

ANALYSIS AND IMPLICATIONS OF THE SPATIO-SPECTRAL MORPHOLOGY OF THE FERMI BUBBLES

AMI TANK 

Department of Astronomy, Astrophysics and Space Engineering, Indian Institute of Technology Indore, Indore 453552, India

ROLAND CROCKER 

Research School of Astronomy & Astrophysics, Australian National University, 233 Mt Stromlo Rd, Stromlo, ACT 2611, Australia

MARK R. KRUMHOLZ 

Research School of Astronomy & Astrophysics, Australian National University, 233 Mt Stromlo Rd, Stromlo, ACT 2611, Australia

Version November 11, 2025

Abstract

The Fermi Bubbles are γ -ray structures extending from the center of the Milky Way to $\pm 50^\circ$ Galactic latitude that were discovered in data obtained by the *Fermi*/LAT instrument. Their origin and power source remain uncertain. To help address this uncertainty, here we use a template-free reconstruction of ten years of all-sky *Fermi*/LAT data provided by Platz et al. (2023) to carry out a pixel-by-pixel spectral analysis of the Bubbles. We recover the position-dependent spectral shape and normalization that would be required for parent proton or electron cosmic ray populations to produce the Bubbles' observed γ -ray spectra. We find that models in which the γ -ray emission is driven by either hadronic or leptonic processes can explain the data equally well. The cosmic ray population driving the emission must have either broken power-law or exponentially cut-off spectra, with break or cutoff energies that are almost constant with latitude but spectral indices below the break that harden towards the Bubbles' southern tip. For the leptonic channel, reproducing the observed position-dependent γ -ray spectrum also requires a cosmic ray electron energy density that grows with distance from the Galactic plane and increases towards the edges of the Bubbles. This finding disfavors scenarios for the origin of the Bubbles where a population of cosmic ray electrons is accelerated near the Milky Way center and subsequently advected out to the extremities of the Bubbles.

Subject headings: Galactic center (565), Galactic cosmic rays (567), Gamma-ray sources (633), High-energy cosmic radiation (731)

1. INTRODUCTION

Observations of the Milky Way halo by the Large Area Telescope (LAT) onboard the *Fermi* Gamma-ray Space telescope revealed two giant bubble-shaped structures extending $\sim 50^\circ$ above and below the Galactic plane. These structures were called the Fermi Bubbles (FBs) by their discoverers (Su et al. 2010). The origin of the FBs is still an open question. Their morphology suggests that they originate from the Galactic Center (GC), and their sky location and morphology suggest they are plausibly the gamma-ray counterparts to the “microwave haze” first discovered (Finkbeiner 2004) in data from the Wilkinson Microwave Anisotropy Probe (WMAP; also see Dobler 2012; Planck Collaboration 2013).

Since the discovery of the FBs, structures coincident with or clearly morphologically related to them have also been detected at other wavelengths. The eROSITA telescope revealed two X-ray structures, now known as the eROSITA bubbles, extending up to 14 kpc above and below the Galactic plane (Predehl et al. 2020); these

encircle the FBs in projection and subsume previously-detected X-ray structures including the North Polar Spur (NPS) discovered by ROSAT (Snowden et al. 1997) and the X- or hourglass-shaped structure detected near the GC (Bland-Hawthorn & Cohen 2003) that is coincident with the edge of the FBs. In 2.3 GHz radio emission, Carretti et al. (2013) detected two giant, linearly polarized lobes coincident with the FBs at low latitudes, but extending to greater heights ($\sim 60^\circ$). There are also observations of both warm ionized (e.g., Bordoloi et al. 2017), cool atomic (e.g., Noon et al. 2023), and cold molecular (e.g., Di Teodoro et al. 2020; Cashman et al. 2021) material in or around the edges of the FBs, which may represent cool interstellar medium (ISM) that has been swept up by the same mechanisms responsible for producing the FBs. Sarkar (2024) provides a detailed review of the multi-wavelength observations near the FBs.

Despite the large volume of observations available, numerous basic questions about the FBs remain open. What microphysical mechanism generates the Bubbles' γ -ray emission? What activity provides the energy to inflate them in the first place? Relatedly, how old are the Bubbles? Where and how are the cosmic ray (CR) par-

ents of the γ -rays accelerated? A successful model, along with answering these questions, should also be consistent with the full set of multi-wavelength observations.

The question of mechanism arises because astrophysical production of γ -rays occurs via two primary channels: (a) the hadronic (p-p or proton) channel¹ and (b) the leptonic (electron) channel. In the former, cosmic-ray protons (CRp) collide inelastically with ambient, thermal protons and produce neutral pions, which each further decay into two γ -ray photons. In the latter, γ -ray photons are produced via inverse Compton (IC) up-scattering² of the interstellar radiation field (ISRF) – including the Cosmic Microwave Background (CMB) – by cosmic-ray electrons (CRe). An appealing aspect of the leptonic channel is that synchrotron radiation by the same population of electrons could also generate the microwave haze emission (Su et al. 2010). However, this is only suggestive, and we cannot at this point rule out the possibility that the γ -ray and microwave emission trace different underlying particle populations (e.g., Crocker et al. 2015).

The question of mechanism is tightly tied to the question of the FBs’ ultimate energy source. Models for this fall into two broad categories: those invoking AGN-like activity of the Milky Way’s central super-massive black hole, and those that posit that the FBs are powered by supernovae produced by bursts of star formation near the GC. The link between these scenarios and the emission mechanism is via cooling times. CRe would cool in at most a few Myr due to IC and synchrotron emission (which are comparably important given the $\sim 10 \mu\text{G}$ magnetic field found even at relatively high latitudes within the FBs – Carretti et al. 2013), so if FB emission is driven by CRe accelerated near the GC, those CRe must be transported outward at $\gg 1000 \text{ km s}^{-1}$ to reach the FB edge before cooling. This is well above the typical speeds of supernova-driven winds (e.g., Vijayan et al. 2024), so in this scenario the CRe must either originate from a relativistic jet from the central black hole that was active in the last few Myr (e.g., Guo & Mathews 2012), or there must be some mechanism within the bubbles capable of re-accelerating a cooled electron population (e.g., Mertsch & Sarkar 2011; Mertsch & Petrosian 2019). By contrast, the cooling times for protons are much longer ($\sim \text{Gyr}$ for ambient gas densities $\sim 10^{-2} \text{ cm}^{-3}$), easing the timescale constraint but conversely requiring a much larger CR energy density within the bubble to power the observed emission. This favors a scenario where the CRp population within the FBs is gradually built-up by the cumulative effects of supernova injection over $\gg \text{Myr}$ timescales (e.g., Crocker & Aharonian 2011; Crocker et al. 2014, 2015), though in principle CRp could also be supplied (and much more quickly) by

¹ In principle, this channel also encompasses heavier nuclei in either or both of the target and beam particles; in the absence of strong constraint on the abundance distribution for either of these, we ignore this complication. For Solar metallicity, the presence of heavy nuclei implies a $\mathcal{O}(50\%)$ upwards correction (Mori 1997) to the efficiency of γ -ray production compared to a gas of pure hydrogen, which is small in comparison to other uncertainties, in particular, the target gas number density.

