Gamma-ray emission from the Sagittarius dwarf spheroidal galaxy due to millisecond pulsars

Roland M. Crocker ^[0,2,15], Oscar Macias ^{3,4,15]}, Dougal Mackey¹, Mark R. Krumholz¹, Shin'ichiro Ando^{3,4}, Shunsaku Horiuchi ^{4,5}, Matthew G. Baring ⁶, Chris Gordon ⁷, Thomas Venville⁸, Alan R. Duffy ⁸, Rui-Zhi Yang^{9,10,11}, Felix Aharonian ^{2,12}, J. A. Hinton², Deheng Song⁵, Ashley J. Ruiter ¹³ and Miroslav D. Filipović ¹⁴

The Fermi bubbles are giant, γ -ray-emitting lobes emanating from the nucleus of the Milky Way discovered in ~1-100 GeV data collected by the Large Area Telescope on board the Fermi Gamma-Ray Space Telescope. Previous work has revealed substructure within the Fermi bubbles that has been interpreted as a signature of collimated outflows from the Galaxy's supermassive black hole. Here we show via a spatial template analysis that much of the γ -ray emission associated with the brightest region of substructure—the so-called cocoon—is probably due to the Sagittarius dwarf spheroidal galaxy (dSph). This large Milky Way satellite is viewed through the Fermi bubbles from the position of the Solar System. As a tidally and ram-pressure stripped remnant, the Sagittarius dSph has no ongoing star formation, but we nevertheless demonstrate that the dwarf's millisecond pulsar population can plausibly supply the γ -ray signal that our analysis associates with its stellar template. The measured spectrum is naturally explained by inverse Compton scattering of cosmic microwave background photons by high-energy electron-positron pairs injected by millisecond pulsars belonging to the Sagittarius dSph, combined with these objects' magnetospheric emission. This finding plausibly suggests that millisecond pulsars produce significant γ -ray emission among old stellar populations, potentially confounding indirect dark-matter searches in regions such as the Galactic Centre, the Andromeda galaxy and other massive Milky Way dSphs.

arly analysis of data from the Fermi Large Area Telescope (Fermi-LAT; ref.¹) identified two counterpropagating, collinear γ -ray substructures within the Fermi bubbles (refs. ^{2,3}; Fig. 1a): a jet in the northern Galactic hemisphere and cocoon in the south⁴. Subsequent, independent analyses^{3,5} have only confirmed the existence of the latter. As the cocoon is contained within the solid angle of the surrounding Fermi bubbles and exhibits a similar γ -ray spectrum, it is natural to propose they share a common origin. However, the cocoon is also spatially coincident with the core of the Sagittarius dwarf spheroidal galaxy (Sgr dSph (ref.⁶); Fig. 1b), a satellite of the Milky Way that is in the process of being accreted and destroyed, as tidal forces gradually strip stars out of its core into elongated streams⁷. The chance probability of such an alignment is low, $\sim 1\%$ (Supplementary Section 1), even before accounting for the fact that the cocoon and the Sgr dSph have similar shapes and orientations and that the Sgr dSph is both one of the nearest and most massive (distance d = 26.5 kpc, mass $M \approx 10^8 \text{ solar masses } (M_{\odot})$; refs.^{8,9}) Milky Way satellites, with the largest mass divided by distance squared of any astronomical object not yet detected in γ -rays.

We therefore considered emission from the Sgr dSph as an alternative origin for the cocoon. To test this possibility, we fitted the γ -ray emission observed by Fermi-LAT over a region of

interest (ROI) containing the cocoon via template analysis. In our baseline model these templates include only known point sources and sources of Galactic diffuse γ -ray emission. We contrasted the baseline with a baseline + Sgr dSph model that invoked these same templates plus an additional template constructed to be spatially coincident with the bright stars of the Sgr dSph (Extended Data Fig. 1 and Supplementary Fig. 1); full details of the fitting procedure are provided in the Methods and Supplementary Section 3. Choosing the best-motivated background templates (i.e., those that deliver the best goodness-of-fit we found in our analysis), we found that the baseline + Sgr dSph model was preferred at 8.1σ significance over the baseline model. We also repeated the analysis for a wide range of alternative templates for both Galactic diffuse emission and for the Sgr dSph (Table 1) and obtained $>5\sigma$ detections for all combinations but one. Moreover, even this is an extremely conservative estimate, because our baseline model uses a structured template for the Fermi bubbles that absorbs some of the signal that is spatially coincident with the Sgr dSph into a structure of unknown origin. If we follow the method recommended by the Fermi collaboration³ and use a flat Fermi bubble template in our analysis, the significance of our detection of the Sgr dSph is always >14 σ . Despite this, for the remainder of our analysis we followed the most conservative choice

¹Research School of Astronomy and Astrophysics, Australian National University, Canberra, Australian Capital Territory, Australia. ²Max-Planck-Institut für Kernphysik, Heidelberg, Germany. ³GRAPPA–Gravitational and Astroparticle Physics Amsterdam, University of Amsterdam, Amsterdam, the Netherlands. ⁴Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwa, Japan. ⁵Center for Neutrino Physics, Department of Physics, Virginia Tech, Blacksburg, VA, USA. ⁶Department of Physics and Astronomy, Rice University, Houston, TX, USA. ⁷School of Physical and Chemical Sciences, University of Canterbury, Christchurch, New Zealand. ⁸Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, Victoria, Australia. ⁹School of Astronomy and Space Science, University of Science and Technology of China, Hefei, China. ¹⁰Department of Astronomy, School of Physical Sciences, University of Science and Technology of China, Hefei, China. ¹¹CAS Key Labrotory for Research in Galaxies and Cosmology, University of Science and Technology of China, Hefei, China. ¹²Dublin Institute for Advanced Studies, Dublin, Ireland. ¹³School of Science, University of New South Wales Canberra, Australian Defence Force Academy, Canberra, Australian Capital Territory, Australia. ¹⁴Western Sydney University, Penrith, New South Wales, Australia. ¹⁵These authors contributed equally: Roland M. Crocker, Oscar Macias. ⁵⁶e-mail: rcrocker@fastmail.fm; o.a.maciasramirez@uva.nl



Fig. 1 The Fermi bubbles, including the cocoon substructure, and the Sgr dSph galaxy. a, c, The γ -ray spatial template for the Fermi bubbles³ in arbitrary units with a linear colour scale, highlighting the cocoon. **b**,**d**, The angular density of RR Lyrae stars with line-of-sight distances >20 kpc from Gaia Data Release 2, in arbitrary units with logarithmic scaling. The Sgr dSph, Sgr stream and the Large and Small Magellanic (Mag) Clouds are clearly visible. The proper motion of the Sgr dSph is upwards. The dashed ellipses in **a**-**d** mark the same coordinates in each panel and highlight both the cocoon and the Sgr dSph. **e**, Contours of RR Lyrae surface density overlaid on the Fermi bubble template shown as the coloured background. **a** and **b** are all-sky views in Mollweide projection in Galactic coordinates of longitude ℓ and latitude *b* with east to the left, with the ROI marked by the dotted box; **c**-**e** are in a cylindrical projection and zoom in on the ROI.

by using the structured template in our baseline model. In the Methods we also show that our analysis passed a series of validation tests: the residuals between our best-fitting model and the data were consistent with photon counting statistics (Extended Data Figs. 2 and 3), our pipeline reliably recovered synthetic signals superimposed on a realistic background (Extended Data Fig. 4), fits using a template tracing the stars of the Sgr dSph yielded significantly better results than fits using purely geometric templates (Supplementary Table 1) and if we artificially rotated the Sgr dSph template on the sky, the best-fitting position angle was very close to the actual angle (Extended Data Fig. 5). By contrast, when we displaced the Sgr dSph template, we found moderate (4.5σ significance) evidence that the best-fitting position was ~4° from the true position, in a direction very closely aligned with the dwarf galaxy's direction of travel (Extended Data Fig. 5); this plausibly represents a small, but real and expected (as explained below), physical offset between the stars and the γ -ray emission.

