The Fermi Haze: A Gamma-Ray Counterpart to the Microwave Haze

Citation

Dobler, Gregory, Douglas P. Finkbeiner, Ilias Cholis, Tracy Slatyer, and Neal Weiner. 2010. "The Fermi Haze: A Gamma-Ray Counterpart to the Microwave Haze." *The Astrophysical Journal* 717 [2] (June 18): 825–842. doi:10.1088/0004-637x/717/2/825.

Published Version

doi:10.1088/0004-637X/717/2/825

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:33462895>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

THE *FERMI* HAZE: A GAMMA-RAY COUNTERPART TO THE MICROWAVE HAZE

GREGORY DOBLER^{1,2}, DOUGLAS P. FINKBEINER^{1,3}, ILIAS CHOLIS⁴, TRACY SLATYER^{1,3}, AND NEAL WEINER⁴

¹ Institute for Theory and Computation, Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, MS-51 Cambridge, MA 02138, USA; dobler@kitp.ucsb.edu

² Kavli Institute for Theoretical Physics, University of California, Santa Barbara Kohn Hall, Santa Barbara, CA 93106, USA

³ Physics Department, Harvard University, Cambridge, MA 02138, USA

⁴ Center for Cosmology and Particle Physics, Department of Physics, New York University, New York, NY 10003, USA

Received 2009 October 26; accepted 2010 April 26; published 2010 June 18

ABSTRACT

The *Fermi* Gamma-ray Space Telescope reveals a diffuse inverse Compton (IC) signal in the inner Galaxy with a similar spatial morphology to the microwave haze observed by *WMAP*, supporting the synchrotron interpretation of the microwave signal. Using spatial templates, we regress out π^0 gammas, as well as IC and bremsstrahlung components associated with known soft-synchrotron counterparts. We find a significant gamma-ray excess toward the Galactic center with a spectrum that is significantly harder than other sky components and is most consistent with IC from a hard population of electrons. The morphology and spectrum are consistent with it being the IC counterpart to the electrons which generate the microwave haze seen at *WMAP* frequencies. In addition, the implied electron spectrum is hard; electrons accelerated in supernova shocks in the disk which then diffuse a few kpc to the haze region would have a softer spectrum. We describe the full-sky *Fermi* maps used in this analysis and make them available for download.

Key words: gamma rays: general – gamma rays: ISM

Online-only material: color figures

1. INTRODUCTION

The most detailed and sensitive maps of diffuse microwave emission in our Galaxy have been produced by the *Wilkinson Microwave Anisotropy Probe* (*WMAP*). An analysis of the different emission mechanisms in these maps uncovered a microwave “haze” toward the Galactic center (GC) that has roughly spherical morphology and radius ~ 4 kpc (Finkbeiner 2004a). Since its discovery, Finkbeiner (2004b) and Dobler & Finkbeiner (2008) have argued that the microwave haze is hard synchrotron emission due to the fact that alternative hypotheses such as free-free (thermal bremsstrahlung of the ionized gas), spinning dust, or thermal dust have difficulty explaining its morphology, spectrum, or both.

However, the 23–33 GHz spectrum of the haze synchrotron is harder than that expected from diffused electrons originally accelerated by supernova (SN) shocks in the plane (Dobler & Finkbeiner 2008). Some variation is expected in the synchrotron spectrum, but generally in the sense that it should be *softer* at higher frequencies since the electrons lose energy preferentially at high energies as they diffuse from their source. While there are significant uncertainties in the spectrum of the haze (see the discussion in Dobler & Finkbeiner 2008), the data require that the diffused spectrum be roughly as hard as the expected *injection* spectrum from first-order *Fermi* acceleration at SN shock fronts (number density $dN/dE \propto E^{-2}$). That is, if the electrons were produced in shocks in the disk, then they would have to have undergone no diffusive energy losses over a $\sim 4\pi/3$ (4 kpc)³ volume, which seems unlikely. Furthermore, there is significant emission in *WMAP* 23, 33, and 41 GHz bands from electrons that *were* generated in SN shocks; this emission has a very disk-like morphology (and softer spectral index which is consistent with shock acceleration), while the haze has a more spherical morphology.⁵

Together, the haze spectrum and morphology imply either (1) a new class of objects distributed in the Galactic bulge

and largely missing from the disk; (2) significant acceleration from shocks several kpc off the plane toward the GC; or, perhaps most intriguingly, (3) a new electron component from a new physical mechanism. The claim that novel physics or astrophysics is required to explain the *WMAP* data is called the *haze hypothesis* to distinguish it from the two null hypotheses: (1) that the microwave haze is not synchrotron, but rather some combination of free-free and spinning dust and (2) that the haze is synchrotron, but the electron spectrum required is not unusual. In this work we do not address the origin of the electrons, but instead consider what the data from the *Fermi* Gamma-ray Space Telescope⁶ imply for their existence.

Electron cosmic rays at 10–100 GeV primarily lose energy in the diffuse interstellar medium (ISM) by producing synchrotron microwaves and inverse Compton (IC) scattered gammas. Synchrotron losses are proportional to magnetic field energy density, $U_B = B^2/8\pi$, while IC losses are proportional to the interstellar radiation field (ISRF) energy density, U_{ISRF} , in the Thomson limit, and less in the Klein–Nishina limit. Bremsstrahlung off the ambient gas also occurs, but is expected to be sub-dominant in the regions of interest. Therefore, the best test of the haze hypothesis is to search for IC gammas in the *Fermi* data, which was studied in the context of dark matter (DM) signals by Cholis et al. (2009a), Zhang et al. (2009), Borriello et al. (2009), Cirelli & Panci (2009), Regis & Ullio (2009), Belikov & Hooper (2010), and Meade et al. (2010).

Previous studies of high latitude gamma-ray emission have reported measurements of an excess of emission above the Galactic plane, most notably SAS-2 (Fichtel et al. 1975; Kniffen & Fichtel 1981), COS-B (Strong 1984), and EGRET (Smialkowski et al. 1997; Dixon et al. 1998). However, in each case the experiments either covered an insufficient energy range (SAS-2 and COS-B) or did not have sufficient sensitivity and angular resolution (EGRET) to permit a spatial correlation with the *WMAP* haze. *Fermi* overcomes both of these obstacles and allows us to search for the gamma-ray counterpart to the microwave haze for the first time.

⁵ Hereafter “spherical” is taken to mean “not disk-like”—if anything, the haze is non-spherical in the direction perpendicular to the disk (see Section 3.2).

⁶ See <http://fermi.gsfc.nasa.gov/ssc/data/>

Figure 1. Top and bottom left: full-sky *Fermi* γ -ray maps in various energy bins. The mask includes the 3-month *Fermi* point-source catalog as well as the LMC, SMC, Orion-Barnard’s Loop, and Cen A. Bottom right: exposure time map with our mask overlaid and stretched to 0%–100% of the peak exposure time. The variation in the exposure is a small modulation—even setting it to unity does not change our qualitative results.

The presence of an IC signal at the expected level can confirm that the microwave haze is indeed synchrotron, ruling out the first null hypothesis. From the spectrum of the IC, we can estimate the electron spectrum required to make the signal, addressing the second null hypothesis. For example, the presence of ~ 50 GeV IC photons requires electrons of $E > 50$ GeV, perhaps much greater. Furthermore, IC photons provide valuable information about the spatial distribution (disk versus bulge) of the source of these particles, which in turn can constrain hypotheses about their origin.⁷

In Section 2, we briefly review the *Fermi* data, describe our map-making procedure, and display full-sky maps at various energy ranges. In Section 3, these maps are analyzed with template correlation techniques and resultant residual maps and spectra are shown. Finally, Section 4 presents our interpretation of the signals and discusses potential sources of contamination. Estimates for the possible contamination from unresolved point sources are given in Appendix A. Appendix B details the creation and processing of the gamma-ray sky maps used in this analysis, and provides instructions for downloading them. Appendix C contains a discussion of Poisson likelihood analysis on smoothed maps.