² Leptonic γ -ray emission due to bremsstrahlung is also important in some astrophysical environments, but is not competitive with IC for the ISRF and gas number conditions pertaining close to the FBs.

AGN activity (e.g., Mou et al. 2015).

These cooling timescale arguments can be sharpened by considering variation in the emission across the FBs. The FBs have approximately uniform surface brightness and spectral shape (Su et al. 2010), which at first might seem difficult to reconcile with a leptonic model, since both the energy density and total intensity of the ISRF are expected to be non-uniform over such a large extent. Despite this expectation, when Narayanan & Slatyer (2017) divided the FBs into five latitude bands and fit the spectra in each of them with a leptonic model, they found a surprisingly good fit even with a constant CRe spectrum. However, they also found that such a model does not naturally reproduce the microwave haze, and that while it is possible to produce both the γ -ray and microwave emission with a single electron population, doing so requires significant fine-tuning of the position-dependent spectral shape. Similarly, Yang & Ruszkowski (2017) investigate a leptonic jet scenario by tracking the evolution of CRe using a 3D hydrodynamical simulation. They find that CRe in such a scenario develop a high energy cutoff due to fast synchrotron and IC cooling of the CRe when the jets were launched, but that the CRe population inside the FBs is then homogenized by the subsequent fast transport of CRe within the bubbles. They also find that the CRe have a slight gradient towards higher energies at higher latitudes, which is essential to compensate for the decrease in the ISRF strength.

However, all of the analyses to date have made use of relatively coarse representations of the variation of γ -ray emission across the FBs, for example simply taking means in latitude slices. As we discuss below, the combination of several additional years of *Fermi*/LAT data with improved analysis techniques have yielded data sets with considerably better spatial and spectral resolution. This situation prompts us to revisit the question of what we can learn about the underlying particle population from a spatially-resolved analysis of the FBs’ γ -ray spectra. In this paper we perform a pixel-by-pixel analysis over $\approx 37,000$ pixels covering the FBs, obtaining position-dependent constraints on the properties of the particle populations that must be present to explain the observed γ -ray spectra under either a leptonic or hadronic scenario³.

The remainder of this paper is organized as follows. In Section 2 we discuss the γ -ray data reconstruction and the ISRF model data that we use. In Section 3 we outline our methodology and discuss the assumptions we make about the structure of the FBs for this work. In Section 4 and Section 5 we set out our results and discussion, respectively.

2. DATA

Our analysis relies on measurements and models for the all-sky γ -ray spectrum and the interstellar radiation field; we discuss these in Section 2.1 and Section 2.2, respectively.

2.1. γ -ray data

We based our analysis on the all-sky γ -ray maps derived by Platz et al. (2023). These authors apply a

³ The dataset is released at <https://doi.org/10.5281/zenodo.17568107>




FIG. 1.— Template-free reconstruction of the non-dust associated all-sky diffuse γ -ray photon flux, I^{nd} , taken from the M2 model of Platz et al. (2023); the data shown are for the energy bin centered at 133 GeV. The data are shown in Mollweide projection using Galactic coordinates. The colorbar is plotted on a log scale. The black line is a by-eye tracing of the FBs’ apparent edges, and the white dashed line shows the region $-40^\circ < \ell < 40^\circ$, $-60^\circ < b < 60^\circ$ on which we focus in this work. The four white points labeled a, b, c, and d are chosen to show example results at different points in the sky.

Bayesian variational inference framework to 10 years of *Fermi*/LAT data, which they use to obtain two different decompositions of the all-sky diffuse photon flux density as a function of direction I : a template-free model M1, and a template-informed model M2. The latter divides diffuse emission into two components: I^{dust} , which uses the *Planck* 545 GHz thermal dust emission map as a template, and I^{nd} , which is template-free and therefore accounts for the processes which are uncorrelated with the dust map. We expect the dust-associated component to trace processes such as pp emission from the Galactic disk, which strongly correlates with the location of dense interstellar gas and thus dust, while the nd (no dust) component should capture processes such as IC scattering from CRE or hadronic emission from regions not traced by the dust map.

Most importantly for our purposes, while the FBs are visible in both M1 and M2 reconstructions, in M1 they are somewhat obscured by the strong foreground (and background) emission from the Galactic disk and halo. In contrast, the I^{dust} component in M2 effectively separates out the foreground largely attributable to the emission from the quasi-uniform ‘sea’ of Galactic plane CR hadrons colliding with denser gas, revealing the FBs clearly in the I^{nd} component. For our analysis we therefore employ the template-free component, I^{nd} , of the M2 model⁴. The map we use pixelizes the

sky using the HEALPix RING pixelization scheme (Górski et al. 2005) with an `nside` parameter of 128; it therefore decomposes the sky into 196608 pixels of 6.5×10^{-5} sr each, and within each pixel the map provides the photon flux in 11 logarithmically-spaced energy bins from 0.56–316 GeV; it also provides corresponding uncertainties on these fluxes. Following the advice of Platz et al. (2023), we discard the lowest energy bin from our analysis due to large systematic uncertainties arising from poor knowledge of the low-energy instrumental response function.

Figure 1 shows the γ -ray sky traced by the I^{nd} component of M2 for an energy bin centered (geometrically) at 133 GeV. The black contour in this figure shows our by-eye tracing of the apparent boundary of the FBs; we use this edge in some of our analyses and when discussing some of our results.

2.2. The interstellar radiation field

In leptonic models for the FBs, emission depends on the ISRF, and we therefore also require a model for the ISRF as a function of position. For this purpose we use the model produced by Popescu et al. (2017), which is an axisymmetric radiative transfer (RT) model derived using the *COBE*, *IRAS* and *Planck* maps of the all-sky emission.⁵ The model is spatially discretized onto a cylindrical grid of 23×22 points covering radii $r \in (0, 24)$

⁴ The model is available at <https://doi.org/10.5281/zenodo.7970865>

⁵ The model data are available at <https://doi.org/10.26093/cds/vizier.74702539>.