The directly measured flux from the Sgr dSph, derived from our fiducial choice of templates, corresponds to a luminosity of $3.8 \pm 0.6 \times 10^{36} \text{ erg s}^{-1} (1\sigma \text{ error})$ for γ -ray photons in the range from 0.5 to 150 GeV (equivalently $\sim 4 \times 10^{28} \text{ erg s}^{-1} M_{\odot}$). Over this range, the spectrum is approximately described by a hard power law $dF_{\gamma}/dE_{\gamma} \approx E_{\gamma}^{-2.1}$, where E_{γ} denotes photon energy and dF_{γ}/dE_{γ} is the number flux of photons due to the source, differential in photon energy (Fig. 2). There is no evidence for a cutoff at high energies. We show in Extended Data Fig. 6 that this spectral shape is qualitatively insensitive to the choice of foreground templates, and Extended Data Fig. 7 demonstrates that the spectra we recovered for the various foregrounds within the ROI remain physically plausible when we introduced a Sgr dSph template.

As our template fits plausibly suggested that there is a real γ -ray emission component tracing the Sgr dSph, a natural next question was what mechanism could be responsible for producing it. The core of the Sgr dSph is the remnant of a once much more

massive galaxy. Tidal and ram-pressure stripping removed its gas and caused it to cease forming stars 2-3 Gyr ago¹⁰, although it did experience punctuated bursts of star formation¹¹—triggered by its crossings through the Galactic plane¹²—up to that time. In the Milky Way, the dominant source of diffuse γ -ray emission is collisions between (hadronic) cosmic rays and ambient interstellar medium gas nuclei13, but this mechanism cannot operate in the Sgr dSph, which lacks both 'target' gas with which cosmic rays could interact, and supernova explosions from young, massive stars to accelerate hadronic cosmic rays in the first place. Stellar y-ray emission is also ruled out: while our Sun is a source of ~100 GeV γ -rays, this emission is again dominantly due to collisions between hadronic cosmic rays from the wider Galaxy and solar gas; γ -ray emission from non-thermal particles accelerated by the Sun itself only extends to 4 GeV (ref. ¹⁴). This leaves two possibilities for the γ -ray signal our template analysis associates with the Sgr dSph: it is created from the self-annihilation of dark-matter particles in the dwarf's dark-matter halo, or by millisecond pulsars (MSPs) deriving from the stars of the Sgr dSph. The former is unlikely because the γ -ray signal largely traces the stars of the dwarf, while N-body simulations¹² show that the Milky Way's tidal field will have overwhelmingly dispersed the progenitor galaxy's original dark-matter halo into the stream over its orbital history.

MSPs, by contrast, should follow the same spatial distribution as the rest of the stellar population, have a spin-down timescale $\gtrsim \mathcal{O}$ (Gyr), long enough to be compatible with the most recent episodes of Sgr dSph star formation, and radiate some part of their magnetic dipole luminosity into γ -rays. However, there are two significant challenges to this scenario: first, the inferred γ -ray luminosity per unit stellar mass is much larger ($\gtrsim 10\times$) for the Sgr dSph than for some other systems whose detected γ -ray emission is plausibly dominated by MSPs including the Galactic Bulge¹⁵⁻¹⁹ and Andromeda^{20,21} (M31), the giant spiral galaxy nearest to the Milky Way (although it is smaller than that observed for globular



Fig. 2 | Measured γ -ray spectral brightness distributions of the signal associated with the Sgr dSph template and the surrounding Fermi bubbles. The dashed black line shows a differential number flux obeying $dN_{\gamma}/dE_{\gamma} \propto E_{\gamma}^{-2.1}$. These data are as obtained by us in our Fermi-LAT data analysis described in the Methods. We have converted luminosities to surface brightnesses, adopting source solid angles of $\Omega_{\rm sgrdSph} = 9.6 \times 10^{-3} {\rm sr}$ and $\Omega_{\rm FB} = 0.49 {\rm sr}$ with the latter set by the 40° × 40° ROI, not the intrinsic sizes of the bubbles (which are larger than the ROI). Error bars show 1 σ errors; for the Sgr dSph, the error bars incorporate both statistical and systematic errors added in quadrature. The blue curves show the best-fit combined (magnetospheric + IC; solid) and best-fit magnetospheric (dashed) spectra.

clusters; see Fig. 3). Second, the hard, $\stackrel{\sim}{\sim} E_{\gamma}^{-2.1}$ spectrum of the Sgr dSph (Fig. 2) does not resemble the classic few-gigaelectronvolt bump (in the spectral energy distribution) of the magnetospheric γ -ray signal detected from individual MSPs or the globular clusters that host populations of MSPs (for example, ref. ²²).

However, both of these challenges can be overcome by considering how the stellar population and interstellar environment of the Sgr dSph differ from other systems. With regard to stars, those in the Sgr dSph are both younger and more metal-poor than those of M31 or the Galactic Bulge; metal-poor stellar systems are expected to produce more MSPs per stellar mass²³, and ~7–8-Gyr-old MSPs (the rough age of the Sgr dSph population) are expected to be significantly brighter than 10–12-Gyr-old ones (the ages of stellar populations in the Galactic Bulge and the core of M31)¹⁹. In Supplementary Section 5 we show that the best-fit value for the γ -ray luminosity of the Sgr dSph is fully consistent with both theoretical predictions and with observations of other γ -ray-emitting old stellar populations once age and metallicity are taken into account. On the basis of stellar population synthesis models, we estimate that the γ -ray luminosity of the Sgr dSph is produced by ~650 MSPs.

With regard to environment, note that while the spectrum of the Sgr dSph does not resemble an MSP magnetospheric signal, it does resemble inverse Compton emission from the upscattering (by a cosmic ray electron–positron population; e^{\pm}) of ambient light that, for the Sgr dSph, is dominated by the cosmic microwave background. We also know that MSPs produce e^{\pm} with energies of at least a few teraelectronvolts, as these are the particles that ultimately drive the observed gigaelectronvolt MSP γ -ray photospheric emission. Some of these e^{\pm} will give up all their energy within the magnetosphere of the MSP. However, given the expected absence of wind nebulae or supernova remnants surrounding these old, low-luminosity objects²⁴, many will freely escape both the magnetosphere and MSP environs into the larger Sgr dSph environment^{25,26}, where they can inverse Compton upscatter cosmic microwave background photons. In an environment like Andromeda or the Galactic



Fig. 3 | Gamma-ray luminosity L₇ **normalized to stellar mass** M_★ **for various structures whose emission is plausibly dominated by MSPs.** The Sgr magneto. data point shows our best-fit magnetospheric luminosity per stellar mass (the spectrum shown as the dashed blue curve in Fig. 2) while the Sgr tot data point is the total measured luminosity associated with the Sgr dSph template. Globular cluster measurements are from ref.²², while the remaining data (collated by ref.²²) are from ref.¹⁸ (nuclear bulge of the Milky Way (NB)), ref.²⁰ (M31) and ref.¹⁷ (Milky Way disk). Error bars show 1 σ errors. The horizontal dashed grey lines show the predicted total γ -ray luminosity per unit stellar mass at the nominated efficiencies, $f_{\gamma,tot} = \{0.1, 0.9\}$, given an MSP spin-down power per unit stellar mass of 2×10^{28} erg s⁻¹M₀⁻¹ as we inferred from ref.²⁴.