2. DATA

The Large Area Telescope (LAT) on *Fermi* (see Gehrels & Michelson 1999 as well as the *Fermi* home page⁸) is a pair-conversion telescope consisting of 16 layers of tungsten on top of a calorimeter with a thickness of 7 radiation lengths. The entire instrument is wrapped in a scintillating anti-coincidence

detector to provide a particle veto. The spacecraft occupies a low Earth orbit with an inclination of $25^\circ 6'$. The field of view is so wide that the entire sky may be covered in two orbits by rocking the spacecraft north of zenith on one orbit and south of zenith on the other. Several times per month, *Fermi* interrupts this pattern to point the LAT at a gamma-ray burst (GRB), though this has little impact on the integrated exposure map. When the spacecraft passes through the South Atlantic Anomaly (SAA), cosmic ray (CR) contamination increases and significant data must be discarded, reducing the mean exposure at southern declinations. Beyond a zenith angle of 105° the data are significantly contaminated by atmospheric gammas. We excise such data, and select only events designated “Class 3” (diffuse class) by the LAT pipeline. The LAT collaboration plans to release a cleaner class of events in the future; however, at the time of this writing, the Class 3 events are the most likely to be real diffuse gamma-ray events.

The events are then binned into a full-sky map using HEALPix, a convenient iso-latitude equal-area full-sky pixelization widely used in the cosmic microwave background (CMB) community.⁹ Spherical harmonic smoothing is straightforward in this pixelization, and we smooth the maps to a Gaussian point-spread function (PSF), usually of 2° FWHM. The full-sky *Fermi* maps are displayed in Figure 1 along with an exposure map. See Appendix B for more details on map construction, smoothing, masking, and for instructions on how to download the maps.

⁷ The synchrotron haze depends on the Galactic magnetic field while the IC haze depends on the Galactic ISRF, and so the two morphologies should not be identical and could be quite different.

⁸ <http://fermi.gsfc.nasa.gov/>

⁹ HEALPix software and documentation can be found at <http://healpix.jpl.nasa.gov>, and the IDL routines used in this analysis are available as part of the IDLUTILS product at <http://sdss3data.lbl.gov/software/idlutils>.

Figure 2. GALPROP model illustrating the three primary gamma-ray emission mechanisms (see Section 3) and their relative amplitudes in the Galactic plane ($|\ell| \leq 30$, $|b| \leq 5$).

3. ANALYSIS

In this work, our goal is to test our general preconceptions about what gamma-ray signals should be present and identify any unexpected features in the *Fermi* data; we avoid detailed comparisons between the data and specific theoretical models for the Galactic gamma-ray emission. Our approach is to compute linear combinations of *Fermi* maps at several energies and perform template analyses with maps of the ISM, radio maps, etc. to see what emerges. This sort of open-minded analysis is flexible enough to find the unexpected.

An alternate approach would be to attempt to fit the data with a sophisticated physical model, in the context of some simulation code (e.g., GALPROP). Such a physical model can be quite detailed, including the three-dimensional distributions of gas and dust, the three-dimensional distribution of optical and FIR photons' density and direction, the three-dimensional magnetic field, and a three-dimensional model of p , e^- injections. By propagating these primary particles with GALPROP, the resulting π^0 , bremsstrahlung, and IC signals may be predicted and compared with *Fermi*. However, while a detailed physical model will certainly be crucial to a full understanding of the *Fermi* data (and such modeling is currently underway within the *Fermi* collaboration), this approach may lack the flexibility to identify new emission components that cannot be absorbed by modifying parameters in the model. On the other hand, such a model may also have too much freedom, so that meaningful patterns are absorbed into the fit and left unnoticed. In future work, we will take the signals revealed by our initial analysis and fold them back into a full physical model.

3.1. Diffuse Gamma Templates

There are three well-known mechanisms for generating gamma rays at the energies observed by *Fermi*. First, at low (~ 1 GeV) energies, gamma-ray emission is dominated by photons produced by the decay of π^0 particles generated in the collisions of cosmic ray protons (which have been accelerated by SNe) with gas and dust in the ISM. Second, relativistic electrons colliding with nuclei (mostly protons) in the ISM produce bremsstrahlung radiation. Finally, those same electrons interact with the ISRF and IC scatter CMB, infrared, and optical photons up to gamma-ray energies. A schematic of the relative importance of these emission mechanisms in the Galactic plane

($|\ell| \leq 30$, $|b| \leq 5$) generated by the GALPROP code, version 50p (Strong et al. 2000; Porter & Strong 2005; Strong et al. 2007), is shown in Figure 2.

Since π^0 gammas and bremsstrahlung are produced by interactions of protons and electrons (respectively) with the ISM, these emission mechanisms should be morphologically correlated with other tracers of the ISM, such as the Schlegel, Finkbeiner, & Davis (SFD) dust map based on 100 μm far IR data (Schlegel et al. 1998). The π^0 gamma-ray intensity scales with the ISM volume density times the proton CR density, integrated along the line of sight. In the limit where the proton cosmic ray spectrum and density is spatially uniform, the ISM column density is a good tracer of π^0 emission. Likewise, for a uniform electron spectrum, it is a good tracer of bremsstrahlung. Because our analysis is limited to $|b| > 5^\circ$, much of the emission we see is within a few kpc, so the assumption of uniform CR density is more valid than it would be for the entire Galaxy, particularly for protons, which have much larger propagation lengths than electrons.

3.2. Residual Maps

Since our goal is to search for an IC emission component with a morphology which roughly matches the microwave haze (i.e., centered on the GC, *roughly* spherical, and about 20° – 40° in radius), we now attempt to remove the π^0 emission from the maps shown in Figure 1, using the same template fitting technique used in Finkbeiner (2004a) and Dobler & Finkbeiner (2008). We perform multiple types of template fits. Type 1 uses only the *Fermi* map itself at 1–2 GeV which roughly traces π^0 emission because the gamma-ray sky at those energies is dominated by π^0 gammas (with subdominant contributions from bremsstrahlung and IC). Type 2 uses only the SFD dust map which is a direct tracer of ISM density and so should roughly map where the π^0 gammas and much of the bremsstrahlung are produced—again, up to some uncertainty involving the line-of-sight distribution. In each case, a uniform background is included in the fit, making our results insensitive to zero-point offsets in the maps.

We model the *Fermi* map at energy E , $F(E)$, as a linear combination of template maps, $F_{\text{model}} = T c_T$, where F_{model} is a column vector of N_{pix} unmasked pixel values, T is the $N_{\text{template}} \times N_{\text{pix}}$ template matrix, and the correlation coefficients c_T are chosen to minimize the mean squared residual,

$$\langle (F - F_{\text{model}})^2 \rangle = \langle (F - T c_T)^2 \rangle, \quad (1)$$

averaging over pixels. The least-squares solution,

$$c_T(E) = (T^T T)^{-1} \times (T^T F(E)), \quad (2)$$

yields the template correlation coefficients at each energy. In this fit, we mask out the *Fermi* three-month point-source catalog as well as the LMC, SMC, Orion-Barnard's Loop, and Cen A (see Appendix B for details). We also mask all pixels with Galactic latitude $|b| < 5^\circ$. Cross-correlations are done over unmasked pixels and for several different ranges in ℓ : eight longitudinal slices that have $|\ell| \geq 75^\circ$ (to avoid the GC) and width 30° (regions 1–7) as well as one region with $75^\circ < \ell < 285^\circ$ (region 8). Note that region 8 is the union of regions 1–7; regions 1–7 are fit individually to show the variation in the fit spectrum, and region 8 is fit to obtain the mean outer Galaxy signal. With these correlation coefficients, we define the residual map to be

$$R_T(E) = F(E) - c_T(E) \times T. \quad (3)$$

Figure 3. Residual maps after subtracting the SFD dust template from *Fermi* maps at various energies. The mask is described in Appendix B. Cross-correlations are done over unmasked pixels and for $75 \leq \ell \leq 285$. Although the template removes much of the emission, there is a clear excess toward the GC. This excess also includes a disk-like component which is likely due to IC and bremsstrahlung from softer electrons (see Figure 5).

To the extent that the templates in T match the morphology of the π^0 and bremsstrahlung gammas, the residual map will include only IC emission.

Figure 3 shows the resultant residual maps using the SFD map as a morphological tracer of π^0 emission for the region 8 fits. The most striking feature of the difference maps is the extended emission centered around the GC and extending roughly 40° in b . The morphological correlation between the *WMAP* synchrotron and the $R_{\text{SFD}}(5\text{--}10\text{ GeV})$ is striking as is shown in Figure 4. Here the 41 GHz synchrotron (haze plus Haslam-correlated emission, see Haslam et al. 1982; Dobler & Finkbeiner 2008) is shown side by side with the 5–10 GeV *Fermi* residual map with the mask used in the Dobler & Finkbeiner (2008) microwave analysis overlaid for visualization.