FIG. 2.— Our model of the specific energy density u_λ of the ISRF (including stellar, dust, and CMB components) as a function of wavelength λ , drawn primarily from Popescu et al. (2017), at two sample positions: the Galactic center ($r = 0$, $z = 0$; blue) and at a point 8 kpc above the Galactic center ($r = 0$, $z = 8$ kpc; orange).

kpc and heights $z \in (0, 10)$ kpc, with symmetry assumed about the $z = 0$ plane. At each point it describes the specific energy density of the radiation field u_λ as a function of wavelength λ as the sum of a direct starlight component covering wavelengths from 91 to 5000 nm (discretized into 15 bins), and a dust-reprocessed component from 3 μm to 1 mm (discretized into 120 bins); to avoid an artificial gap between the components, we extrapolate the stellar component to longer wavelengths by fitting a powerlaw to the three longest wavelength points, following Niederwanger et al. (2019). The total ISRF also contains a contribution from the cosmic microwave background (CMB), a spatially uniform blackbody field with temperature $T = 2.73$ K. We show the total ISRF at two sample points in Figure 2.

As we describe below, our analysis requires that we be able to compute the ISRF at arbitrary spatial positions. For positions within the range of (r, z) covered by the Popescu et al. grid, we use bilinear interpolation on the grid they provide. We also require the ISRF at $z > 10$ kpc, beyond the provided grid, for which purpose we assume that the starlight and dust components of the radiation energy density follow a powerlaw functional form in z at each radius r . For $r < 4$ kpc, where the Popescu et al. model grid behaves smoothly at large z , we perform a powerlaw fit to the three data points with the largest z to generate $u_\lambda(r, z)$ for $z > 10$ kpc. For $r > 4$ kpc the behavior of the radiation field at large z is more complex and a powerlaw fit is poorly-behaved, so we instead simply assume a z^{-1} scaling, typical of the best-fit values at smaller r , and extrapolate from the largest z value at each r under this assumption. Note that, while formally necessary for us to cover our full target region up to Galactic latitudes $b = \pm 60^\circ$, these extrapolations have limited impact because almost all of the volume of the FBs lies below $z = 10$ kpc.

Finally, it is worth noting that many previous studies of the FBs (e.g., Yang et al. 2013, 2022; Narayanan & Slatyer 2017; Yang & Ruszkowski 2017; Herold & Malyshchev 2019), relied on estimates of the spatial variation of the ISRF from GALPROP⁶. Popescu et al. (2017) point out that their ISRF model differs systematically from the GALPROP model v.54 (the pre-2017 model), particu-

larly in the UV and MIR range. The GALPROP model of this generation obtains the stellar and dust ISRF components from star counts and gas measurements but Popescu et al. (2017) point out that this leads to an inconsistency with direct, inner Galaxy measurements by COBE/IRAS/Planck; Popescu et al. (2017) optimize their stellar and dust ISRF component models such that they reproduce the COBE, IRAS and Planck maps. Consequently, their estimate of the radiation field should be more accurate near the Galactic center, an important region for our application.

3. METHODOLOGY

Our goal in this section is to fit every pixel in the FBs with a range of potential models for the high-energy particle population driving the γ -ray emission. Given our aim, we restrict our attention to the region between $-60^\circ < b < 60^\circ$ and $-40^\circ < \ell < 40^\circ$, where (ℓ, b) are the usual Galactic longitude and latitude coordinates, indicated by dashed white line in Figure 1. This region contains 37,714 pixels. We first describe how we compute the auxiliary quantities necessary for fits in Section 3.1 and then detail the fitting method itself in Section 3.2.

3.1. ISRF and gas density estimation

Leptonic models for the FBs require knowledge of the ISRF. In order to estimate this, we make the simplifying geometric assumption that the emitting volumes that contribute to all pixels of interest are prisms whose centers lie on a plane that passes through the GC and is oriented normal to the line from the Sun to the GC, in which case we can map from Galactic coordinates (ℓ, b) to a cylindrical coordinate system centered on the GC with coordinates (r, z) as $r = d_{\text{GC}} \tan(\ell)$ and $z = d_{\text{GC}} \tan(b) / \cos(\ell)$, where we adopt $d_{\text{GC}} = 8.2$ kpc as the distance to the GC (Bland-Hawthorn & Gerhard 2016). The distance from the Sun to each pixel is then $\sqrt{d_{\text{GC}}^2 + r^2 + z^2}$.

The thickness of the FBs is unknown, and likely varies with projected position (ℓ, b) . However, we must make some choice for thickness, because this choice influences the mean ISRF intensity that we must know in order to model the CRE population that produces the observable emission in each pixel – the Popescu et al. (2017) radiation field that we use is a function of (r, z) coordinates, so if the emitting volume that contributes to the observable γ -ray emission from a given pixel extends to a distance d_{thick} on either side of the plane, then the mean ISRF inside the emitting volume is

$$\langle u_\lambda \rangle(\ell, b) = \frac{1}{2d_{\text{thick}}} \int_{-d_{\text{thick}}}^{d_{\text{thick}}} u_\lambda(r', z) dx, \quad (1)$$

where $r' = \sqrt{r(l, b)^2 + x^2}$. In the absence of a better option, we assume that d_{thick} is comparable to the width of the FBs along the plane of the sky, and thus adopt $d_{\text{thick}} = 2$ kpc. We use this assumption to compute $\langle u_\lambda \rangle$ for each pixel.

Hadronic models of the FBs in turn require knowledge of the gas density. Unfortunately this is completely unknown. For this reason, we will adopt a fiducial uniform number density of H nuclei $n_{\text{H}} = 10^{-2} \text{ cm}^{-3}$. Because this is so uncertain, we will report all quantities for hadronic models normalized to this density.

⁶ <https://galprop.stanford.edu>

3.2. Model fits

We use the *Naima* Python package (Zabalza 2015) to forward model the non-thermal γ -ray emission from relativistic proton and electron distributions with prescribed spectral parameters. For hadronic models we use the *Naima* `PionDecay` class, which calculates the γ -ray spectrum resulting from the interaction of a relativistic CRp population with ambient protons using the parameterization of Kafexhiu et al. (2014); as noted above, for hadronic models we normalize to a fiducial proton density $n_H = 10^{-2} \text{ cm}^{-3}$. For leptonic models, we use the *Naima* `InverseCompton` class to calculate the γ -ray spectrum produced by IC scattering of seed photon fields by an assumed CRe distribution. This calculation uses the differential cross-section provided by Khangulyan et al. (2014) for blackbody photon fields and a cross-section developed by Aharonian & Atoyan (1981) for non-thermal photon fields. We use the former capability to compute γ -ray production by IC scattering of the CMB, and the latter to compute production by IC scattering of the ISRF, using the mean ISRF energy density $\langle u_\lambda \rangle$ for each pixel derived as described in Section 3.1.

We consider three models for the energy distributions of parent CRp and CRe populations: (i) power-law (PL), (ii) power-law with an exponential cut-off (EPL), and (iii) broken power-law (BPL). The functional forms and free parameters for these models are

PL: $N(E) = A(E/E_0)^{-\alpha}$, where α is a free spectral index.