Bulge, this inverse Compton signal will be weak (albeit detectable in the case of the Galactic Bulge according to ref.¹⁹), because much of the escaping e^{\pm} energy will be lost to synchrotron (rather than inverse Compton) radiation. In an ultra-gas-poor system like the Sgr dSph, however, we expect the magnetic field of the interstellar medium to be far weaker than in a gas-rich galaxy (ref. 27; also see Methods) with an energy density significantly smaller than that in the cosmic microwave background; thus radiative losses from MSP-escaping e^{\pm} are overwhelmingly into hard-spectrum inverse Compton γ -rays, rather than (radio to X-ray) synchrotron radiation. Consistent with this explanation, globular clusters-which are also gas-poor and weakly magnetized-represent another environment where MSP-driven γ -ray emission sometimes seems to include a significant inverse Compton component²². We formalize this intuitive argument in the Methods, where we show that the spectrum of the Sgr dSph is extremely well fitted by a combination of inverse Compton and magnetospheric radiation with self-consistently related spectral parameters. This scenario also explains why the γ -ray signal is displaced ~4°, or about 1.9 kpc (right panel of Extended Data Fig. 5), from the centre of the Sgr dSph; the dwarf's northward proper motion²⁸ means that this displacement is backwards along its path. As the Sgr dSph plunges through the Milky Way halo, the magnetic field around it will be elongated into a magnetotail oriented backwards along its trajectory, and e^{\pm} emitted into the dwarf will be trapped by these magnetic field lines, leading them to accumulate and emit in a position that trails the Sgr dSph, exactly as we observe. We offer a more quantitative evaluation of this scenario in Supplementary Section 4.

There are some caveats to our results that the reader should note. First, in common with other Fermi-LAT data analyses of diffuse emission from extended regions, it is evident that our model, although very good, does not reproduce the data accurate down to the level of Poisson noise over the entire ROI (see discussion in Supplementary Section 3). Indeed, Extended Data Fig. 3 shows that

NATURE ASTRONOMY

Table 1 | Template analysis results comparing the baseline and baseline + Sgr dSph models

Template choices				Results			
Hadronic/bremsstrahlung emission	Inverse Compton emission	Fermi bubble	Sgr dSph model	$-\log(\mathcal{L}_{\text{Base}})$	$-\log(\mathcal{L}_{\text{Base}+\text{Sgr}})$	TS _{Source}	Significance
Default model							
HD	3D	S	Model I	866,680.6	866,633.0	95.2	8.1 <i>o</i>
Alternative background templates							
HD	2D A	S	Model I	866,847.1	866,810.9	72.3	6.9 <i>o</i>
HD	2D B	S	Model I	867,234.9	867,192.1	85.8	7.8 <i>o</i>
HD	2D C	S	Model I	866,909.4	866,868.5	81.7	7.4 <i>o</i>
Interpolated	3D	S	Model I	867,595.4	867,567.4	56.0	5.8 <i>o</i>
GALPROP	3D	S	Model I	866,690.5	866,640.8	99.5	8.3 <i>o</i>
Flat Fermi bubble template							
HD	3D	U	Model I	867,271.7	867,060.1	423.2	19.1 <i>o</i>
HD	2D A	U	Model I	867,284.2	867,122.9	322.5	16.5 <i>o</i>
HD	2D B	U	Model I	867,624.3	867,464.0	320.7	16.4 <i>o</i>
HD	2D C	U	Model I	867,322.7	867,158.2	329.0	16.6 <i>o</i>
Interpolated	3D	U	Model I	867,287.4	867,081.2	412.4	18.9 <i>o</i>
GALPROP	3D	U	Model I	868,214.6	868,040.9	347.6	17.2 <i>σ</i>
Alternative Sgr dSph templates							
HD	3D	S	Model II	866,680.6	866,626.3	108.5	8.7 <i>o</i>
HD	3D	S	Model III	866,680.6	866,647.5	66.1	6.4 <i>o</i>
HD	3D	S	Model IV	866,680.6	866,678.2	4.8	0.4σ
HD	3D	S	Model V	866,680.6	866,644.9	71.5	6.7 <i>o</i>
HD	3D	U	Model II	867,271.7	866,970.7	602.1	23.2 <i>o</i>
HD	3D	U	Model III	867,271.7	866,994.1	555.3	22.2 <i>o</i>
HD	3D	U	Model IV	867,271.7	867,152.2	239.1	14.0 <i>o</i>
HD	3D	U	Model V	867,271.7	866,993.3	556.9	22.2 <i>o</i>

The source templates describing the Sgr dSph are details in the Methods. The log likelihoods \mathcal{L} for the baseline model (without the Sgr dSph) and the baseline + Sgr dSph model are given. TS_{source} is the test statistic with which the baseline + Sgr dSph model was preferred and the corresponding statistical significance of that preference is provided. The improvement in TS from {HD, 3D, U, Model } to {HD, 3D, S, Model I} to {Show and Sgr dSph model was preferred and the corresponding statistical significance of that preference is provided. The improvement in TS from {HD, 3D, U, Model I} to {HD, 3D, S, Model I} to {Show and Sgr dSph model was preferred and the corresponding statistical significance of that preference is provided. The improvement in TS from {HD, 3D, U, Model I} to {HD, 3D, to that Sgr dSph Model IV (which generated a statistically insignificant improvement to the baseline for one particular combination in the last cluster) is the sparsest stellar template, containing only 675 stars.

there are structured residuals within the ROI, although we note that the strongest of these are at the edges of the ROI and not coincident with the Sgr dSph. We do not believe, therefore, that these residuals indicate that the detection of the signal connected to the Sgr dSph stellar template made in our γ -ray analysis is spurious, nor that the spectrum we measured is likely to be in significant error (see Extended Data Fig. 4). Instead, we suspect that the structured residuals point to the existence of still-mismodelled substructure in the Fermi bubbles that is completely unrelated to the Sgr dSph. Thus, while we argue on the basis of our analysis that much of the cocoon substructure is probably emission from the Sgr dSph, we do not claim to explain all Fermi bubble substructure.

This point connects to a second caveat: we are aware of no independent, multi-wavelength (non γ -ray) evidence for the existence of a well-defined nuclear jet or jets on angular scales comparable to the Fermi bubbles. Thus, in distinction to the case presented by the Sgr dSph (where we could construct from independent, multi-wavelength data a spatial template to incorporate into our γ -ray analysis), we cannot construct any definitive, a priori jet template. Although we argue that this is actually a weakness of the jet hypothesis, it nevertheless is true that we cannot via a formal statistical analysis rule out the presence of γ -ray substructure in the Fermi bubbles that is connected to a nuclear jet.