Figure 5 shows residual maps using the 1–2 GeV *Fermi* maps as a template for the π^0 emission. This sort of “internal linear combination” has the advantage that π^0 emission cancels out as long as the *shape* of the proton CR spectrum is the same everywhere—it does not rely on the proton CR density to be uniform. The residual maps look largely similar to the case with the SFD template regressed out, but there are some notable differences. In particular, using this 1–2 GeV template, the IC haze has a slightly “taller” appearance. This seems to be due to a disk-like component that is present in the 1–2 GeV maps but not in the dust map. This is probably because the 1–2 GeV *Fermi* map is not entirely π^0 emission, but also contains bremsstrahlung and IC components generated by electrons accelerated by SNe in the disk. The result is that when this lower energy map is regressed out

(i.e., cross-correlated and subtracted) from higher energy maps, this emission component is subtracted along with the dust-correlated emission. Conversely, the SFD map contains only emission from dust grains and not relativistic electrons; while the bremsstrahlung from those electrons largely traces the gas distribution, the IC does not, and so is not regressed out when using the SFD template.

Despite the significant shot noise in the *Fermi* map, Figure 4 shows that there is a clear morphological correlation between the microwave haze and the gamma-ray haze. Of course, we do not expect the morphologies to agree perfectly since the microwave haze is generated by interactions of electrons with the magnetic field while this *IC haze* is due to interactions of the electrons with the ISRF. Nevertheless, this is evidence that the microwave haze seen in the *WMAP* data is indeed *synchrotron* and is not some other component such as free–free emission or spinning dust.

3.3. Four-component Template Fits

Because the solution to Equation (2) minimizes the variance of the (zero mean) residual map $R(E)$, the presence of the IC haze affects the coefficients $c_T(E)$, since there is non-zero spatial correlation between the templates and the IC haze. To relax the stress on the fit, we expand our template analysis with a third type. Type 3 uses four templates in T : the SFD map to trace π^0 and bremsstrahlung emission, the 408 MHz Haslam et al. (1982) map which is dominated by radio wavelength synchrotron and thus roughly traces soft spectrum electrons

Figure 4. Top: the *WMAP* *Q*-band (41 GHz) total synchrotron (left) and haze (right). Bottom: the residual map obtained with the Type 1 template fit using the SFD map (left) and the 1–2 GeV *Fermi* map (right) as a template for π^0 decay emission. These residual maps are the same as the right column second row panel of Figures 3 and 5, respectively, though with a different stretch. Both maps are shown with the mask used in the Dobler & Finkbeiner (2008) microwave analysis for comparison. The upper panel represents cosmic ray electrons interacting with the Galactic magnetic field to produce synchrotron while the lower panel represents cosmic ray electrons interacting with the ISRF to produce IC emission, and colliding with the ISM to produce bremsstrahlung. The morphological similarities between the microwave haze and gamma-ray haze (right) are striking, indicating a correspondence between the two electron populations.

Figure 5. Same as Figure 3 but using the *Fermi* 1–2 GeV map for cross-correlations instead. Unlike the SFD dust map which should trace π^0 (and subdominant bremsstrahlung) emission only, the low-energy *Fermi* map includes the soft IC associated with lower energy electrons. In fact, comparing the residuals in this figure with those in Figure 3, it is clear that the disk-like component has been subtracted leaving only the IC haze. Furthermore, the IC haze is more prominent in the high-energy maps indicating a harder spectrum than π^0 emission which is the dominant emission mechanism at ~ 1 GeV energies.

Figure 6. Templates and fit solutions used in the Type 3 (see Section 3.2) template fits. Upper left: the SFD dust map, upper right: the Haslam 408 MHz map, middle left: the bivariate Gaussian haze template, middle right: the *Fermi* map at 10–20 GeV (same as the first column, third row of Figure 5 but with a different stretch to show the detailed morphological structure), lower left: the best-fit template solution for the observed emission, lower right: the residual map. Note the very small residuals indicating that the template fit is a remarkably good representation of the data over large areas of the sky.

which produce soft IC and bremsstrahlung, and a bivariate Gaussian of width $\sigma_\ell = 15^\circ$ and $\sigma_b = 25^\circ$. We note that this template is chosen to roughly match the morphology in Figure 5 and has no other physical motivation. We also use a uniform template to fit out the isotropic background signal in the maps, again, making our results insensitive to zero points. Lastly, for this fit we use *all* values in ℓ (region 9).

Note that since the bremsstrahlung originates from interactions of the electrons with the ISM, its spatial distribution depends on both the gas density and the cosmic ray electron density; consequently, some contribution from bremsstrahlung will be present in both the SFD-correlated and Haslam-correlated emission.

The previous fits were done with uniform weighting and assuming Gaussian errors, minimizing χ^2 . For the Type 3 fit we do a more careful regression, maximizing the Poisson likelihood of the four-template model in order to weight the *Fermi* data properly. In other words, for each set of model parameters, we compute the log likelihood

$$\ln \mathcal{L} = \sum_i k_i \ln \mu_i - \mu_i - \ln(k_i!), \quad (4)$$

where μ_i is the synthetic map (i.e., linear combination of templates) at pixel i , and k is the map of observed counts. Note that the last term does not depend on the model parameters. It may appear strange at first to compute a Poisson likelihood on smoothed maps, however, the smoothing is necessary to match

PSFs at different energies and with various templates (some of which have lower resolution than *Fermi* in the energy range of interest). The smoothing itself does not pose any problems for relative likelihoods, as we show in Appendix C.

However, we must keep in mind that the uncertainties derived in this way are the formal errors corresponding to $\Delta \ln \mathcal{L} = 1/2$, which would be 1σ in the case of Gaussian errors. The error bars plotted are simply the square root of the diagonals of the covariance matrix. This estimate of the uncertainty should be accurate at high energies, where photon Poisson noise dominates. At low energies, although the formal errors properly reflect the uncertainty in the fit coefficients for this simple model, the true uncertainty is dominated by the fact that the four-template model is not an adequate representation of the data.

Figure 6 shows the sky maps and best-fit solution including the residual map at 10–20 GeV while Figure 7 shows residual maps at other energies. It is clear from these residuals that the template fitting produces a relatively good approximation of the gamma-ray data over large areas of sky. Furthermore, Figure 7 shows that *not* including the bivariate Gaussian template for the IC haze yields a statistically significant residual toward the center indicating that a model including an IC haze is a better match to the data than one without. The prominent North Polar spur feature in the Haslam map, which is thought to originate from synchrotron emission from electrons in Loop I (Large et al. 1962), is oversubtracted in each case, because the North

Figure 7. Residual maps using the Type 3 template fit. The right column is the same as the lower right map in Figure 6 but for maps at different energy bands. The left column performs the same fit *without* including a bivariate Gaussian template for the IC haze. It is clear that not including the haze template results in a significant residual toward the GC in each energy band, but particularly at high energies. Including the haze template improves $\ln \mathcal{L}$ by 504, 215, 78, and 54, respectively, for the four energy bins shown.

Polar spur is brighter in the Haslam map than in the gamma-ray maps (i.e., the ratio of synchrotron microwaves to IC gamma rays in the North Polar spur is larger than in the rest of the Haslam map).¹⁰ This may be due to different ISRF, B -field, and ISM density values in Loop I relative to the inner Galaxy (since Loop I is thought to be quite nearby), or may be due to a softer-than-usual electron spectrum in Loop I, since the electrons producing the synchrotron measured in the Haslam map are much lower energy than those producing IC gamma rays at energies measured by *Fermi*.

Figure 8 shows these same residual maps, but with a smoothing of 10° which is on the order of the scale of the haze emission. With this large smoothing, smaller scale variations (due to individual photons at high energies) are smoothed over and the residual maps clearly show that the haze is a robust feature at all energies.

In Figure 9, we show the results of a four-template fit using the 1–2 GeV map instead of SFD to trace the π^0 emission. This figure shows that the haze is *not* due to the SFD template being an imperfect tracer of π^0 emission. If the proton cosmic

ray density is higher toward the GC, then SFD may systematically underestimate the π^0 emission there and perhaps the gamma-ray haze is the result. However, Figure 9 shows that this is clearly not the case. The 1–2 GeV template includes the effects of proton cosmic ray density variations (as well as line-of-sight gas density effects) and the haze remains as a robust residual. That is, the haze is *not* due to imperfect templates as suggested by Linden & Profumo (2010).

In Figure 10, we show the haze amplitude (residual maps from the Type 3 fit plus the correlation-coefficient-weighted bivariate Gaussian template) as a function of Galactic latitude and for the longitudinal bin: $-15^\circ < \ell < 15^\circ$. Although the data are noisy, the figure shows that the *Fermi* haze dies off by roughly $b \sim 50^\circ$ in all energy bands. For comparison, we show the same plot for just the residual map. The figure also shows the amplitude as a function of longitude for the latitudinal bin $-20^\circ < b < -10^\circ$. The more rapid falloff of the haze emission with ℓ compared to b indicates a haze morphology elongated in the b direction.