EPL: $N(E) = A(E/E_0)^{-\alpha} \exp(-E/E_{\text{cut}})$, where α is a free spectral index and E_{cut} is a free cut-off energy.

BPL: $N(E) = A(E/E_0)^{-\alpha_1}$ for $E < E_{\text{break}}$ and $N(E) = A(E_{\text{break}}/E_0)^{\alpha_2 - \alpha_1} (E/E_0)^{-\alpha_2}$ for $E > E_{\text{break}}$, where α_1 and α_2 are free spectral indices and E_{break} is a free break energy.

In these expressions, $N(E)$ is the differential number density of CR particles as a function of energy E , $E_0 = 10$ GeV is an arbitrary reference energy, and A is a free amplitude. Thus the PL model has two free parameters (A and α), the EPL model has three (A , α , E_{cut}), and the BPL model has four (A , α_1 , α_2 , E_{break}).

We find the best fit model parameters for each pixel by fitting the γ -ray spectrum calculated using *Naima* to the data using the *SciPy* `curve_fit` routine, which uses non-linear least squares fitting. We evaluate the goodness of fit of each model by computing the reduced χ^2 value, given by $\chi_{\text{red}}^2 = (1/10) \sum_{j=1}^{10} (\Phi_{\text{obs},j} - \Phi_{\text{mod},j})^2 / \sigma_j^2$, where $\Phi_{\text{mod},j}$ and $\Phi_{\text{obs},j}$ are the model-predicted and observed (taken from the M2 non-dust component – cf. Section 2.1) photon fluxes in the j th energy bin (out of 10 total) for each pixel, and σ_j is the corresponding uncertainty. However, this quantity should be interpreted with considerable caution, because the errors σ_j are *not* predominantly random measurement errors; instead they are dominated by uncertainties in the decomposition, which means that uncertainties in adjacent energy bins are highly-correlated rather than independent. Consequently, we cannot give the reduced χ^2 its usual interpretation in terms of an absolute goodness of fit; we will instead simply use it to characterize relative goodness of fit between alternative models.




FIG. 3.— The reduced χ^2 for each model and source particle population (left panel for protons and right panel for electrons, PL, BPL, and EPL levels from top to bottom) for each pixel. The boundary of the FBs is marked by a yellow line, which is the same in each panel and the same as the black line in Figure 1. We see that a power law model with a break or an exponential cutoff provides a better fit than a simple PL model in the FBs region.

4. RESULTS

Having carried out fits to each pixel using each parent population model as described in Section 3, we are now in a position to examine what the fits tell us.

4.1. Which models can explain the data?

We start by discussing how well each model fits the data. Figure 3 shows the value of χ_{red}^2 for our six models – three functional forms (PL, EPL, BPL) times two particle types (protons and electrons) – in every pixel. We remind readers that we do not expect values around




FIG. 4.— Best-fit model spectra versus energy-binned data for the six different models evaluated at the point $b = 45^\circ$ and $l = 0.38^\circ$ (labeled ‘a’ in Figure 1) within the FBs. The blue points with error bars show the γ -ray data, and the computed SED is shown by solid black (PL), green (EPL), and red (BPL) lines. The EPL and BPL models perform better than the PL models for both hadronic (left) and leptonic (right) cases.




FIG. 5.— Best-fit parameters for the EPL models (protons on the left and electrons on the right). The plots in the top panel shows the model amplitude and the middle panel shows the power-law index. We see a spectral hardening towards the Southern tip for both. The bottom panel shows the cut-off energy. It is of the order of a few TeV and does not vary much with latitude.

unity for well-fitting models because the uncertainties are correlated, and so the main quantity of interest in this feature is the relative value of χ^2_{red} , not the absolute value. With this caveat in mind, we find that in general the BPL and EPL models are noticeably better fits

within the FBs than a simple PL model for both electrons and protons, but that there are no significant differences in fit quality between hadronic and leptonic models.

We show a typical fit in Figure 4, which compares the best-fit model spectra to the data for the pixel at $b = 45^\circ$ and $l = 0.38^\circ$ (labeled as ‘a’ in Figure 1), a point within the FBs. We see that the γ -ray data are not well accommodated by radiation from power-law distributions of either protons or electrons. We henceforth discard PL models. On the other hand, our results do not obviously prefer EPL over BPL models or vice versa, nor do they show obvious preferences for hadronic versus leptonic models; we therefore discuss best-fit parameters for all these cases next.

4.2. Inferences from acceptable models

Figure 5 shows best-fit amplitudes, power-law indices, and cutoff energies for the EPL models. Not unexpectedly, best-fit CRp models have harder spectra than best-fit CRe models. Best-fit spectra harden towards the tip of the Southern bubble for all models. This mirrors a spectral hardening that was previously found in the *Fermi* γ -ray data (Yang et al. 2014; Selig et al. 2015; Platz et al. 2023). Cutoff energies are of the order of a few TeV within the FBs for both hadronic and leptonic cases. In contrast to spectral indices, best-fit break energies do not show much variation with latitude inside the FBs. Unlike in the projected interior, outside the FBs, break energies are not robustly determined. In particular, for regions outside the FBs best-fit cutoff energies tend to very high values and fits return a large standard deviations, consistent with the γ -ray sky outside the FBs being well explained by simple power law distributions of either electrons or protons over the 1-200 GeV energy range. Best-fit model amplitudes increase with latitude for CRe models while, in contrast, they are maximum in the Galactic plane for CRp models. We will shortly see that the particle energy density follows a similar pattern as the amplitude. However, we caution that the trend for the CRp models is dependent on our assumption of a con-




FIG. 6.— Best-fit parameters for the BPL models (protons on the left and electrons on the right). The top panel shows the model amplitude. The second and fourth panel show the spectral index before and after the break respectively. The break energy does not show much variation along the length of the FBs as shown in the third panel.

stant target proton number density, and would change if we were to instead assume a density gradient.

Figure 6 shows the corresponding best-fit parameters for the BPL models. Here, we find break energies of the order of ~ 100 GeV for both cases with little latitudinal variation. Best fit spectral indices below the break are more or less consistent with spectral indices in the EPL model showing qualitatively the same hardening for large negative latitudes in the southern Bubble. In contrast, spectral indices after the break do not show variation with latitude. Figure 7 shows the spatial variation of the difference between spectral indices before and after the break. The difference is almost zero outside the Bubble boundary for the electrons, again indicating that, while CRe models require a break inside the FBs, outside, simple PL models adequately explain the data. Within the bubbles, the difference in spectral index varies from $\approx 0.1 - 0.6$. Interestingly, the fits strongly rule out spectral index changes as large as unity, the value that would be expected for a synchrotron or inverse Compton cooling break – this will become important for our




FIG. 7.— The difference between spectral indices between the two powerlaw components in the BPL model (protons on left and electrons on right).

discussion of the implications of our findings below.