Taking note of all the above, there are a number of potential implications of the discovery of a γ -ray signal associated with the

Sgr dSph stellar template to follow up. First, our results motivate the introduction of stellar templates into the analysis of data from all γ -ray resolved galaxies (M31 and the Large and Small Magellanic Clouds) to probe the contribution of MSPs. Such studies may confirm (or not) that the relatively strong signal our analysis associated with the Sgr dSph stellar template can be explained reasonably via MSP emission (see Supplementary Fig. 3 and Supplementary Section 6). Second, our study lends support to the argument²⁴ that MSPs contribute significantly to the energy budget of CR e^{\pm} in galaxies with low specific star formation rates. Third, we show in the Supplementary Information that a direct extrapolation of the Sgr dSph MSP γ-ray luminosity per unit mass to other nearby dSph galaxies suggests that they could have considerably larger astrophysical γ -ray signatures than previous estimates; we report our revised estimates in Supplementary Table 3 for the sample of ref.²⁹. These signals are large enough that some are potentially detectable via careful analysis of Pass 8 (15-yr) Fermi-LAT data. Conversely, these brighter astrophysical signatures represent a larger-than-expected background with which searches for dark-matter annihilation signals (due to putative weakly interacting massive particles in the tens of gigaelectronvolts mass range) must contend, and potentially swamp dark-matter signals in some nearby dwarfs. We emphasize that these are not predictions per se, but naive extrapolations that do not account for peculiarities of the Sgr dSph with respect to other dSphs that may render it anomalously γ -ray efficient (for example,

its relatively recent star formation). These extrapolations do, nevertheless, motivate further work to pin down in detail how the γ -ray luminosity of an MSP population scales with gross parameters of the host stars (mass, age, metallicity and so on).

Methods

Our analysis pipeline consists of three steps: (1) data and template selection, (2) fitting and (3) spectral modelling.

Data and template selection. We used 8 years of LAT data, selecting Pass 8 UltraCleanVeto class events in the energy range from 500 MeV to 177.4 GeV. We chose the limit at low energy to mitigate both the impact of γ -ray leakage from the Earth's limb and the increasing width of the point spread function at lower energies. We spatially binned the data to a resolution of 0.2° and divided it into 15 energy bins; the 13 lowest-energy bins were equally spaced in log energy, while the 2 highest-energy bins are twice that width to improve the signal-to-noise ratio. We selected data obtained over the same observation period as that used in the construction of the Fourth Fermi Catalogue³⁰ (4 August 2008 to 2 August 2016). The ROI of our analysis was a square region defined by $-45^{\circ} \le b \le -5^{\circ}$ and $30^{\circ} \ge \ell \ge -10^{\circ}$ (Fig. 1). This sky region fully contains the Fermi cocoon substructure but avoids the Galactic plane ($|b| \le 5^{\circ}$) where uncertainties are largest. Because the ROI is of modest size, we allowed the Galactic diffuse emission templates greater freedom to reproduce potential features in the data. We carried out all data reduction and analysis using the standard FERMITOOLS V1.0.1 software package (available from https://github.com/fermi-lat/ Fermitools-conda/wiki). We modelled the performance of the LAT with the P8R3_ULTRACLEANVETO_V2 instrument response functions.

We fitted the spatial distribution of the ROI data as the sum of a series of templates for different components of the emission. For all the templates we considered, we defined a baseline model that included only known point and diffuse emission sources, to which we compared a baseline + Sgr dSph model that included those templates plus the Sgr dSph. Our baseline models, following the approach of ref.³¹, contained the following templates: (1) diffuse isotropic emission, (2) point sources, (3) emission from the Sun and Moon, (4) Loop I, (5) the Galactic centre excess, (6) Galactic cosmic ray-driven hadronic and bremsstrahlung emission, (7) inverse Compton emission and (8) the Fermi bubbles; baseline + Sgr dSph models also included a Sgr dSph template.

Our templates for the first five emission sources were straightforward, and we adopted a single template for each of them throughout our analysis. As our data selection was identical to that used to construct the Fourth Ferrni Catalogue, we adopted the standard isotropic background and point-source models provided as part of the catalogue³⁰, iso_P8R3_ULTRACLEANVETO_V2_v1.txt and gll_psc_v20.fit, respectively; the latter includes 177 γ -ray point sources within our ROI. We similarly adopted the standard Sun and Moon templates provided. For the foreground structure Loop I, we adopted the model of ref.³². Finally, given that the low-latitude boundary of our ROI overlaps with the spatial tail of the Galactic centre excess, we included the 'Boxy Bulge' template of ref.³³, which has been shown^{16–18} to provide a good description of the observed Galactic centre excess away from the nuclear Bulge region (which is outside our ROI). The inclusion of this template in our ROI model had only a small impact on our results.

The remaining templates required more care. The dominant source of y-rays within the ROI is hadronic and bremsstrahlung emission resulting from the interaction of Milky Way cosmic ray protons and electrons with interstellar gas; the emission rate is proportional to the product of the gas density and the cosmic ray flux. We modelled this distribution using three alternative approaches. Our preferred approach followed that described in ref.¹⁶. We assumed that the spatial distribution of γ -ray emission traces the gas distribution from the hydrodynamical model of ref.³⁴, which gave a more realistic description of the inner Galaxy than alternatives. To normalize the emission, we divided the Galaxy into four rings spanning the radial ranges 0-3.5 kpc, 3.5-8.0 kpc, 8.0-10.0 kpc and 10.0-50.0 kpc, within which we treated the emission per unit gas mass in each of our 15 energy bins as a constant to be fitted. We refer to the template produced in this way as the HD model. Our first alternative was to use the same procedure of dividing the Galaxy into rings but describe the gas distribution within those rings using a template constructed from interpolated maps of Galactic H I and H₂, following the approach described in appendix B of ref. ³⁵; we refer to this as the Interpolated approach. Our third alternative, the GALPROP model, is the SA50 model described by ref.³⁶, which prescribes the full-sky hadronic cosmic ray emission distribution

We similarly needed a model for diffuse, Galactic inverse Compton emission the second largest source of background —which is a product of the cosmic ray electron flux and the interstellar radiation field (ISRF). As with hadronic emission, we considered four alternative distributions. Our default choice was the SA50 model described by ref.³⁶, which includes 3D models for the ISRF³⁷. We therefore refer to this as the 3D model. However, unlike in ref.³⁶, we used this model only to obtain the spatial distribution of the emission, not its normalization or energy dependence. Instead, we obtained these in the same way as for our baseline hadronic emission model; that is, we divided the Galaxy into four rings and left the total amount of emission in each ring at each energy as a free parameter to be fitted to the data; this approach reduced the sensitivity of our results to uncertainties in the electron injection spectrum and ISRF normalization. Our three alternatives to this are models 2D A, 2D B and 2D C, which correspond to models A, B and C as described by ref.³⁸, which model inverse Compton emission over the full sky under a variety of assumptions about cosmic ray injection and propagation, but rely on a 2D model for the ISRF.

The final component of our baseline template was a model for the Fermi bubbles themselves, which are one of the strongest sources of foreground emission in high-latitude regions of the ROI. The Fermi bubbles are themselves defined as highly statistically significant and spatially coherent residuals in the inner Galaxy that remain once other sources are modelled out in all-sky γ-ray analyses. The Fermi bubbles are not reliably traced by emission at any other wavelength, so we do not have an a priori model with which to guide the construction of a spatial template of these structures. However, one characteristic that renders the Fermi bubbles distinct from other large-angular-scale diffuse γ -ray structures is their hard γ -ray spectrum. Indeed, the state-of-the-art, structured spatial template for them generated by the Fermi Collaboration³-the templates one would normally employ in large-ROI, inner-Galaxy Fermi-LAT analyses-were constructed using a spectral component analysis. That study recovered a number of regions of apparent substructure within the solid angle of the Fermi bubbles, most notably substructure overlapping the previously discovered^{4,5} 'cocoon', which (as we have discussed here) is largely coincident with the Sgr dSph. Of course, a potential issue with constructing a phenomenological, spectrally defined model for the Fermi bubbles is that, if there happens to be an extended, spectrally similar source coincident with the Fermi bubbles, it will tend to be incorporated into the template. For this reason, ref.3 suggested using a flat Fermi bubble template when searching for new structures. Despite this proposal, our default analysis used the more conservative choice of a structured Fermi bubble template. However, we also ran tests using an unstructured template for comparison, and to understand the systematic uncertainties associated with the choice of template. We refer to these two cases as the U (Unstructured) and S (Structured) Fermi bubble templates, respectively,