There are two important features to note about Figures 3 through 7. First, the IC haze has a spectrum which is *harder* than the other IC in the Galaxy. This is evidenced by the fact that using the 1–2 GeV *Fermi* map as a template removes the IC emission from disk electrons but the haze IC fades more

¹⁰ See Casandjian et al. (2009) for a discussion of gamma rays from Loop I seen by *Fermi*.

Figure 8. Same as Figure 7 but for an FWHM smoothing of 10° . Smoothing the residuals at this scale demonstrates that the haze is a robust feature and, in particular, is not the consequence of single photon fluctuations at high energies.

slowly with energy. Second, the IC haze is morphologically distinct from either the π^0 emission or IC and bremsstrahlung from the disk electrons. These two facts taken together strongly suggested that the electrons responsible for the microwave and gamma-ray haze are from a *separate* component with a *harder spectrum* than SN shock-accelerated electrons.

3.4. Fermi Galactic Diffuse Emission Model

So far, we have only done template fits of maps of data to other maps of data. Using the SFD dust map as a tracer of π^0 emission is equivalent to assuming that the proton CR spectrum and density are spatially uniform, and the dust/gas ratio is constant. Using the Haslam map to estimate IC emission is equivalent to assuming that the B -field and ISRF have similar spatial variation, and neglecting the anisotropies in both IC and synchrotron emission. Consequently, there have been concerns (Linden & Profumo 2010) that this template-based analysis might introduce spurious large-scale residuals. To address these concerns, we investigate whether the *Fermi* diffuse model contains a structure similar to the haze emission.

The *Fermi* team provides a model of diffuse gamma-ray emission consisting of maps sampled at 30 energy bins from 50 MeV to 100 GeV.¹¹ These maps are based on template fits to

the gamma-ray data and also include an IC component generated by the GALPROP cosmic ray propagation code.

The difference between the model and data is shown in Figure 11. We have interpolated the model to the energy ranges of interest and performed the simple one-template fit to the data (analogous to our Type 1 and 2 fits described above). This fit allows for any error in the normalization of the diffuse model; the fit coefficients are within a few percent of unity in every case.

At each energy range, the haze is clearly visible in the residual maps. GALPROP uses the standard inhomogeneous ISRF model (Porter & Strong 2005), making it unlikely that the observed residual is due to the expected spatial variation of the ISRF. Furthermore, the diffuse model includes a disk-like injection of primary electrons and estimates for the line-of-sight density variations of the ISM. The fact that the haze residual remains suggests that, given the propagation model included in GALPROP, disk-like sources will not produce the observed IC emission, nor will expected variations in the gas density and proton CR density along the line of sight.

3.5. Spectra

While the morphology of the gamma-ray haze is indicative of IC emission from the microwave haze electrons, we now attempt to estimate the spectrum of this emission. This is difficult

¹¹ The background models can be downloaded from <http://fermi.gsfc.nasa.gov/ssc/data/access/lat>.

Figure 9. Same as Figure 7 but using the 1–2 GeV map instead of the SFD dust map to trace π^0 emission. The clear haze residual seen, particularly at high energies, indicates that the gamma-ray haze is not due to shortcomings of the SFD template resulting from variations in the proton cosmic ray density. These variations, as well as line-of-sight gas density effects, are automatically included in the 1–2 GeV template.

for several reasons. First, π^0 emission is dominant (or nearly so) at most energies in *Fermi*’s energy range. Thus, in a given region, we must estimate the spectral shape and amplitude of the π^0 emission in order to subtract it from the total. Second, the total number of photons measured by *Fermi* decreases rapidly with increasing energy. For example, in the inner Galaxy ($|\ell| \leq 30$, $|b| \leq 5$), there are only ~ 4300 photons between 10 and 100 GeV, resulting in substantial Poisson errors. Lastly, there is also significant isotropic background due to extragalactic emission and charged particle contamination, including heavy nuclei at high energies. We will show that the isotropic backgrounds are manageable, and present two types of spectra: template coefficients $c_T(E)$ and total fluxes dN/dE .

3.6. Isotropic Background

The *Fermi* data contain gamma rays from an unresolved extragalactic signal with $dN/dE \sim E^{-2.45}$ (Ackermann 2010) as well as particle contamination. We make no attempt to separate the extragalactic gamma-ray signal from the particle contamination. Instead, we measure the spectrum of this (nearly) isotropic background in eight regions at high latitude and test the assumption of isotropy. Specifically, we take combinations of longitude ranges of $-180^\circ < \ell < -90^\circ$, $-90^\circ < \ell < 0^\circ$,

$0^\circ < \ell < 90^\circ$, and $90^\circ < \ell < 180^\circ$ together with latitude ranges of $-90^\circ < b < -60^\circ$ and $60^\circ < b < 90^\circ$ for the $4 \times 2 = 8$ regions. The regions are chosen to be far from the plane to avoid contamination from Galactic emissions. We use the point-source masked maps described above, binned in 12 energy bins from 0.3 to 300 GeV. The energy bins, mean background, etc. are given in Table 1 and the eight spectra, along with the inverse-variance weighted mean, are plotted in Figure 12. We note that the error bars are the standard deviation of the eight regions, *not* the standard deviation of the mean. To be conservative, we use this standard deviation as the uncertainty in the background in the remaining stages of the analysis. Future work by the *Fermi* team to understand both the particle contamination and the gamma-ray sky will likely reduce this error in the background substantially.

3.7. Template-correlated Spectra

Figure 13 shows $c_T(E) \times \langle T \rangle$ for the two templates and regions 1–7 used in the Type 1 and Type 2 fits along with the model π^0 spectrum from GALPROP. It is clear from the figure that the cross-correlation technique produces π^0 spectra that are remarkably similar to the model spectrum at low energies, while at high energies the cross-correlation spectrum is slightly

Figure 10. Top left: the haze amplitude (Type 3 residual map plus haze template) as a function of Galactic latitude for four different energy ranges. The data are binned in steps of 0.055 in $\sin b$, i.e. roughly 4° bins near the plane. The plot shows that haze amplitude has roughly the same falloff with b ($\sim 40^\circ$) in all energy ranges. Top right: The same but for the Type 3 residual map. Bottom left and right: the same as the top left and right, but as a function of ℓ for a fixed range of b . The haze profile clearly falls off more quickly with ℓ than with b indicating a profile elongated in the b direction.

Table 1
Isotropic Background

E Range (GeV)	Energy (GeV)	Background ($\times 10^{-6} \text{ GeV cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$)	χ^2_v
0.3–0.5	0.4	2.235 ± 0.187	86.467
0.5–0.9	0.7	1.901 ± 0.190	76.812
0.9–1.7	1.3	1.657 ± 0.181	51.885
1.7–3.0	2.2	1.340 ± 0.157	27.169
3.0–5.3	4.0	1.093 ± 0.119	10.981
5.3–9.5	7.1	0.886 ± 0.069	2.720
9.5–16.9	12.7	0.687 ± 0.098	4.190
16.9–30.0	22.5	0.632 ± 0.068	1.357
30.0–53.3	40.0	0.695 ± 0.078	0.936
53.3–94.9	71.1	0.624 ± 0.125	1.183
94.9–168.7	126.5	0.704 ± 0.199	1.594
168.7–300.0	225.0	0.420 ± 0.205	1.949

Notes. The background is tabulated in each energy bin for eight polar ($|b| > 60^\circ$) regions after point-source masking. The uncertainty is the rms of the eight regions, *not* the uncertainty in the mean, which is $\sqrt{7}$ times smaller. For each region we also compute the Poisson uncertainty, σ_p , which for low E is much smaller. To test the significance of background variation, we compute $\chi^2_v = \langle \sigma^2 / \sigma_p^2 \rangle$, averaging over the eight regions. There is no strong indication of anisotropy at high latitude for $E > 5 \text{ GeV}$.

Table 2
Region Definitions

Region	ℓ Range	b Range
Inner galaxy	$ \ell < 30^\circ$	$ b < 5^\circ$
Haze	$ \ell < 15^\circ$	$-30^\circ < b < -10^\circ$
Four corners	$5^\circ < \ell < 10^\circ$	$5^\circ < b < 10^\circ$

Notes. Regions for which the total gamma-ray intensity is evaluated (see also Figure 15).

background events such as heavy nuclei, or uncertainties in the π^0 emission model. Of these, the first is most likely since the cross-correlation between the templates and a nearly isotropic background is likely small and since the spectrum of π^0 gammas is quite well known.