One important byproduct of our fits, which will be important for our interpretation, is an understanding of which radiation field components drive which parts of the observed γ -ray spectrum in leptonic models. Figure 8 shows the contribution of the different components of the ISRF – stellar, dust, and CMB – to the total γ -ray emission at the four points labeled ‘a’ through ‘d’ in Figure 1. Since the stellar and dust components both contribute about equally at μm wavelengths, for this purpose we roughly divide them at $3 \mu\text{m}$, i.e., we define the stellar component as the portion of the ISRF at wavelengths $< 3 \mu\text{m}$, and the dust component as emission at $3 - 300 \mu\text{m}$, where $300 \mu\text{m}$ is roughly the wavelength at which the CMB begins to dominate. The energy-weighted mean photon energies of the stellar, dust and CMB components defined in this way are 1.6 eV, 0.1 eV, and 9.3×10^{-4} eV respectively. We find that the dust component dominates at all γ -ray energies observable by *Fermi*/LAT, followed by the starlight component and lastly the CMB. Interestingly this is true even at position ‘c’, where the starlight field overall has a significantly larger energy density than the dust field (cf. Figure 2). This is because interaction between CRe and the starlight field is strongly suppressed by Klein-Nishina effects, particularly for the higher initial photon and CRe energies that drive emission towards the top of the observed γ -ray band. By contrast, Klein-Nishina effects are essentially never important for the much lower-energy dust photons.

We can understand the difference in behavior by computing the Klein-Nishina parameter $\Gamma_e = 4E_i E_e / m_e^2 c^4$, where E_i and E_e are the initial photon and CRe energies, respectively; Klein-Nishina suppression is important for $\Gamma_e \gtrsim 1$. For the few eV photons typical of the starlight field, Γ_e is greater than unity for CRe energies $\gtrsim 30$ GeV, well below the break or cutoff energies of our best-fitting BPL or EPL models; by contrast, for the few $\times 10^{-2}$ eV photons that characterize the dust field, $\Gamma_e \lesssim 1$ for electron energies up to $\sim \text{TeV}$, above the break energy and comparable to the cutoff energies for our best-fitting BPL and EPL models. An important implication of this finding is that we can to good approximation assume that IC losses are in the Thomson regime throughout the FBs, since the dominant losses are due to dust-produced pho-




FIG. 8.— Contributions of the stellar (green), dust (red), and CMB (purple) components of the ISRF to the total γ -ray emission (orange best-fit line) for the best-fitting leptonic model at the locations marked ‘a’ - ‘d’ in Figure 1. Blue points with error bars show the observed γ -ray emission. The dust component has the highest contribution at all locations, followed by the stellar component (which is suppressed by the Klein-Nishina effect); the CMB component is always sub-dominant.

tons rather than starlight ones.

5. DISCUSSION

We next discuss some of the implications of our results, as well as comparing to earlier work.

5.1. The energy budget of the Fermi bubbles

To understand the implications of our analysis for the origins of the FBs, a useful first step is to estimate the total energies and energy densities of CRs inside the FBs in either the hadronic or leptonic scenarios. Such an estimate will necessarily be approximate: our best-fitting powerlaw indices at low energy are steeper than two, and thus the total energy depends on the minimum energy below which the powerlaw breaks. We can obtain only a weak upper limit on the presence of such a break from the fact that no break is visible in the γ -ray data, for which our lowest energy bin is $E_\gamma = 0.56$ GeV. In a hadronic model, γ -rays of this energy will be produced primarily by protons of $\approx 5 - 6$ GeV (Kafexhiu et al. 2014). In a leptonic model where the γ -rays come primarily from IC scattering in the Thomson regime, the maximum γ -ray energy $E_{\gamma,\max}$ that can be produced by up-scattering of a photon with initial energy E_i by a CR

electron of energy E_e is $E_{\gamma,\max} = 4E_i(E_e/m_e c^2)^2$ (Blumenthal & Gould 1970); using $E_{\gamma,\max} = 0.56$ GeV and $E_i \approx 0.03$ eV, typical of our dominant dust-produced photon field, gives $E_e \approx 30$ GeV. Given the weakness of these limits, and the fact that they are different between the leptonic and hadronic models, we instead elect to calculate the energy budgets for a nominal minimum CR energy $E_{\min} = 1$ GeV. For our measured powerlaw indices – $\alpha \approx \alpha_1 \approx 2.2 - 2.3$ for protons, and ≈ 3 for electrons – our estimate for the total energy in a hadronic model is weakly dependent on this choice, $E_{\text{tot}} \propto E_{\min}^{-k}$ with $k \approx 0.2 - 0.3$, while our estimate in the leptonic case is more strongly dependent, $k \approx 1$. On the other hand, in the hadronic case the total energy will scale with our assumed number density as $E_{\text{tot}} \propto n_H^{-1}$.

We use these assumed lower limits together with our best-fit model parameters in the EPL and BPL models to compute the total CR energy in each pixel, and we sum over all the pixels inside the FBs (those within the white contours in Figure 1) to obtain total energies of $1.4 \times 10^{56}/n_{H,-2}$ erg, $1.2 \times 10^{56}/n_{H,-2}$ erg, 2.4×10^{53} erg, and 2.3×10^{53} erg in the hadronic EPL, hadronic BPL, leptonic EPL, and leptonic BPL scenarios, respectively,




FIG. 9.— The line of sight depth of each pixel, calculated by assuming the 3D structure of the FBs formed by rotating its boundary about the z -axis.

where $n_{\text{H},-2}$ is the number density of H nuclei normalized by 0.01 cm^{-3} . With some additional assumptions about geometry, we can also use this procedure to estimate the energy density within the FBs. For this purpose, we assume that the FBs are cylindrically-symmetric, and compute thickness of every pixel by rotating the projected boundary shown in Figure 1 about the z -axis. We show the resulting thicknesses in Figure 9. We treat each pixel as a prism with this thickness and an area equal to its area on the sky computed at a distance equal to the distance between the Sun and the pixel (Section 3.1),⁷ and from this plus the total energy content of the pixel as deduced from our fits we compute the energy density. We show the result in Figure 10 for the EPL model; the results for the BPL model are qualitatively similar.