Finally, our baseline + Sgr dSph models required a template for the Sgr dSph. Our templates traced the distribution of bright stars in the dwarf, which we constructed from five alternative stellar catalogues, all based on different selections from Gaia Data Release 2. We refer to the resulting templates as models I–V and show them in Extended Data Fig. 1. Full details of how we constructed each of these templates are provided in Supplementary Section 2. Model I, our default choice, came from the catalogue of 2.26×10^5 Sgr dSph candidate member stars from ref.⁸; the majority of the catalogue consists of red clump stars. Model II used the catalogue of RR Lyrae stars in the Sagittarius Stream from ref.¹⁹, of which we have selected a sample of 2,369 stars whose kinematics are consistent with being members of the Sgr dSph itself. Model III used the catalogue of 1.31 × 10⁴ RR Lyrae stars belonging to the Sgr dSph provided by ref.⁴⁰. Finally, models IV and V come from the nGC3 and Strip catalogues of RR Lyrae stars from ref.⁴¹; the former contains 675 stars with higher purity but lower completeness, while the latter contains 4,812 stars of higher completeness but lower purity.

Fitting procedure. Our fitting method followed that introduced in refs. 16,18 and treated each of the 15 energy bins as independent, thereby removing the need to assume any particular spectral shape for each component and allowing the spectra to be determined solely by the data. Our data to be fitted consisted of the observed γ -ray photon counts in each spatial pixel *i* and energy bin *n*, which we denote $\Phi_{n,iobs}$, where *n* ranges from 1 to 15, and the index *i* runs over the positions (ℓ_i,b_i) of all spatial pixels within the ROI. For a given choice of template, the corresponding model-predicted γ -ray counts were $\Phi_{n,i,\text{mod}} = \sum_{c} \mathcal{N}_{n,c} R_{n,i} \Phi_{c,i}$ where $\mathcal{N}_{n,c}$ is defined below, $R_{n,i}$ is the instrument response for each pixel and energy bin (computed assuming an E^{-2} spectrum within the bin), and Φ_{ci} is the value of template component c evaluated at pixel i; for baseline models, we had a total of eight components, while for the baseline + Sgr dSph models, we had nine. Note that Φ_{ci} is a function of *i* but not of *n*; that is, we assumed that the spatial distribution of each template component was the same at all energies, except for the inverse Compton templates, for which an energy-dependent morphology was predicted by our GALPROP simulations. Without loss of generality, we further normalized each template component as $\sum_{i} \Phi_{c,i} = 1$, for which $\mathcal{N}_{n,c}$ is simply the total number of photons contributed by component *c* in energy bin *n* integrated over the full ROI; the values of $\mathcal{N}_{n,c}$ were the parameters to be fitted. We found the best fit by maximizing the usual Poisson likelihood function:

$$\ln \mathcal{L}_n = \sum_{i} \frac{\Phi_{n,i,\text{mod}}^{\Phi_{n,i,\text{mod}}} e^{-\Phi_{n,i,\text{mod}}}}{\Phi_{n,i,\text{obs}}!},\tag{1}$$

using the pylikelihood routine, the standard maximum likelihood method in FermiTools. Note that as each energy bin n is independent, we carried out the likelihood maximization bin-by-bin.

We performed all fits in pairs, one for a baseline model containing only known emission sources, and one for a baseline + Sgr dSph model containing the same known sources plus a component tracing the Sgr dSph. The set of paired fits we performed in this manner is shown in Table 1. We compared the

quality of these baseline and baseline + Sgr dSph fits by defining the test statistic $TS_n = -2 \ln(\mathcal{L}_{n,Base} / \mathcal{L}_{n,Base} + Sgr)$; the total test statistic for all energy bins is simply $TS = \sum_n TS_n$. We could assign a *P* value to a particular value of the TS by noting that baseline + Sgr dSph models have 15 additional degrees of freedom compared with the baseline models: the value of $\mathcal{N}_{n,c}$ for the component *c* corresponding to the Sgr dSph, evaluated at each of the 15 energy bins. In this case, the mixture distribution formula gives¹⁶:

$$P(\mathrm{TS}) = 2^{-N} \left[\delta(\mathrm{TS}) + \sum_{n=1}^{N} \binom{N}{n} \chi_n^2(\mathrm{TS}) \right], \qquad (2)$$

where N = 15 is the difference in the number of degrees of freedom, $\binom{N}{n}$ is the

binomial coefficient, δ is the Dirac delta function and χ_n^2 is the usual χ^2 distribution with *n* degrees of freedom. The corresponding statistical significance (in σ units) is¹⁶:

Number of
$$\sigma \equiv \sqrt{\text{InverseCDF}\left(\chi_1^2, \text{CDF}\left[P(\text{TS}), \hat{\text{TS}}\right]\right)},$$
 (3)

where (InverseCDF) CDF is the (inverse) cumulative distribution function and the first argument of each of these functions is the distribution function, the second is the value at which the CDF is evaluated, and the total TS is denoted by \hat{TS} . For 15 extra degrees of freedom, a 5σ detection corresponds to TS = 46.1 (additional details of these formulae are given in supplementary section 2 of ref.¹⁶.) We report values of \mathcal{L}_{Base} , $\mathcal{L}_{Base+Sgr}$, TS and the significance level for all the templates we used in Table 1.

The final step in our fitting chain was to assess the uncertainties. For our default choice of baseline + Sgr dSph model (first row in Table 1), our maximum likelihood analysis returned the central value \mathcal{N}_n^{def} on the total γ -ray flux in the *n*th energy bin attributed to the Sgr dSph and also yielded an uncertainty $\sigma_{\mathcal{N},n}^{def}$ on this quantity. This represents the statistical error arising from measurement uncertainties. However, there are also systematic uncertainties stemming from our imperfect knowledge of the templates characterizing the other emission sources. To estimate these, we examined the five alternative models listed in Table 1 as 'Alternative background templates', for which we used different templates for the hadronic plus bremsstrahlung and inverse Compton backgrounds. Each of these models *m* also returned a central value \mathcal{N}_n^m and an uncertainty $\sigma_{\mathcal{N},n}^m$ on the Sgr dSph flux. We used the uncertainty-weighted dispersion of these models as an estimate of the systematic uncertainty (for example, ref.⁴²):

$$\delta \mathcal{N}_n = \sqrt{\frac{1}{\sum_m \left(\sigma_{\mathcal{N},n}^m\right)^{-2}} \sum_m \left(\sigma_{\mathcal{N},n}^m\right)^{-2} \left(\mathcal{N}_n^{\text{def}} - \mathcal{N}_n^m\right)^2},\tag{4}$$

where the sums run over the m = 6 - 1 alternative models. We took the total uncertainty on the Sgr dSph flux in each energy bin to be a quadrature sum of the systematic and statistical uncertainties; that is, $(\sigma_{\mathcal{N},n}^{\text{defind}})^2 = (\sigma_{\mathcal{N},n}^{\text{def}})^2 + \delta \mathcal{N}_n^2$. We plotted the central values and uncertainties of the fluxes for the default model derived in this manner in Fig. 2.