Template-correlated spectra for the Type 3 template fit are shown in Figure 14. Here the correlation coefficients are weighted by the mean of each template in the “haze” region (see Table 2). As shown in the figure, the spectra for the SFD and Haslam maps reasonably match the model expectations in that region.¹² That is, the SFD-correlated emission roughly follows the model π^0 spectrum while the Haslam-correlated spectrum resembles a combination of IC and bremsstrahlung

higher than the model spectrum. This could be due to a number of reasons such as non-zero spatial correlation between the templates and the harder spectrum haze IC, contamination from

¹² The GALPROP model here was tuned to match locally measured protons and anti-protons as well as locally measured electrons at ~ 20 – 30 GeV .

Figure 11. Same as Figure 3 but using the official *Fermi* team Galactic diffuse model. The strong haze residual in the right-hand panels shows that the haze is *not* included in the model. Since the IC emission in the model was obtained with GALPROP, this indicates that variation in the ISRF with position in the Galaxy cannot account for the haze emission. In addition, the haze morphology is not reproduced by disk-like injection of electrons, nor by the cosmic ray propagation model employed by GALPROP.

emission. However, the *haze*-correlated emission is clearly significantly harder than either of these components. This fact coupled with the distinct spatial morphology of the haze indicates that the IC haze is generated by a *separate* electron component.

3.8. Total Intensity Spectra

While the template-correlated spectra and residual maps are useful for identifying separate components, for the purposes of comparing the map intensities to a model for the physical mechanisms, we now generate total intensity spectra in several regions of interest. We define three key regions: a “haze region”

south of the GC,¹³ the “four corners” region from Cholis et al. (2009b), and a Galactic plane region used by Porter (2009).

The spectrum of each region is shown (Figure 15), with the background spectrum from Figure 12 subtracted from the other spectra to remove any isotropic component. The inner Galaxy region clearly shows the low-energy behavior characteristic of

¹³ The region south of the GC is of greatest interest for studying the microwave and gamma-ray haze. This is because, north of the GC, the microwave maps from *WMAP* include bright free-free emission from ζ Oph, and spinning dust emission from ρ Oph, both of which are bright enough to leave substantial residuals, even though they are relatively well subtracted by the template fitting described in Dobler & Finkbeiner (2008). South of the GC there are some small dust and gas features that should provide some signal in the gamma-ray map, but the situation is much simpler.

Figure 12. Background spectrum, including gamma rays and particle backgrounds. The background is tabulated in each energy bin for 8 high-latitude polar ($|b| > 60^\circ$) regions after point-source removal (dotted gray lines). The mean of the eight regions is shown with a solid black line and the uncertainty is the rms of the eight regions, *not* the uncertainty in the mean, which is $\sqrt{7}$ times smaller. For reference, a $dN/dE \propto E^{-2.45}$ model for the unresolved background is shown with a dot-dashed line.

π^0 emission (though it is important to note that there may be significant emission from *unresolved* point sources such as pulsars in this region which have a similar spectral shape; see Appendix A and Abdo et al. 2009b). At high energies however, there is a significant excess which is not expected for π^0 emission (cf. Figure 2). The spectrum of this excess is similar to the spectrum derived for the haze template from Figure 14 and is consistent with an IC signal from a hard electron population with energies >10 GeV. The evidence for this excess is more pronounced in the haze and four corners region. Lastly, we take the template estimate of π^0 emission in the haze region from the Type 3 fits and subtract it from the total emission. This reveals two clear features: a bump centered on roughly 1–2 GeV that is likely due to either bremsstrahlung from a low-energy electron component or emission from unresolved pulsars, and a hard tail above ~ 10 GeV that is the IC signal from the haze electrons.

Figure 13. Left: π^0 spectrum derived using the cross-correlation technique defined in Section 3.2 and with the *Fermi* 1–2 GeV map as a π^0 template. The dotted gray lines are the cross-correlation spectra in several different regions of the sky and the solid black line is the mean of those spectra. Error bars on the cross-correlation spectrum are defined as the variance in the values for the different regions. The dot-dashed line is the π^0 spectrum output from a GALPROP model (shown here with arbitrary normalization). Right: same but using the SFD dust map as a template.

3.9. Comparison to the Microwave Haze

Both the morphology and the relatively hard spectrum of the gamma-ray haze motivate a common physical origin with the *WMAP* haze. We now provide a simple estimate of the microwave and gamma-ray signals from a population of hard electrons in the inner Galaxy, to demonstrate that the magnitudes and spectral indices of the two signals are consistent for reasonable parameter values.

We consider a steady-state electron spectrum described by a power law, $dN/dE \propto E^{-\alpha}$, with a high-energy cutoff at 1 TeV (here the cutoff is implemented as a step function, not an exponential falloff; of course this is only an approximation to the true spectrum). This choice is motivated by the local measurement of the cosmic ray electron spectrum by *Fermi* (Abdo et al. 2010). We consider a region ~ 2 kpc above the GC, as an example (and since both hazes are reasonably well measured there), and employ the model for the ISRF used in GALPROP (Porter & Strong 2005) at 2 kpc above the center. We normalize the synchrotron to the approximate value measured by *WMAP* in the 23 GHz *K* band (Hooper et al. 2007), $\sim 15^\circ$ below the Galactic plane, and compute the corresponding synchrotron and IC spectra. The *WMAP* haze was estimated to have a spectrum $I_\nu \propto \nu^{-\beta}$, $\beta = 0.39\text{--}0.67$ (Dobler & Finkbeiner 2008), corresponding approximately to an electron spectral index of $\alpha \approx 1.8\text{--}2.4$; Figure 16 shows our results for a magnetic field of $10 \mu\text{G}$, and electron spectral indices $\alpha = 2\text{--}3$. This field strength is appropriate for an exponential model for the Galactic *B*-field intensity, $|B| = |B_0|e^{-z/z_s}$, with $B_0 \approx 20\text{--}30 \mu\text{G}$ (which seems reasonable; see, e.g., Ferriere 2009 and references therein) and scale height $z_s \approx 2$ kpc (the default value in GALPROP). We find good agreement in the case of $\alpha \approx 2\text{--}2.5$, consistent with the spectrum of the *WMAP* haze.

3.10. Comments on Haze Morphology

Although a detailed analysis of the possible sources of the *Fermi* haze are beyond the scope of this paper, a few simple comments are in order. First, the profile is not well described by a disk source. While quantifying this is a subtle task, the success of the template makes this point clear. There are many possibilities to explain the oblong shape, should that persist in future data. For instance, active galactic nucleus (AGN) jets and

Figure 14. Correlation coefficients for the templates used in the Type 3 template fit (see Section 3.2). Upper left: SFD-correlated spectrum which roughly traces π^0 emission. Upper right: Haslam-correlated emission which traces the soft IC and bremsstrahlung component. Lower left: haze template-correlated emission. This component has a notably harder spectrum than both the SFD- and Haslam-correlated spectra or their model shapes (dashed lines, cf. Figure 2), indicating a separate component. The dashed lines in the upper two panels are a GALPROP estimate of π^0 decay (left) and IC emission (right). The GALPROP estimate of bremsstrahlung emission is shown in both panels. Lower right: the uniform template-correlated spectrum which traces the isotropic background. Here the dashed line is the result from Figure 12. This high latitude estimate is higher than the uniform template estimate likely because the π^0 emission is non-zero at high latitudes and leaks into our measured background. This is less of an issue for the uniform template which uses the morphological (i.e., uniform) information.

triaxial DM profiles could both produce signals of this shape. Even approximately spherical injection could yield such a signal, should the diffusion be considerably anisotropic. That said, while one might have attempted to invoke, e.g., rapid and significant longitudinal variation of the magnetic field to explain the microwave haze with a disk-like electron injection profile, such approaches are no longer tenable (and moreover already had significant tension with an understanding of the Haslam synchrotron maps). The presence of this feature in both gamma rays and microwaves demonstrates that *the electrons themselves* do not follow a disk-like injection profile. Furthermore, we emphasize that the angular scale of the haze is very large, roughly 25° – 40° in b . This region of the sky is far from the Galactic disk or the very GC where complications from the central black hole, other point sources, or disk-like pulsars could complicate the analysis.