We can see that the inferred energy densities follow the pattern of the model amplitudes in Figure 5. Here we caution that CRp energy density plot should not be over-interpreted given we adopt a fixed target gas density (while the true ambient gas number density distribution is likely strongly position-dependent). That said, it is interesting to compare the pressures implied by the results in Figure 5 to gas pressures. For our nominal $n_{\text{H}} = 10^{-2} \text{ cm}^{-3}$, the CR pressure is one third of the value shown in the right panel of this figure; the range we find, $U_{\text{CR}} \approx (1-3)n_{\text{H},-2} \text{ eV cm}^{-3}$, corresponds to $P_{\text{CR}}/k_{\text{B}} \approx (0.4-1)n_{\text{H},-2} \times 10^4 \text{ K cm}^{-3}$, while the pres-

sure of fully-ionized⁸ thermal gas with temperature T is $P_{\text{th}}/k_{\text{B}} = 2.3n_{\text{H}}T = 2.3n_{\text{H},-2}T_6 \times 10^4 \text{ K cm}^{-3}$, where $T_6 = T/(10^6 \text{ K})$. Thus for our nominal n_{H} , thermal gas pressure is likely dominant over CR pressure, though not by a huge amount. However, since $P_{\text{CR}}/P_{\text{th}} \propto n_{\text{H}}^{-2}$, CR pressure will dominate if the gas density is lower than our fiducial estimate.

For leptonic models we can be considerably more confident about the variation of energy density with position, since we use a self-consistent, spatially-dependent ISRF model. For this case, a striking feature is that the CRe energy density peaks towards the edges of the FBs, qualitatively consistent with limb brightening, one interpretation of which is emission from a shell structure (cf. Su et al. 2010). Thus not only does the CRe spectrum harden toward the FB edges, the total energy content increases as well.

5.2. Cooling timescales

In addition to energy budgets, we can also calculate characteristic cooling timescales for the non-thermal particle population powering the FBs, and the combination of these two provides useful constraints on their possible origin. For the hadronic case, we calculate the luminosity at each pixel by multiplying the observed intensity of the non-dust-associated emission component by the area of each pixel and numerically integrating over energy bins. For our fiducial distance $d_{\text{GC}} = 8.2 \text{ kpc}$, this gives a total luminosity of $3.5 \times 10^{37} \text{ erg s}^{-1}$. The statistical error on this is very small (order $10^{33} \text{ erg s}^{-1}$) due to the large number of pixels, but as noted above the dominant errors are systematics in the decomposition of the γ -ray sky into components, which render the errors highly-correlated; thus a more realistic uncertainty estimate is of order of the sizes of the correlated error bars, which is tens of percent. We have to apply two corrections to this luminosity for full energy accounting. First, we have to make a bolometric correction ~ 2 (given the inferred CRp distribution) to account for production of γ -rays above and below our sensitivity band. Second, given that hadronic γ -ray emission involves production of approximately equal numbers of π^+ and π^- and π^0 , we have to factor in a further factor of ~ 3 to correct for the power ultimately liberated in neutrinos and secondary electrons and positrons in addition to gamma-rays. We find, therefore, that the total pp power loss within the FBs is $\approx 2.1 \times 10^{38} \text{ erg s}^{-1}$. To estimate the time required for the current CR proton population to cool under this loss, we compute the CRp energy in protons with energies $> 5.6 \text{ GeV}$ – the energy range that contributes to the observable emission – using the same method as described in Section 5.1. Doing so gives a total energy budget of $6.2 \times 10^{55} \text{ erg}$ over this energy range, and thus a characteristic hadronic cooling timescale

$$t_{\text{had}} \sim \frac{6.2 \times 10^{55} n_{\text{H},-2}^{-1} \text{ erg}}{2.1 \times 10^{38} \text{ erg s}^{-1}} = 9.4 n_{\text{H},-2}^{-1} \text{ Gyr} \quad (2)$$

Note that this cooling time is equivalent to the time to reach saturation at the given n_{H} defined by Crocker & Aharonian (2011), and that, while the absolute energies

⁷ In principle we should adjust this calculation to account for the fact that we are not located at infinity relative to the FBs, so the path length of a given line of sight through the FBs is not identical to the thickness. However, given the crudeness of our geometric assumption, such an increase in complexity does not seem warranted.

⁸ Meaning that there are 2.3 free particles per H nucleon for the standard He/H ratio.




FIG. 10.— CR energy density inferred for the best-fitting leptonic (left) and hadronic (right) EPL models. Note the normalization of the CRp population with respect to the fiducial value of the ambient hydrogen number density (n_H).

we have quoted depend on the distance to the FBs, the cooling timescale does not, since the numerator and denominator both depend on distance in the same way.

For leptonic models, given the strong decrease in cooling time with particle energy, it is more useful to estimate the cooling time for the highest-energy electrons than for the population as a whole. The question therefore becomes how to identify the highest energies that contribute to the observed emission, which is model-dependent. In the BPL model, we cannot identify the break energy with a cooling break, since the difference in spectral indices between the upper and lower powerlaws is $\lesssim 0.5$, far smaller than the value of unity expected for an IC cooling break in the Thomson regime (cf. Figure 7), and so the population must be un-cooled up to the highest energies that contribute to any of the *Fermi*/LAT data. For the EPL model, we can identify the highest electron energy with the cutoff energy $E_{\text{cut}} \approx 1 - 2$ TeV. Since the latter is more permissive regarding the cooling timescale, we calculate the cooling time in the EPL model. We use *GAMERA*⁹ to compute the total IC energy loss rate $\dot{E}_{\text{IC}}$ for a CRe with energy E_{cut} in each pixel, and define the cooling time as $E_{\text{cut}}/\dot{E}_{\text{IC}}$. We plot this cooling time in Figure 11. We emphasise that this IC cooling time actually defines an upper limit to the total loss time given the existence of magnetic fields suffusing the Bubbles (on which CRe can synchrotron radiate) with characteristic energy densities similar to those of the dust field (Carretti et al. 2013).

5.3. Implications of the cooling timescales

As has long been known (Crocker & Aharonian 2011; Sarkar 2024), the CRp power required to sustain the γ -ray emission from the FBs in a hadronic scenario –

⁹ <https://libgamera.github.io/GAMERA/docs/documentation.html>




FIG. 11.— Inverse Compton cooling time of CRe with energy E_{cut} .

$2.1 \times 10^{38} \text{ erg s}^{-1}$ – is rather modest and could be sustained by the current rate of nuclear star formation, even if we have to make upwards corrections of order a few to account for additional Coulombic/ionisation losses. On the other hand, as has also long been appreciated, the timescale to reach a hadronic steady state (as set by the pp loss time) is long given the low density of the volume-filling phase of the Bubbles; given that, one must either countenance that the Bubbles are very old (multi-Gyr) structures and/or quite far from steady state (and or arrange that the CRp's are concentrated into regions of sys-

tematically higher density than the volume-filling phase; cf. Crocker et al. 2014, 2015). In the latter case, energy/power demands become increasingly difficult to be accommodated by star formation the smaller the Bubbles’ age. Another concern would be the additional adiabatic losses suffered by the CRp’s in the case that they are rapidly advected out of the nucleus co-mingled with expanding thermal plasma and magnetic fields.