We carried out several validation tests of this pipeline, which we describe in the Supplementary Information.

Spectral modelling. We modelled the observed Sgr dSph γ -ray spectrum as a combination of prompt magnetospheric MSP emission and inverse Compton emission from e^{\pm} escaping MSP magnetospheres. We constructed this model as follows. The prompt component was due to curvature radiation from e^{\pm} within MSP magnetospheres. The e^{\pm} energy distribution can be approximated as an exponentially truncated power law^{22,43}:

$$\frac{\mathrm{d}N_{\mathrm{MSP},e^{\pm}}}{\mathrm{d}E_{e^{\pm}}} \propto E_{e^{\pm}}^{\gamma_{\mathrm{MSP}}} \exp\left(-\frac{E_{e^{\pm}}}{E_{\mathrm{cut},e^{\pm}}}\right),\tag{5}$$

and curvature radiation from these particles has a rate of photon emission per unit energy per unit time:

$$\frac{\mathrm{d}\dot{N}_{\gamma,\mathrm{prompt}}}{\mathrm{d}E_{\gamma}} = \mathcal{N}\left(L_{\gamma,\mathrm{prompt}}\right) E_{\gamma}^{\alpha} \exp\left(-\frac{E_{\gamma}}{E_{\mathrm{cut},\mathrm{prompt}}}\right),\tag{6}$$

where the superscript dot on the left hand side indicates a quantity differential with respect to time, E_{γ} is the photon energy, $\mathcal{N}(L_{\gamma,\text{prompt}})$ is a normalization factor chosen so that the prompt component has total luminosity $L_{\gamma,\text{prompt}}$, the index α is related to that of the e^{\pm} distribution by $\alpha = (\gamma_{\text{MSP}} - 1)/3$ and the photon cutoff energy is related to the e^{\pm} cutoff energy by²⁵:

$$E_{\text{cut,prompt}} = \frac{3\hbar c}{2\rho_{\text{c}}} \left(\frac{E_{\text{cut,e}\pm}}{m_{\text{e}}}\right)^3 \approx 2.0 \text{ GeV} \left(\frac{\rho_{\text{c}}}{30 \text{ km}}\right)^{-1} \left(\frac{E_{\text{cut,e}\pm}}{3 \text{ TeV}}\right)^3 \tag{7}$$

where m_e is the electron mass, ρ_e is the radius of curvature of the magnetic field lines, c the velocity of light and \hbar the reduced Planck's constant. Given the relatively

small magnetospheres, we expect ρ_c to be a small multiple of the ~10 km neutron star characteristic radius; henceforth, we set $\rho_c = 30$ km. Empirically, $L_{\chi,prompt}$ is ~10% of the total MSP spin-down power⁴³.

A larger proportion of the spin-down power goes into a wind of e^{\pm} escaping the magnetosphere. In the ultra-low-density environment of the Sgr dSph, ionization and bremmstrahlung losses for this population (which occur at a rate proportional to the gas density) are negligible. Synchrotron losses, which scale as the magnetic energy density, will also be negligible; as noted in the main text, observed magnetic fields in dwarf galaxies are very weak²⁷, and we can also set a firm upper limit on the Sgr dSph magnetic field strength simply by noting that the magnetic pressure cannot exceed the gravitational pressure provided by the stars as, if it did, that magnetic field and the gas to which it is attached would blow out of the galaxy in a dynamical time. The gravitational pressure is $P \approx (\pi/2)G\Sigma^2$, where $\Sigma = M/\pi R^2$ is the surface density and G the gravitational constant, and using our fiducial numbers $M = 10^8 M_{\odot}$ and R = 2.6 kpc gives an upper limit on the magnetic energy density 0.06 eV cm-3; non-zero gas or cosmic ray pressure would lower this estimate even further. This is a factor of four smaller than the energy density of the cosmic microwave background, implying that synchrotron losses comprise at most 20% of losses, and can therefore be neglected.

This analysis implies that the only significant loss mechanism for these e^{\pm} is inverse Compton emission, resulting in a steady-state e^{\pm} energy distribution:

$$\frac{\mathrm{d}N_{e\pm}}{\mathrm{d}E_{e\pm}} \propto E_{e\pm}^{\gamma} \exp\left(-\frac{E_{e\pm}}{E_{\mathrm{cut},e\pm}}\right),\tag{8}$$

where $\gamma = \gamma_{\rm MSP} - 1$. We computed the IC photon distribution produced by these particles following ref. ⁴⁴, assuming that the ISRF of the Sgr dSph is the sum of the cosmic microwave background and two subdominant contributions, one consisting of light escaping from the Milky Way and the other a dilute stellar blackbody radiation field due to the stars of the dwarf. We estimated the Milky Way contribution to the photon field at position of the dwarf using GALPROP³⁷, which predicted a total energy density of 0.095 eV cm⁻³ (compared with 0.26 eV cm⁻³ for the cosmic microwave background), comprised of five dilute black bodies with colour temperatures and dilution factors { $T_{\rm rad}$, κ } as follows: {40 K, 1.4×10^{-6} }, {430 K, 3.0×10^{-11} }, {3.400 K, 4.3×10^{-14} }, {6.400 K, 4.0×10^{-15} } and {26,000 K, 8.0×10^{-18} }. We characterized the intrinsic light field of the dwarf as having a colour temperature of 3.500 K and dilution factors 7.0×10^{-15} (giving an energy density of 0.005 eV cm^{-3} ; these choices are those expected for a spherical region of radius 2.6 kpc and stellar luminosity $2 \times 10^8 L_{\odot}$, the approximate parameters of the Sgr dSph). This yields an inverse Compton spectrum:

$$\frac{\mathrm{d}\dot{N}_{\gamma,\mathrm{IC}}}{\mathrm{d}E_{\gamma}} = \mathcal{N}\left(L_{\gamma,\mathrm{IC}}\right)F\left(\gamma, E_{\mathrm{cut},e\pm}\right),\tag{9}$$

where $\mathcal{N}(L_{\gamma,\text{IC}})$ is again a normalization chosen to ensure that the total inverse Compton luminosity is $L_{\gamma,\text{IC}}$, and $F(\gamma, E_{\text{cut},e^{\pm}})$ is the functional form given by equation (14) of ref.⁴⁴, which depends on the e^{\pm} spectral index γ and $E_{\text{cut},e^{\pm}}$.

Combining the prompt and inverse Compton components, we can therefore write the complete emission spectrum as:

$$\frac{\mathrm{d}\dot{N}_{\gamma}}{\mathrm{d}E_{\gamma}} = \mathcal{N}\left(L_{\gamma,\mathrm{prompt}}\right)E_{\gamma}^{\alpha}\exp\left(-\frac{E_{\gamma}}{E_{\mathrm{cut},\mathrm{prompt}}}\right) + \mathcal{N}\left(L_{\gamma,\mathrm{IC}}\right)F\left(\gamma,E_{\mathrm{cut},e^{\pm}}\right). \quad (10)$$

This model is characterized by four free parameters: the total prompt plus inverse Compton luminosity $L_{\gamma,\text{tot}} = L_{\gamma,\text{prompt}} + L_{\gamma,\text{LC}}$ the ratio of the prompt and inverse Compton luminosities $f = L_{\gamma,\text{prompt}} / L_{\gamma,\text{LC}}$ the spectral index α of the prompt component (which in turn fixes the other two spectral indices γ_{MSP} and γ) and the cutoff energy for the prompt component $E_{\text{cut,prompt}}$ (which then fixes $E_{\text{cut},e^{\pm}}$). Note that we made the simplest assumption that α and $E_{\text{cut,prompt}}$ are uniform across the MSP population. In reality, there may be a distribution of these properties but the parameteric form of equation (5) provides a good description, in general, of both individual MSP spectra and the aggregate spectra of globular cluster MSP populations²².