4. DISCUSSION/CONCLUSIONS

We have presented full-sky maps generated from photon events in the first year data release of the *Fermi* Gamma-ray Space Telescope (see Appendix B for data processing details).

Using a template fitting technique, we have approximated both the spectrum and morphology of the well-known gamma-ray emission components at *Fermi* energies. The SFD dust map was used to trace the π^0 decay gammas generated by collisions of cosmic ray protons with the ISM, while the Haslam 408 MHz map was used to trace IC scattered photons from interactions of SN shock-accelerated electrons (~ 1 – 10 GeV) with the ISRF. Bremsstrahlung radiation, generated by interactions of these electrons with the ISM, should be approximately traced by some combination of these two maps. Although our template fitting technique is subject to significant uncertainties due to uncertain line-of-sight gas and CR distributions, a robust positive residual has been identified.

This excess diffuse emission is centered on the GC, and can be parameterized by a simple two-dimensional Gaussian template ($\sigma_\ell = 15^\circ$, $\sigma_b = 25^\circ$). The template-correlated spectrum of this emission is significantly *harder* than either π^0 emission or IC from softer electrons, whose fitted spectra agree well with models. This harder spectrum coupled with the distinct spatial morphology of the gamma-ray and microwave haze are evidence that these electrons originate from a *separate component* than the softer SN shock-accelerated electrons.

Figure 15. Total intensity spectra in three regions of interest (solid, dashed, dot-dashed; see Table 2). All three show the characteristic π^0 spectrum at low energies as well as an excess at higher energies with roughly the same spectrum as the haze spectrum derived from the Type 3 template fits. Furthermore, subtracting the template π^0 spectrum in the haze region (triple-dot-dashed) yields two features, one centered on ~ 2 GeV and one centered on ~ 50 GeV. The lower energy feature is likely either bremsstrahlung from a low-energy electron population (like those accelerated by SN shocks) or emission from unresolved pulsars. The higher energy tail is the IC signal from the haze electrons.

The gamma-ray excess is almost certainly the IC counterpart of the microwave haze excess described by Finkbeiner (2004a) and Dobler & Finkbeiner (2008). Although it is still possible that a significant fraction is prompt photons from weakly interacting, massive particle (WIMP) annihilations (e.g., the 200 GeV wino advocated by Grajek et al. 2009) such explanations are difficult to reconcile with the spatial similarity to the *WMAP* haze (see Figure 4). The simplest hypothesis is that the signal is mainly IC from the same electrons that produce the *WMAP* haze synchrotron.

This addresses the stubborn question about the origin of the *WMAP* haze. Until recently, it has been argued that the *WMAP* haze had alternative explanations, such as free-free emission from hot gas or spinning dipole emission from rapidly rotating dust grains. However, the existence of this IC signal proves that the microwave haze is indeed synchrotron emission from a hard electron spectrum.

Fermi-LAT photon data are contaminated by particle events, especially at high energies. We have taken care to account for the isotropic background resulting from extragalactic sources, cosmic ray contamination, and heavy nuclei contamination and found that this background, though significant, is below the observed IC excess even up to 100 GeV. Particle contamination is extremely unlikely to mimic the observed signal.

The LAT collaboration continues to refine the cuts used to define “diffuse class” events, and plans to release a cleaner class of events in coming months. This, along with a new public version of GALPROP, including updated ISRF models, will allow a more sophisticated analysis than that presented in this paper. We eagerly await the release of these software and data products.

The spectrum and morphology of both the microwave and gamma-ray haze constrain explanations for the source of these electrons. There have been speculations that the microwave haze could indicate new physics, such as the decay or annihilation of

Figure 16. Estimated spectrum of IC gamma rays (upper panel) and synchrotron radiation (lower panel) originating from a hard electron spectrum along a line of sight 2 kpc above the GC (i.e., $b \approx 15^\circ$). The steady-state electron spectrum is taken to be a power law, $dN/dE \propto E^{-\alpha}$, with index $\alpha = 2$ (solid), 2.5 (dashed), and 3 (dotted). In all cases the spectrum has a cutoff at 1 TeV. The ISRF model is taken from GALPROP, and the magnetic field is set to be $10 \mu\text{G}$. The data points in the upper panel show the magnitude of the *Fermi* haze averaged over $|b| = 10$ – 18 , for $|l| < 15$, as a function of energy, taken from Figure 10. The highest two bins contain 3σ upper limits rather than data points with 1σ error bars, due to the large uncertainties in the haze amplitude at those energies. The data point in the lower panel shows the magnitude of the *WMAP* haze averaged over $b = -10$ to -18 , for $|l| < 10$, in the 23 GHz *K* band (the overall normalization is chosen to fit this value), and the gray area indicates the range of synchrotron spectral indices allowed for the *WMAP* haze by Dobler & Finkbeiner (2008).

(A color version of this figure is available in the online journal.)

DM, or new astrophysics, such as a GRB explosion, an AGN jet, or a spheroidal population of pulsars emitting hard electrons. We do not speculate in this paper on the origin of the haze electrons, other than to make the general observation that the roughly spherical morphology of the haze makes it difficult to explain with any population of disk objects, such as pulsars. The search for new physics—or an improved understanding of conventional astrophysics—will be the topic of future work.

We acknowledge helpful conversations with Elliott Bloom, Jean-Marc Casandjian, Carlos Frenk, Isabel Grenier, Igor Moskalenko, Simona Murgia, Troy Porter, Andy Strong, and Kent Wood. This work was partially supported by the Director, Office of Science, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. N.W. is supported by NSF CAREER grant PHY-0449818, and I.C. and N.W. are supported by DOE OJI grant number DE-FG02-06ER41417. I.C. is supported by the Mark Leslie Graduate Assistantship. D.F. and T.S. are partially supported by NASA grant NNX10AD85G. T.S. is

partially supported by a Sir Keith Murdoch Fellowship from the American Australian Association. This research made use of the NASA Astrophysics Data System (ADS) and the IDL Astronomy User's Library at Goddard.¹⁴

APPENDIX A

UNRESOLVED POINT SOURCES

In this section, we explore the possibility that the diffuse excess discussed in this work (the *Fermi* haze) could originate from a large number of unresolved point sources. The limit where the contribution from unresolved point sources is dominated by emission from many very faint ($\ll 1$ count) sources cannot be distinguished from purely diffuse emission; however, if a smaller number of brighter (~ 1 count) point sources dominate, photon events coincide within the PSF at a greater rate than would be the case for purely diffuse emission.

We model the luminosity function of unresolved point sources by a simple power law for the expected number of counts, S : $dN/dS \propto S^{-\alpha}$, with $S_{\min} < S < S_{\max}$. Smaller values of α , and larger values of $S_{\max}$ and $S_{\min}$, correspond to a larger fraction of bright point sources and thus strengthen any upper bounds on the flux fraction from point sources. Examining known point-source populations (of AGN and X-ray binaries) generally yields spectral indices $\alpha \sim 1.5$ – 2.2 , although at high luminosities steeper power laws with $\alpha \sim 2.5$ have been observed in flat spectrum radio quasars, and at low luminosities very shallow power laws with $\alpha \sim 1.2$ have been measured in low-mass X-ray binary (LMXB) populations (Grimm et al. 2003; Gilfanov 2004; Kim & Fabbiano 2004; Voss & Gilfanov 2006; Padovani et al. 2007; Abdo et al. 2009a). We take $\alpha = 2.2$ and $S_{\max} = 10$, as pessimistic benchmark parameters to provide a robust upper bound.

This power law must break at some low-luminosity cutoff to avoid divergences if $\alpha > 2$; we will display the statistical limits obtained for $S_{\min} = 0.1$ and 1 count. It is possible that any remaining excess emission could originate from point sources with expected counts year⁻¹ below 0.1, but this would require the number of point sources to exceed the number of unexplained excess photons by greater than a factor of 10.