In the case that the timescale for the formation of the Bubbles becomes much less than, say, ~ 100 Myr, nuclear or inner Milky Way star formation is increasingly strongly excluded and one is pushed towards invoking the Galaxy’s supermassive black hole as their ultimate generator. An AGN scenario origin for the Bubbles could be either leptonic or hadronic (or both), of course.

For the case that the γ -ray emission from the Bubbles is primarily leptonic, we find that one needs an increasingly large CRe energy density towards the Bubbles’ edges; *a priori* this seems hard to accommodate within scenarios where one advects CRe’s out of the nucleus, but it can be accommodated accounting for non-steady state effects as demonstrated, e.g., by the model of Yang & Ruszkowski (2017). Probably the most severe test of any scenario invoking leptonic emission from advected, nuclear CRe’s is the short cooling time of electrons radiating at ~ 200 GeV. This is $\lesssim 1$ Myr at the top of the Bubbles accounting only for IC losses; factoring in synchrotron losses (certainly present at some level), bremsstrahlung and ionisation (important at lower latitudes/earlier in the expansion for an explosion), even stronger IC losses in the higher energy density fields closer to the plane (through which the CRe’s must pass), and adiabatic losses (which, again, should be important in an explosion), the real, effective loss time is likely significantly smaller. Transport to 8 kpc (the top of the Bubbles) within 1 Myr implies an expansion speed ~ 8000 km/s, or $\sim 0.03c$, a highly supersonic flow that seems difficult to reconcile with the much lower Mach numbers of expansion inferred from X-ray phenomenology near the Bubbles’ edges (see the extended discussion in Sarkar 2024). Overall, to accommodate these transport time constraints and the peaking of the putative CRe’s energy density towards or at the Bubbles’ edges, it seems more natural to us – as it has also seemed to many previous authors (Sarkar 2024, and references therein) – that there be distributed, *in situ* acceleration concentrated towards the Bubbles’ periphery, however that is achieved.

6. SUMMARY AND CONCLUSIONS

In this work, we investigate the cosmic ray distributions required to produce the measured γ -ray emission in the Fermi Bubbles (FBs) on a pixel-by-pixel basis, leveraging a recent template-free reconstruction of *Fermi* data by Platz et al. (2023) together with a modern model for the position-dependent interstellar radiation field from Popescu et al. (2017). We consider both leptonic models in which IC scattering is the primary emission mechanism and hadronic ones in which pion production dominates, and derive pixel-by-pixel fits for the cosmic ray

electron (CRe) or proton (CRp) spectra over 37,714 pixels, thereby tracing out the variations in these spectra across the FBs.

We show that simple power law models for both CRe and CRp energy distributions predict γ -ray spectra that provide poor fits to the data, enabling us to rule them out, but that power laws with exponential cutoffs and broken power law models both provide similarly good fits to the data; the fits are indistinguishably good for both CRe and CRp scenarios. Best-fit spectral indices at the low-energy end for either class of model show hardening towards the caps of the FBs, particularly in the South, for both CRp and CRe models. This result disagrees with the earlier findings of Narayanan & Slatyer (2017), who bin the data by latitude and do not find any spectral hardening. Best-fit cutoff energies in exponential cutoff models are of the order of a few TeV, while break energies in broken powerlaw models are a few hundred GeV; neither exhibit much variation with latitude. The change in spectral index for broken powerlaw models is $\lesssim 0.5$, inconsistent with an inverse Compton or synchrotron cooling break.

Our fits allow us to calculate the total energy, the position-dependent energy densities, and the cooling times for CRs in the FBs. While these are significantly uncertain for hadronic models due to the unknown distribution of thermal background gas, for leptonic models – where the uncertainty is smaller because the radiation field is better known – the fits impose important constraints. The data require the CRe energy density to increase with distance from the plane, particularly towards the edge of the bubbles, with a variation as a function of position that is morphologically similar to the spectral hardening. Moreover, the cooling times are extremely short – ≈ 1 Myr at the bubble caps, but falling to ≈ 0.1 Myr near the Galactic center. These findings are difficult to accommodate in a model where the CRe driving the emission are injected at the nucleus and then transported to the bubble edge, and instead are most naturally explained by either *in situ* acceleration of leptonic CRs at the bubble caps, or by a hadronic scenario that does not suffer from such severe cooling limits.

ACKNOWLEDGMENTS

AT acknowledges the support of the ANU Future Research Talent (FRT) program, which provided the funding and the opportunity to undertake part of this research at the Australian National University. RMC and MRK acknowledge support from the Australian Research Council through the *Discovery Projects* scheme, award DP230101055. This research was undertaken with the assistance of resources from the National Computational Infrastructure (NCI Australia), an NCRIS enabled capability supported by the Australian Government, through award jh2.

DATA AVAILABILITY

Our best-fitting parameters for the leptonic and hadronic EPL and BPL models for all pixels are available from <https://doi.org/10.5281/zenodo.17568107>.

REFERENCES

- Aharonian, F. A., & Atoyan, A. M. 1981, *Ap&SS*, 79, 321, doi: [10.1007/BF00649428](https://doi.org/10.1007/BF00649428)
- Bland-Hawthorn, J., & Cohen, M. 2003, *ApJ*, 582, 246, doi: [10.1086/344573](https://doi.org/10.1086/344573)