We fitted the observed Sgr dSph spectrum to this model using a standard χ^2 minimization, using the combined statistical plus systematic uncertainty. We obtained an excellent fit: the minimum χ^2 is 7.7 for 15 (data points) – 4 (fit parameters) = 11 degrees of freedom or a reduced χ^2 of 0.70. We report the best-fitting parameters in Supplementary Table 2 and plot the result best-fit spectra over the data in Fig. 2; we show the best-fit estimate (with $\pm 1\sigma$ confidence regions) for the magnetospheric luminosity per stellar mass of the Sgr dSph MSPs in Fig. 3.

We also carried out an additional consistency check, by comparing our best-fit parameters describing the prompt emission (α and $E_{cut,prompt}$) to direct measurements of the prompt component from nearby, resolved MSPs^{22,43} and to measurements of globular clusters, whose emission is probably dominated by unresolved MSPs²². We carried out this comparison in Supplementary Fig. 2, where we show joint confidence intervals on α and $E_{cut,prompt}$ from our fit. We constructed confidence intervals for α and $E_{cut,prompt}$ from observations using the sample of ref.²², who fitted the prompt emission from 40 globular clusters and 110 individually

ARTICLES

resolved MSPs. We drew 100,000 Monte Carlo samples from these fits, treating the stated uncertainties as Gaussian, and constructed contours in the ($E_{\rm cut,prompt}\alpha$) plane containing 68%, 95% and 99% of the sample points. As the plot shows, the confidence region from our fit is fully consistent with the confidence regions from the observations, indicating that our best-fit parameters are fully consistent with those typically observed for MSPs and globular clusters.

Data availability

All data analysed in this study are publicly available. Fermi-LAT data are available from https://fermi.gsfc.nasa.gov/ssc/data/ and Gaia data are available from https:// gea.esac.esa.int/archive/. The statistical pipeline, astrophysical templates and gamma-ray observations necessary to reproduce our main results are publicly available via Zenodo at https://doi.org/10.5281/zenodo.6210967.

Code availability

Fermi-LAT data used in our study were reduced and analysed using the standard FERMITOOLS V1.0.1 software package available from https://github.com/fermi-lat/ Fermitools-conda/wiki. The performance of the Fermi-LAT was modelled with the P8R3_ULTRACLEANVETO_V2 instrument response functions. Spectral analysis and fitting were performed using custom MATHEMATICA code created by the authors, which is available from RMC upon reasonable request.

Received: 13 May 2022; Accepted: 29 July 2022; Published online: 5 September 2022

References

- Atwood, W. B. et al. The Large Area Telescope on the Fermi Gamma-Ray Space Telescope mission. Astrophys. J. 697, 1071–1102 (2009).
- Su, M., Slatyer, T. R. & Finkbeiner, D. P. Giant gamma-ray bubbles from Fermi-LAT: active galactic nucleus activity or bipolar galactic wind? *Astrophys. J.* 724, 1044–1082 (2010).
- Ackermann, M. et al. The spectrum and morphology of the Fermi bubbles. Astrophys. J. 793, 64 (2014).
- 4. Su, M. & Finkbeiner, D. P. Evidence for gamma-ray jets in the Milky Way. *Astrophys. J.* **753**, 61 (2012).
- Selig, M., Vacca, V., Oppermann, N. & Enßlin, T. A. The denoised, deconvolved, and decomposed Fermi γ-ray sky. an application of the D³PO algorithm. *Astron. Astrophys.* 581, A126 (2015).
- Ibata, R. A., Gilmore, G. & Irwin, M. J. A dwarf satellite galaxy in Sagittarius. Nature 370, 194–196 (1994).
- Belokurov, V. et al. The field of streams: Sagittarius and its siblings. Astrophys. J. Lett. 642, L137–L140 (2006).
- 8. Vasiliev, E. & Belokurov, V. The last breath of the Sagittarius dSph. *Mon. Not. R. Astron. Soc.* **497**, 4162–4182 (2020).
- 9. Vasiliev, E., Belokurov, V. & Erkal, D. Tango for three: Sagittarius, LMC, and the Milky Way. *Mon. Not. R. Astron. Soc.* **501**, 2279–2304 (2021).
- 10. Siegel, M. H. et al. The ACS survey of galactic globular clusters: M54 and young populations in the Sagittarius dwarf spheroidal galaxy. *Astrophys. J. Lett.* **667**, L57–L60 (2007).
- Weisz, D. R. et al. The star formation histories of Local Group dwarf galaxies.
 I. Hubble Space Telescope/Wide Field Planetary Camera 2 observations. Astrophys. J. 789, 147 (2014).
- 12. Tepper-García, T. & Bland-Hawthorn, J. The Sagittarius dwarf galaxy: where did all the gas go? *Mon. Not. R. Astron. Soc.* **478**, 5263–5277 (2018).
- Strong, A. W. et al. Global cosmic-ray-related luminosity and energy budget of the Milky Way. *Astrophys. J. Lett.* **722**, L58–L63 (2010).
- 14. Linden, T. et al. Evidence for a new component of high-energy solar gamma-ray production. *Phys. Rev. Lett.* **121**, 131103 (2018).
- Abazajian, K. N. The consistency of Fermi-LAT observations of the galactic center with a millisecond pulsar population in the central stellar cluster. J. Cosmol. Astropart. Phys. 2011, 010 (2011).
- Macias, O. et al. Galactic bulge preferred over dark matter for the Galactic centre gamma-ray excess. *Nat. Astron.* 2, 387–392 (2018).
- 17. Bartels, R., Storm, E., Weniger, C. & Calore, F. The Fermi-LAT GeV excess as a tracer of stellar mass in the Galactic bulge. *Nat. Astron.* **2**, 819–828 (2018).
- 18. Macias, O. et al. Strong evidence that the galactic bulge is shining in gamma rays. *J. Cosmol. Astropart. Phys.* **2019**, 042 (2019).
- Gautam, A. et al. Millisecond pulsars from accretion-induced collapse as the origin of the Galactic Centre gamma-ray excess signal. *Nat. Astron.* 6, 703–707 (2022).
- Ackermann, M. et al. Observations of M31 and M33 with the Fermi Large Area Telescope: a galactic center excess in Andromeda? *Astrophys. J.* 836, 208 (2017).
- Eckner, C. et al. Millisecond pulsar origin of the galactic center excess and extended gamma-ray emission from Andromeda: a closer look. *Astrophys. J.* 862, 79 (2018).
- Song, D., Macias, O., Horiuchi, S., Crocker, R. M. & Nataf, D. M. Evidence for a high-energy tail in the gamma-ray spectra of globular clusters. *Mon. Not. R. Astron. Soc.* 507, 5161–5176 (2021).