To bound the fractional flux originating from point sources, we apply a statistical test of isotropy described in detail by Slatyer & Finkbeiner (2010), designed for the case where the density of events is $\lesssim 1$ count per PSF, to the *Fermi*-LAT data in the haze region ($|l| < 15$, $-30 < b < -10$). In brief, we compute (1) the fraction of “isolated events” with no neighbors within some test radius r , and (2) the fraction of 10^6 randomly distributed circles of radius r which contain no events (“empty circles”). If r is chosen appropriately (i.e., $r \sim \text{PSF}$) then the ratio of these two quantities, $R \equiv \text{“isolated”}/\text{“empty”}$, is related to the fraction of counts arising from unresolved point sources. If the average number of events per circle of radius r , denoted λ , is greater than 1, then we redefine “isolated” events to be those with fewer than λ neighbors, and “empty circles” to be those containing fewer than λ events. Neighbors are found efficiently using the publicly available IDL routine `spherematch`, and we take the test radius r to be equal to the estimated 1σ value for the LAT PSF, corresponding to 39% flux containment¹⁵ (our

results are not very sensitive to this choice). We apply this test separately to the diffuse class data in each of the energy bins in Figure 1, removing events with a large zenith angle as described previously, and calibrate our results via Monte Carlo simulations. The smaller the PSF, the better the bounds (the PSF is an input to our Monte Carlo calibration); since the PSF size decreases with increasing energy, and is greater for back-converting events compared to front-converting events, for each energy bin we take the PSF for back-converting events at the minimum energy of the bin, in order to set robust limits.

At low energies we find values of R consistent with a significant point-source component (although the point-source fraction cannot be precisely determined without knowing the luminosity function). Above 10 GeV, however, the value of R is consistent with entirely diffuse emission. In the 50–100 GeV bin, low statistics nonetheless allow the haze-correlated fraction of the emission to be explained entirely by point sources with 0.1+ expected counts year⁻¹ (since in this bin the haze-correlated flux in this region corresponds to ~ 30 photons), but at ~ 10 – 50 GeV there is a clear excess over the 95% confidence limits on the unresolved point-source contribution, for $S_{\min} = 0.1, 1$ (Figure 17). There are 900 diffuse class events in this energy range and spatial region which pass the zenith angle cut, of which $\sim 1/3$ are attributed to the haze: in order for point sources providing an expected flux of less than 0.1 counts year⁻¹ to make up the difference, there would have to be 1000+ such sources emitting 10–50 GeV gamma rays in this 600 deg² region of the sky (in addition to similar source populations in the other regions where the haze-correlated emission is bright). Therefore, even for these conservative assumptions, it seems very unlikely that the hard spectral shape of the *Fermi* haze could be due to point-source contamination over the entire relevant energy range.

APPENDIX B

FERMI FULL-SKY MAP PROCESSING AND DATA RELEASE

In this section, we describe full-sky maps generated from photon events in the first year data release of the *Fermi Gamma-ray Space Telescope*. These maps are corrected for exposure, point-source masked, and are smoothed to a Gaussian PSF, usually of 2° FWHM. The lowest energy maps are smoothed to 3° or 4° . For convenience, we provide two sets of energy bins: 12 logarithmically spaced bins for use in spectral analyses (“specbin”) and 8 somewhat larger bins with better signal/noise for visual inspection (“imbin”).

Although the maps were generated from a public *Fermi* data release, the maps are not an official LAT data release, and the procedure used to make them has not been endorsed by the LAT collaboration.

B.1. Event Selection and Binning

The LAT is a pair-conversion telescope, in which incoming photons strike layers of tungsten and convert to e^+e^- pairs, which are then tracked (to determine direction) on the way to a calorimeter (to determine energy). The first year P6_V3_DIFFUSE data file contains records of 15,878,650 events, providing the time, energy, arrival direction (with respect to the spacecraft, and in celestial coordinates), and zenith angle for each event. Events are divided into three classes, using a number of cuts to separate photon events from particle background.

¹⁴ Available at <http://idlastro.gsfc.nasa.gov>

¹⁵ The energy-dependent PSF for the LAT is taken from http://www-glast.slac.stanford.edu/software/IS/glast_lat_performance.htm. The radius of 68% containment used by the *Fermi* Collaboration would in the case of a Gaussian PSF be $r_{68} = 1.51\sigma$.

Figure 17. Upper limits on the fraction of diffuse class events originating from unresolved point sources, in the *Fermi* haze region ($|l| < 15$, $-30 < b < -10$), at 95% confidence. The dark shaded region indicates bounds on sources with an expected contribution of 1+ counts year⁻¹, while the light shaded region includes sources with 0.1+ counts year⁻¹; both cases assume a luminosity function with spectral index $\alpha = 2.2$ and $S_{\max} = 10$. The dotted and dashed lines indicate the bounds with a steeper luminosity function $\alpha = 1.8$, for $S_{\min} = 1$ and $S_{\min} = 0.1$ respectively. Data points indicate the fraction of emission correlated with the “Fermi Haze,” obtained from the spectra summarized in Figure 14. The table shows the value of 1σ for the detector PSF used in the Monte Carlo calibration, as conservatively estimated for each energy bin.

“Class 3” rejects the largest fraction of background contamination, and these events are used to study diffuse emission. We also require the zenith angle be less than 105° to exclude most atmospheric gammas. By choosing to use these cuts, we study the events most likely to be real gamma rays, at the expense of a smaller effective area. These cuts discard roughly 3/4 of the signal, but vastly reduce the noise.

The Class 3 events are then binned in energy and into spatial pixels to produce counts maps. We use the hierarchical equal-area isolatitude pixelization (HEALPix), a convenient equal-area full-sky pixelization widely used in the CMB community.¹⁶

B.2. Exposure Maps

Because the exposure on the sky is non-uniform, we generate an exposure map using the `gtexpcube` tool developed by the LAT team.¹⁷ The exposure for each pixel is the LAT effective area at each θ (angle with respect to the LAT axis) summed over the lifetime of the LAT at that θ and has units of cm² s. The exposure map spatially modulates the signal $\pm 20\%$ from raw photon counts and is slightly energy dependent for photon energies $E > 1$ GeV. To avoid systematic errors in generating the exposure maps we use the same energy grid used for the spectral plots, which is given in Table 1. The default spatial binning of 1° in latitude and longitude is adequate, as neighboring bins have exposure differences of $< 0.5\%$. The LAT contains 12 thin layers of tungsten designated “front” and 4 thicker layers designated “back.” Photons may convert to an e^+e^- pair in any of the layers, and events are labeled “front” or “back” accordingly. The effective area as a function of energy is different for front and back events, so we use the “P6_V3_DIFFUSE: :FRONT” and

“P6_V3_DIFFUSE: :BACK” instrumental response functions, respectively. Maps of counts are divided by the exposure and pixel solid angle to produce intensity maps (counts s⁻¹ cm⁻² sr⁻¹). The “combined” maps are simply an exposure-weighted linear combination of the front and back intensity maps. The full-sky *Fermi* maps are displayed in Figure 1 along with the exposure map and mask.

B.3. Smoothing

In our analysis we wish to compare *Fermi* maps at different energies to each other, and to other templates. In order to match the PSFs of all these maps, we smooth each map by an appropriate kernel. On average, the front and back converting events have different PSFs (by a factor of ~ 2), so we smooth the intensity maps of each to a common PSF (usually 2°) to produce “front” and “back” smoothed maps (which can be averaged to obtain “combined” smoothed maps as above).

For the PSF, we use a simple fit to the radius of 68% flux containment (in degrees):

$$r_{68} = ((c_1 E^\beta)^\alpha + c_2^\alpha)^{1/\alpha}, \quad (\text{B1})$$

where E is in GeV, $\alpha = 1.2$, $\beta = -0.83$, and $(c_1, c_2) = (0.50, 0.04)$ and $(0.90, 0.09)$ for front- and back-converting events, respectively. This yields $r_{68} = c_1 E^\beta$ at low E and $r_{68} = c_2$ at high E . For simplicity, we assume that the PSF is a Gaussian so that the FWHM, $f = 2r_{50} \approx 1.56r_{68}$. The “raw” FWHM of a counts map is then taken to be Equation (B1) evaluated at $E_{\text{mean}} = \sqrt{E_0 E_1}$, where $E_0 < E < E_1$ for the events in the map. In order to smooth to the target PSF (usually $f_{\text{targ}} = 2^\circ$), a Gaussian smoothing kernel of size f_{kern} is used such that

$$f_{\text{kern}} = \sqrt{f_{\text{targ}}^2 - f_{\text{raw}}^2}. \quad (\text{B2})$$

For the lowest energy maps ($E < 1$ GeV), where f_{raw} is large, we take $f_{\text{targ}} = 3^\circ$ or 4° .

¹⁶ HEALPix software and documentation can be found at <http://healpix.jpl.nasa.gov>, and the IDL routines used in this analysis are available as part of the IDLUTILS product at <http://sdss3data.lbl.gov/software/idlutils>.

