysis. It can be used to gain insight into general model comparison problems and is able to produce the same data products as MCMC analysis. However, effort will be made to speed up the process specific to this microflare analysis. This could be achieved by speeding up the code used for the individual component models and testing a range of nested sampling algorithm implementations such as diffusive nested sampling (Brewer, Pártay & Csányi 2011) and

those discussed in Speagle (2020) and Williams, Veitch & Messenger (2023).

Using nested sampling has enabled the robust analysis of the high-energy count excess. Methods to determine the emission mechanism for such features in the past either are primarily subjective in nature and/or fail to give a reliable conclusion. The nested sampling algorithm is able to allow the repeatable and intuitive study of data with low signal and is crucial when investigating microflares at such weak scales. Nested sampling will be used in future *NuSTAR* microflare studies, and other flare studies with ambiguous spectra, to determine the confidence on the emission mechanisms present throughout the event.

DATA AVAILABILITY

All data used are publicly available. SDO data can be obtained from the Joint Science Operations Center (JSOC)[4] using SUNPY's `Fido`[5] object, while the *NuSTAR* data are available from the *NuSTAR* Master Catalog[6] with the OBSID 20618003001. An overview of *NuSTAR* solar observation campaigns is also available on GitHub.[7]

REFERENCES

Aschwanden M. J., Boerner P., Ryan D., Caspi A., McTiernan J. M., Warren H. P., 2015, ApJ, 802, 53
Ashton G. et al., 2022, Nat. Rev. Methods Primers, 2, 39
Astropy Collaboration, 2018, AJ, 156, 123

[8] http://www.astropy.org
[4] http://jsoc.stanford.edu/
[5] https://docs.sunpy.org/en/stable/guide/acquiring_data/fido.html#fido-guide
[6] https://heasarc.gsfc.nasa.gov/db-perl/W3Browse/w3table.pl?tablehead=name=numaster&Action = More + Options
[7] https://ianan.github.io/nsigh_all/

Barnes W. et al., 2020a, aiapy, Zenodo, https://doi.org/10.5281/zenodo.431 5741

Barnes W. et al., 2020b, J. Open Source Softw., 5, 2801

Benz A. O., 2017, Living Rev. Sol. Phys., 14, 2

Brewer B. J., Pártay L. B., Csányi G., 2011, Stat. Comput., 21, 649

Brown J. C., 1971, Sol. Phys., 18, 489

Christe S., Hannah I. G., Krucker S., McTiernan J., Lin R. P., 2008, ApJ, 677, 1385

Cooper K., Hannah I. G., Grefenstette B. W., Glesener L., Krucker S., Hudson H. S., White S. M., Smith D. M., 2020, ApJ, 893, L40

Cooper K. et al., 2021, MNRAS, 507, 3936

Del Zanna G., 2013, A&A, 558, A73

Dere K. P., Landi E., Mason H. E., Monsignori Fossi B. C., Young P. R., 1997, A&AS, 125, 149

Duncan J. et al., 2021, ApJ, 908, 29

Feldman U., Mandelbaum P., Seely J. F., Doschek G. A., Gursky H., 1992, ApJS, 81, 387

Fletcher L. et al., 2011, Space Sci. Rev., 159, 19

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306

Glesener L., Krucker S., Hannah I. G., Hudson H., Grefenstette B. W., White S. M., Smith D. M., Marsh A. J., 2017, ApJ, 845, 122

Glesener L. et al., 2020, ApJ, 891, L34

Grefenstette B. W. et al., 2016, ApJ, 826, 20

Hannah I. G., Christe S., Krucker S., Hurford G. J., Hudson H. S., Lin R. P., 2008, ApJ, 677, 704

Hannah I. G., Hudson H. S., Battaglia M., Christe S., Kašparová J., Krucker S., Kundu M. R., Veronig A., 2011, Space Sci. Rev., 159, 263

Hannah I. G. et al., 2016, ApJ, 820, L14

Hannah I. G., Kleint L., Krucker S., Grefenstette B. W., Glesener L., Hudson H. S., White S. M., Smith D. M., 2019, ApJ, 881, 109

Harris C. R. et al., 2020, Nature, 585, 357

Harrison F. A. et al., 2013, ApJ, 770, 103

Holman G. D. et al., 2011, Space Sci. Rev., 159, 107

Hudson H., 1991, Sol. Phys., 133, 357

Hunter J. D., 2007, Comput. Sci. Eng., 9, 90

Ireland J., Tolbert A. K., Schwartz R. A., Holman G. D., Dennis B. R., 2013, ApJ, 769, 89

Jaynes E. T., 2003, Probability Theory: The Logic of Science. Cambridge Univ. Press, Cambridge

Kass R. E., Raftery A. E., 1995, J. Am. Stat. Assoc., 90, 773

Kester D., Mueller M., 2021, Astron. Comput., 37, 100503

Knuth K. H., Habeck M., Malakar N. K., Mubeen A. M., Placek B., 2015, Digit. Signal Process., 47, 50

Kuhar M., Krucker S., Glesener L., Hannah I. G., Grefenstette B. W., Smith D. M., Hudson H. S., White S. M., 2018, ApJ, 856, L32

Landi E., Feldman U., Dere K. P., 2002, ApJ, 574, 495

Landi E., Young P. R., Dere K. P., Del Zanna G., Mason H. E., 2013, ApJ, 763, 86

Lemen J. R. et al., 2012, Sol. Phys., 275, 17

Lin R. P., Schwartz R. A., Kane S. R., Pelling R. M., Hurley K. C., 1984, ApJ, 283, 421