- Bland-Hawthorn, J., & Gerhard, O. 2016, *ARA&A*, 54, 529, doi: [10.1146/annurev-astro-081915-023441](https://doi.org/10.1146/annurev-astro-081915-023441)
- Blumenthal, G. R., & Gould, R. J. 1970, *Rev. Mod. Phys.*, 42, 237, doi: [10.1103/RevModPhys.42.237](https://doi.org/10.1103/RevModPhys.42.237)
- Bordoloi, R., Fox, A. J., Lockman, F. J., et al. 2017, *ApJ*, 834, 191, doi: [10.3847/1538-4357/834/2/191](https://doi.org/10.3847/1538-4357/834/2/191)
- Carretti, E., Crocker, R. M., Staveley-Smith, L., et al. 2013, *Nature*, 493, 66, doi: [10.1038/nature11734](https://doi.org/10.1038/nature11734)
- Cashman, F. H., Fox, A. J., Savage, B. D., et al. 2021, *ApJ*, 923, L11, doi: [10.3847/2041-8213/ac3cbc](https://doi.org/10.3847/2041-8213/ac3cbc)
- Crocker, R. M., & Aharonian, F. 2011, *Phys. Rev. Lett.*, 106, 101102, doi: [10.1103/PhysRevLett.106.101102](https://doi.org/10.1103/PhysRevLett.106.101102)
- Crocker, R. M., Bicknell, G. V., Carretti, E., Hill, A. S., & Sutherland, R. S. 2014, *ApJ*, 791, L20, doi: [10.1088/2041-8205/791/2/L20](https://doi.org/10.1088/2041-8205/791/2/L20)
- Crocker, R. M., Bicknell, G. V., Taylor, A. M., & Carretti, E. 2015, *ApJ*, 808, 107, doi: [10.1088/0004-637X/808/2/107](https://doi.org/10.1088/0004-637X/808/2/107)
- Di Teodoro, E. M., McClure-Griffiths, N. M., Lockman, F. J., & Armillotta, L. 2020, *Nature*, 584, 364, doi: [10.1038/s41586-020-2595-z](https://doi.org/10.1038/s41586-020-2595-z)
- Dobler, G. 2012, *ApJ*, 750, 17, doi: [10.1088/0004-637X/750/1/17](https://doi.org/10.1088/0004-637X/750/1/17)
- Finkbeiner, D. P. 2004, *ApJ*, 614, 186, doi: [10.1086/423482](https://doi.org/10.1086/423482)
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759, doi: [10.1086/427976](https://doi.org/10.1086/427976)
- Guo, F., & Mathews, W. G. 2012, *ApJ*, 756, 181, doi: [10.1088/0004-637X/756/2/181](https://doi.org/10.1088/0004-637X/756/2/181)
- Herold, L., & Malyshev, D. 2019, *A&A*, 625, A110, doi: [10.1051/0004-6361/201834670](https://doi.org/10.1051/0004-6361/201834670)
- Kafexhiu, E., Aharonian, F., Taylor, A. M., & Vila, G. S. 2014, *Phys. Rev. D*, 90, 123014, doi: [10.1103/PhysRevD.90.123014](https://doi.org/10.1103/PhysRevD.90.123014)
- Khangulyan, D., Aharonian, F. A., & Kelner, S. R. 2014, *ApJ*, 783, 100, doi: [10.1088/0004-637X/783/2/100](https://doi.org/10.1088/0004-637X/783/2/100)
- Mertsch, P., & Petrosian, V. 2019, *A&A*, 622, A203, doi: [10.1051/0004-6361/201833999](https://doi.org/10.1051/0004-6361/201833999)
- Mertsch, P., & Sarkar, S. 2011, *Phys. Rev. Lett.*, 107, 091101, doi: [10.1103/PhysRevLett.107.091101](https://doi.org/10.1103/PhysRevLett.107.091101)
- Mori, M. 1997, *ApJ*, 478, 225, doi: [10.1086/303785](https://doi.org/10.1086/303785)
- Mou, G., Yuan, F., Gan, Z., & Sun, M. 2015, *ApJ*, 811, 37, doi: [10.1088/0004-637X/811/1/37](https://doi.org/10.1088/0004-637X/811/1/37)
- Narayanan, S. A., & Slatyer, T. R. 2017, *MNRAS*, 468, 3051, doi: [10.1093/mnras/stx577](https://doi.org/10.1093/mnras/stx577)
- Niederwanger, F., Reimer, O., Kissmann, R., et al. 2019, *Astropart. Phys.*, 107, 1, doi: [10.1016/j.astropartphys.2018.11.001](https://doi.org/10.1016/j.astropartphys.2018.11.001)
- Noon, K. A., Krumholz, M. R., Di Teodoro, E. M., et al. 2023, *MNRAS*, 524, 1258, doi: [10.1093/mnras/stad1890](https://doi.org/10.1093/mnras/stad1890)
- Planck Collaboration. 2013, *A&A*, 554, A139, doi: [10.1051/0004-6361/201220271](https://doi.org/10.1051/0004-6361/201220271)
- Platz, L. I., Knollmüller, J., Arras, P., et al. 2023, *A&A*, 680, A2, doi: [10.1051/0004-6361/202243819](https://doi.org/10.1051/0004-6361/202243819)
- Popescu, C. C., Yang, R., Tuffs, R. J., et al. 2017, *MNRAS*, 470, 2539, doi: [10.1093/mnras/stx1282](https://doi.org/10.1093/mnras/stx1282)
- Predehl, P., Sunyaev, R. A., Becker, W., et al. 2020, *Nature*, 588, 227, doi: [10.1038/s41586-020-2979-0](https://doi.org/10.1038/s41586-020-2979-0)
- Sarkar, K. C. 2024, *A&A Rev.*, 32, 1, doi: [10.1007/s00159-024-00152-1](https://doi.org/10.1007/s00159-024-00152-1)
- Selig, M., Vacca, V., Oppermann, N., & Enßlin, T. A. 2015, *A&A*, 581, A126, doi: [10.1051/0004-6361/201425172](https://doi.org/10.1051/0004-6361/201425172)
- Snowden, S. L., Egger, R., Freyberg, M. J., et al. 1997, *ApJ*, 485, 125, doi: [10.1086/304399](https://doi.org/10.1086/304399)
- Su, M., Slatyer, T. R., & Finkbeiner, D. P. 2010, *ApJ*, 724, 1044, doi: [10.1088/0004-637X/724/2/1044](https://doi.org/10.1088/0004-637X/724/2/1044)
- Vijayan, A., Krumholz, M. R., & Wibking, B. D. 2024, *MNRAS*, 527, 10095, doi: [10.1093/mnras/stad3816](https://doi.org/10.1093/mnras/stad3816)
- Yang, H.-Y. K., & Ruszkowski, M. 2017, *ApJ*, 850, 2, doi: [10.3847/1538-4357/aa9434](https://doi.org/10.3847/1538-4357/aa9434)
- Yang, H. Y. K., Ruszkowski, M., & Zweibel, E. 2013, *MNRAS*, 436, 2734, doi: [10.1093/mnras/stt1772](https://doi.org/10.1093/mnras/stt1772)
- Yang, H.-Y. K., Ruszkowski, M., & Zweibel, E. 2022, *Nature Astron.*, 6, 584, doi: [10.1038/s41550-022-01618-x](https://doi.org/10.1038/s41550-022-01618-x)
- Yang, R.-z., Aharonian, F., & Crocker, R. 2014, *A&A*, 567, A19, doi: [10.1051/0004-6361/201423562](https://doi.org/10.1051/0004-6361/201423562)
- Zabalza, V. 2015, in *International Cosmic Ray Conference*, Vol. 34, 34th International Cosmic Ray Conference (ICRC2015), 922, doi: [10.22323/1.236.0922](https://doi.org/10.22323/1.236.0922)

provides fast and easy peer review for new papers in the **astro-ph** section of the arXiv, making the reviewing process simpler for authors and referees alike. Learn more at <http://astro.theoj.org>.

This paper was built using the Open Journal of Astrophysics L^AT_EX template. The OJA is a journal which