¹⁷ <http://fermi.gsfc.nasa.gov/ssc/data/analysis/documentation>

Figure 18. Left panel: 2° FWHM mock map fit to a 2° model. Right panel: a 1° mock map is smoothed to a 2° PSF and then fit to the 2° model. In each case, 1000 realizations are computed, and the histograms of χ are shown for the four parameters. In each case, the χ histogram has a mean near zero and rms near 1, indicating that the error analysis is correct, even in the case of a smoothed map. Biases introduced are small ($\ll 1\sigma$) for the case shown, and are relatively even less important at higher energy.

(A color version of this figure is available in the online journal.)

B.4. Point-source Mask

The point-source mask (Figure 1) contains the 3-month 10 σ point-source catalog, plus the LMC, SMC, Orion, and Cen A. The mask radius for point sources is taken to be the 95% containment radius for the lowest energy event in the energy bin. For unsmoothed maps, both the counts map and the exposure map are multiplied by the mask. For smoothed maps, counts and exposure are multiplied by the mask, then smoothed, and then divided. Pixels where the smoothed masked exposure drops below 25% of the mean unmasked exposure are replaced with zeros. In cases where the mask radius is much smaller than the smoothing radius, the mask is “smoothed away.” In other cases, masked pixels are visible in the smoothed maps. Because of the energy dependence of the mask radius, the effective mask changes with energy, so care must be taken in cross comparisons between energy bins, especially within a few degrees of the Galactic plane.

B.5. Data Release

All LAT maps described in this section are available in the HEALPix pixelization on the Web.¹⁸ Maps of front- and back-converting events are available, as well as the combined maps. All three are available smoothed and unsmoothed, and with and without point-source masking, for a total of 12 maps at each energy. Both the imbin and specbin energy binnings are available, for a total of 240 FITS files. Gray scale and color jpegs are also available to provide a quick overview. Software used to make the maps is available on request.

APPENDIX C

POISSON LIKELIHOOD ANALYSIS

Ideally the spatial templates (dust map, synchrotron map, etc.) going in to the Type 3 likelihood analysis (Section 3.3) would be smoothed to the PSF of the gamma-ray data, allowing an event-by-event likelihood to be calculated. However, at some energies, the LAT has a higher angular resolution than our templates. In this case, the LAT data and templates must be smoothed to a common PSF. In this section, we explain the Poisson likelihood analysis in more detail, and check that a Poisson likelihood evaluated on a smoothed map behaves as expected.

For the Type 3 fit we maximize the Poisson likelihood of the four-template model in order to weight the *Fermi* data properly. In other words, for each set of model parameters, we compute the log likelihood

$$\ln \mathcal{L} = \sum_i [k_i \ln \mu_i - \mu_i - \ln(k_i!)] , \quad (\text{C1})$$

where μ_i is the synthetic counts map (i.e., linear combination of templates times exposure and mask) at pixel i , and k is the map of observed counts. The μ_i term depends only on the model parameters, and is not affected by smoothing the data. The $\ln(k_i!)$ term does not depend on the model parameters, and so it cannot affect the best-fit model or uncertainties. Furthermore, there is no problem evaluating the likelihood for fractional k using the gamma function. So the $k_i \ln \mu_i$ is the only potentially problematic term.

In order to investigate the effects of smoothing, we generate each mock map by taking a linear combination of templates

¹⁸ <http://fermi.skymaps.info>

(including the uniform background), multiplying it by the exposure and mask to obtain a map of predicted counts, and Poisson sampling it in HEALPix pixels. This map is then passed to our parameter estimation code to obtain the best-fit values and uncertainties for the four parameters. Repeating this procedure for 1000 mock maps, we compute

$$\chi = \left(\frac{\text{fit} - \text{true}}{\sigma} \right) \quad (\text{C2})$$

for each parameter and plot the histograms (Figure 18). An unbiased fit corresponds to $\langle \chi \rangle = 0$ and a correct uncertainty estimate corresponds to $\text{stdev}(\chi) = 1$. The uncertainties appear to be correctly estimated and the bias is small. Note that this behavior is different from an $\exp(-\chi^2/2)$ likelihood on a smoothed map: in that case the correlated noise induced by the smoothing must be taken account of carefully.

The above analysis was done for a range of energy bins with similar results. The analysis was also completely redone with a χ^2 minimization instead (indeed, this provides the initial guess for the Poisson likelihood minimization) with very similar results. Because we feel the Poisson likelihood is more correct for the problem at hand, we use it for the “official” results.

REFERENCES

- Abdo, A. A., et al. 2009a, *ApJ*, **700**, 597
 Abdo, A. A., et al. 2009b, *ApJS*, **183**, 46
 Abdo, A. A., et al. 2010, *ApJS*, **187**, 460
 Ackermann, M. 2010, arXiv:1002.3603
 Belikov, A. V., & Hooper, D. 2010, *Phys. Rev. D*, **81**, 043505
 Blattnig, S. R., Swaminathan, S. R., Kruger, A. T., Ngom, M., & Norbury, J. W. 2000, *Phys. Rev. D*, **62**, 094030
 Borriello, E., Cuoco, A., & Miele, G. 2009, *ApJ*, **699**, L59
 Casandjian, J., & Grenier, I. for the Fermi Large Area Telescope Collaboration 2009, arXiv:0912.3478
 Cholis, I., Dobler, G., Finkbeiner, D. P., Goodenough, L., & Weiner, N. 2009a, *Phys. Rev. D*, **80**, 123518
 Cholis, I., et al. 2009b, arXiv:0907.3953
 Cirelli, M., & Panci, P. 2009, *Nucl. Phys. B*, **821**, 399
 Dixon, D. D., Hartmann, D. H., Kolaczyk, E. D., Samimi, J., Diehl, R., Kanbach, G., Mayer-Hasselwander, H., & Strong, A. W. 1998, *New Astron.*, **3**, 539
 Dobler, G., & Finkbeiner, D. P. 2008, *ApJ*, **680**, 1222
 Ferriere, K. 2009, *A&A*, **505**, 1183
 Fichtel, C. E., Hartman, R. C., Kniffen, D. A., Thompson, D. J., Ogelman, H., Ozel, M. E., Tumer, T., & Bignami, G. F. 1975, *ApJ*, **198**, 163
 Finkbeiner, D. P. 2004a, *ApJ*, **614**, 186
 Finkbeiner, D. P. 2004b, arXiv:astro-ph/9027
 Gehrels, N., & Michelson, P. 1999, *Astropart. Phys.*, **11**, 277
 Gilfanov, M. 2004, *MNRAS*, **349**, 146
 Grajek, P., Kane, G., Phalen, D., Pierce, A., & Watson, S. 2009, *Phys. Rev. D*, **79**, 043506
 Grimm, H. J., Gilfanov, M., & Sunyaev, R. 2003, *MNRAS*, **339**, 793
 Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, *A&AS*, **47**, 1
 Hooper, D., Finkbeiner, D. P., & Dobler, G. 2007, *Phys. Rev. D*, **76**, 083012
 Kim, D.-W., & Fabbiano, G. 2004, *ApJ*, **611**, 846
 Kniffen, D. A., & Fichtel, C. E. 1981, *ApJ*, **250**, 389
 Large, M. I., Quigley, M. J. S., & Haslam, C. G. T. 1962, *MNRAS*, **124**, 405
 Linden, T., & Profumo, S. 2010, *ApJ*, **714**, L228
 Meade, P., Papucci, M., Strumia, A., & Volansky, T. 2010, *Nu. Ph. B*, **831**, 178
 Padovani, P., Giommi, P., Landt, H., & Perlman, E. S. 2007, *ApJ*, **662**, 182
 Porter, T. 2009, Oral presentation at TeV Particle Astrophysics (TeVPA) Conference, July 13–17, 2009, Menlo Park, CA
 Porter, T. A., & Strong, A. W. 2005, *ICRC (Pune)*, **4**, 77
 Regis, M., & Ullio, P. 2009, *Phys. Rev. D*, **80**, 043525
 Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, **500**, 525
 Slatyer, T. R., & Finkbeiner, D. P. 2010, *MNRAS*, in press (arXiv:0910.0482)
 Smialkowski, A., Wolfendale, A. W., & Zhang, L. 1997, *Astropart. Phys.*, **7**, 21
 Strong, A. W. 1984, *Adv. Space Res.*, **3**, 87
 Strong, A. W., Moskalenko, I. V., & Ptuskin, V. S. 2007, *Annu. Rev. Nucl. Part. Sci.*, **57**, 285
 Strong, A. W., Moskalenko, I. V., & Reimer, O. 2000, *ApJ*, **537**, 763
 Voss, R., & Gilfanov, M. 2006, *A&A*, **447**, 71
 Zhang, J., et al. 2009, *Phys. Rev. D*, **80**, 023007