Lin R. P. et al., 2002, Sol. Phys., 210, 30

Mumford S. J. et al., 2021, SunPy, Zenodo, https://doi.org/10.5281/zenodo.4 421322

Parker E. N., 1988, ApJ, 330, 474

Paterson S., Hannah I. G., Grefenstette B. W., Hudson H. S., Krucker S., Glesener L., White S. M., Smith D. M., 2023, Sol. Phys., 298, 47

Priest E., Forbes T., 2002, A&AR, 10, 313

Schou J. et al., 2012, Sol. Phys., 275, 229

Schwartz R. A., Csillaghy A., Tolbert A. K., Hurford G. J., McTiernan J., Zarro D., 2002, Sol. Phys., 210, 165

Shanmugam M. et al., 2020, Curr. Sci., 118, 45

Sivia D. S., Skilling J., 2006, Data Analysis: A Bayesian Tutorial, 2nd edn. Oxford Science Publications. Oxford Univ. Press, Oxford

Skilling J., 2004, Am. Inst. Phys. Conf. Proc., 735, 395

Skilling J., 2006, Bayesian Anal., 1, 833

Speagle J. S., 2020, MNRAS, 493, 3132

SunPy Community, 2020, ApJ, 890, 68

Vadawale S. et al., 2014, Adv. Space Res., 54, 2021

Vadawale S. V. et al., 2021, ApJ, 912, L13

Verbeeck C., Delouille V., Mampaey B., De Visscher R., 2014, A&A, 561, A29

Virtanen P. et al., 2020, Nat. Methods, 17, 261

Williams M. J., Veitch J., Messenger C., 2023, Mach. Learn.: Sci. Technol., 4, 035011

Wright P. J. et al., 2017, ApJ, 844, 132

APPENDIX A: X-RAY THERMAL AND NON-THERMAL PHOTON MODELS

Fig. A1 displays the photon models used to investigate the null (left panel), thermal (middle panel), and non-thermal (right panel) scenarios for the microflare as initially described in Section 2.3. Two photon models are utilized to represent each scenario, an isothermal model and a thick-target non-thermal model. Both models represent a photon spectrum produced from different electron distributions.

The isothermal model (purple in all panels and red in the middle panel of Fig. A1) represents a photon spectrum produced from a Maxwell–Boltzmann distribution of electrons. The isothermal model is composed of continuum and line emission and is a function of the plasma temperature (T) and emission measure (EM).

The non-thermal photon spectrum (red in the right panel of Fig. A1) results from a power-law electron distribution losing all their energy through Coulomb collisions (Brown 1971; Holman et al. 2011). The thick-target model is a function of electron flux (F_{e^-}), electron spectral index (δ), and low-energy cut-off (E_c).
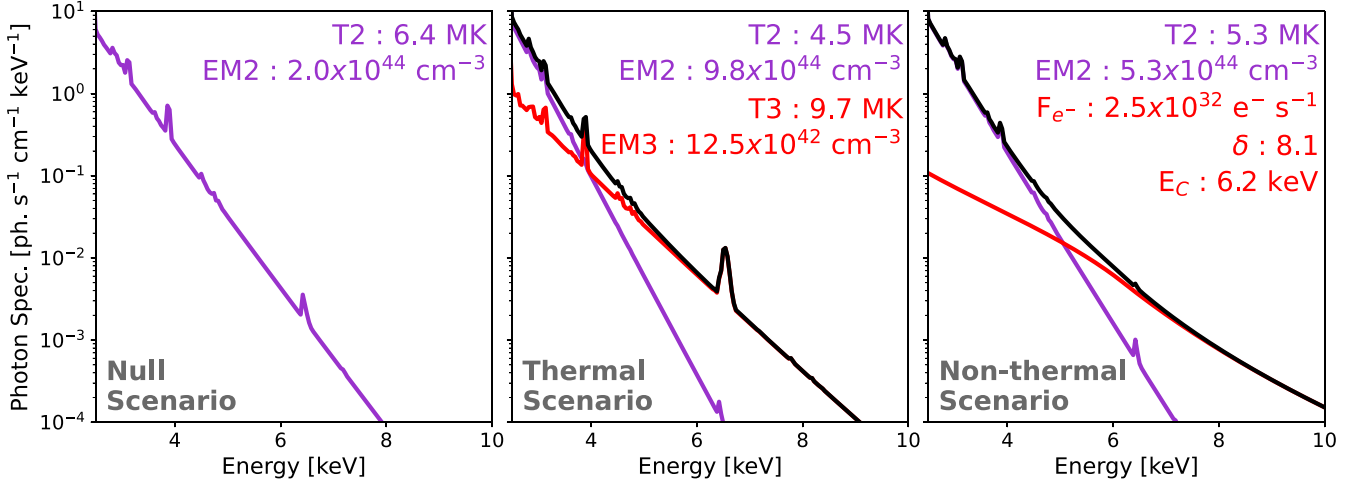
Figure A1. Photon models, obtained from SUNKIT-SPEX, used to fit the *NuSTAR* microflare X-ray count spectra in the *null* (left panel), *thermal* (middle panel), and *non-thermal* (right panel) *scenarios* shown in Fig. 3 (right panel) and Fig. 4. The models shown are calculated using *NuSTAR*'s native energy resolution and the parameter values obtained from MCMC analysis (Section 2.3). Each model component, along with corresponding parameter set, is shown in purple or red with the total model in the thermal and non-thermal cases being shown in black.

This paper has been typeset from a TEX/LATEX file prepared by the author.