- Ruiter, A. J. et al. On the formation of neutron stars via accretion-induced collapse in binaries. *Mon. Not. R. Astron. Soc.* 484, 698–711 (2019).
- Sudoh, T., Linden, T. & Beacom, J. F. Millisecond pulsars modify the radio-star-formation-rate correlation in quiescent galaxies. *Phys. Rev. D* 103, 083017 (2021).
- Baring, M. G. Perspectives on gamma-ray pulsar emission. In AstroPhysics of Neutron Stars 2010: A Conference in Honor of M. Ali Alpar Conference Series Vol. 1379 (eds Göğüş, E. et al.) 74–81 (American Institute of Physics, 2011).
- Venter, C., Kopp, A., Harding, A. K., Gonthier, P. L. & Büsching, I. Cosmic-ray positrons from millisecond pulsars. *Astrophys. J.* 807, 130 (2015).
- Regis, M. et al. Local Group dSph radio survey with ATCA II. Non-thermal diffuse emission. *Mon. Not. R. Astron. Soc.* 448, 3747–3765 (2015).
- del Pino, A. et al. Revealing the structure and internal rotation of the Sagittarius dwarf spheroidal galaxy with Gaia and machine learning. *Astrophys. J.* 908, 244 (2021).
- Winter, M., Zaharijas, G., Bechtol, K. & Vandenbroucke, J. Estimating the GeV emission of millisecond pulsars in dwarf spheroidal galaxies. *Astrophys.* J. Lett. 832, L6 (2016).
- Abdollahi, S. et al. Fermi Large Area Telescope fourth source catalog. Astrophys. J. Suppl. Ser. 247, 33 (2020).
- Abazajian, K. N., Horiuchi, S., Kaplinghat, M., Keeley, R. E. & Macias, O. Strong constraints on thermal relic dark matter from Fermi-LAT observations of the Galactic Center. *Phys. Rev. D* 102, 043012 (2020).
- 32. Wolleben, M. A new model for the Loop I (North Polar Spur) region. Astrophys. J. 664, 349–356 (2007).
- Freudenreich, H. T. A COBE model of the galactic bar and disk. Astrophys. J. 492, 495–510 (1998).
- Pohl, M., Englmaier, P. & Bissantz, N. Three-dimensional distribution of molecular gas in the barred Milky Way. Astrophys. J. 677, 283–291 (2008).
- Ackermann, M. et al. Fermi-LAT observations of the diffuse γ-ray emission: implications for cosmic rays and the interstellar medium. *Astrophys. J.* **750**, 3 (2012).
- Jóhannesson, G., Porter, T. A. & Moskalenko, I. V. The three-dimensional spatial distribution of interstellar gas in the Milky Way: implications for cosmic rays and high-energy gamma-ray emissions. *Astrophys. J.* 856, 45 (2018).
- Porter, T. A., Johannesson, G. & Moskalenko, I. V. High-energy gamma rays from the Milky Way: three-dimensional spatial models for the cosmic-ray and radiation field densities in the interstellar medium. *Astrophys. J.* 846, 67 (2017).
- Ackermann, M. et al. The spectrum of isotropic diffuse gamma-ray emission between 100 MeV and 820 GeV. Astrophys. J. 799, 86 (2015).
- 39. Ibata, R. et al. A panoramic landscape of the Sagittarius stream in Gaia DR2 revealed with the STREAMFINDER spyglass. *Astrophys. J. Lett.* **891**, L19 (2020).
- Iorio, G. & Belokurov, V. The shape of the Galactic halo with Gaia DR2 RR Lyrae. Anatomy of an ancient major merger. *Mon. Not. R. Astron. Soc.* 482, 3868–3879 (2019).
- 41. Ramos, P. et al. Full 5D characterisation of the Sagittarius stream with Gaia DR2 RR Lyrae. *Astron. Astrophys.* **638**, A104 (2020).
- Ackermann, M. et al. The search for spatial extension in high-latitude sources detected by the Fermi Large Area Telescope. *Astrophys. J. Suppl. Ser.* 237, 32 (2018).
- Abdo, A. A. et al. The second Fermi Large Area Telescope catalog of gamma-ray pulsars. Astrophys. J. Suppl. Ser. 208, 17 (2013).
- 44. Khangulyan, D., Aharonian, F. A. & Kelner, S. R. Simple analytical approximations for treatment of inverse Compton scattering of relativistic electrons in the blackbody radiation field. *Astrophys. J.* 783, 100 (2014).

Acknowledgements

R.M.C. acknowledges support from the Australian Government through the Australian Research Council under award number DP190101258 (shared with M.R.K.) and hospitality from the Virginia Institute of Technology, the Max-Planck Institut für Kernphysik and the GRAPPA Institute at the University of Amsterdam supported by the Kavli IPMU at the University of Tokyo. O.M. is supported by the GRAPPA Prize Fellowship and JSPS KAKENHI grant numbers JP17H04836, JP18H04340, JP18H04578 and JP20K14463. This work was supported by World Premier International Research Centre Initiative (WPI Initiative), MEXT, Japan. D.M. acknowledges support from the Australian Government through a Future Fellowship from the Australian Research Council, award number FT160100206. M.R.K. acknowledges support from the Australian Government through the Australian Research Council, award numbers DP190101258 (shared with R.M.C.) and FT180100375. The work of S.A. was supported by MEXT KAKENHI grant numbers JP20H05850 and JP20H05861. The work of S.H. is supported by the US Department of Energy Office of Science under award number DE-SC0020262 and NSF grant numbers AST-1908960 and PHY-1914409 and by the Japan Society for the Promotion of Science KAKENHI grant number JP22K03630. The work of D.S. is supported by the US Department of Energy Office of Science under award number DE-SC0020262. T.V. and A.R.D. acknowledge the support of the Australian Research Council's Centre of Excellence for Dark Matter Particle Physics (CDM)

CE200100008. A.J.R. acknowledges support from the Australian Government through the Australian Research Council under award number FT170100243. R.M.C. thanks E. Berkhuijsen, R. Beck, R. Ekers, M. Roth and T. Siegert for useful communications.

Author contributions

R.M.C. initiated the project and led the spectral analysis and theoretical interpretation. O.M constructed the astrophysical templates, designed the analysis pipeline and performed the data analysis of the γ -ray observations. D.M., M.R.K., S.A., S.H., M.G.B, C.G., F.A., J.A.H., D.S., and A.J.R. provided theoretical insights and interpretation and gave advice about statistical analysis. T.V. and A.R.D. provided insights on the expected distribution of dark matter. R-Z.Y. performed an initial γ -ray data analysis. M.D.F helped with radio data. The main text was written by R.M.C., M.R.K. and O.M., and the Methods was written by O.M., R.M.C. and M.R.K. All authors were involved in the interpretation of the results and reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at https://doi.org/10.1038/s41550-022-01777-x. **Supplementary information** The online version contains supplementary material available at https://doi.org/10.1038/s41550-022-01777-x.

Correspondence and requests for materials should be addressed to Roland M. Crocker or Oscar Macias.

Peer review information *Nature Astronomy* thanks the anonymous reviewers for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature Limited 2022



Extended Data Fig. 1 | The stellar density templates for the Sgr dSph used in this study. Each map has been normalized, so the units are arbitrary; the color scale is logarithmic. Morphological differences among the templates are due to different stellar candidates (red clump or RR Lyrae), search algorithms, and search target (the dwarf remnant or the stream). Data sources are as follows: Model I, ref. ⁸; Model II, ref. ³⁹; Model III, ref. ⁴⁰; Model IV and Model V, ref. ⁴¹. Detailed descriptions of these templates are given in the S.I. sec. 2.

NATURE ASTRONOMY



Extended Data Fig. 2 | Goodness of fit computation for the best-fitting baseline + Sgr dSph model. These use our preferred set of templates (first entry in Table 1). In each of the 15 panels, one for each of the energy bins in our analysis pipeline, the blue histograms show the distribution of - In L values produced in 100 Monte Carlo trials where we use our pipeline to fit a mock data set produced by drawing photons from the same set of templates used in the fit; orange dashed vertical lines show the 68% confidence range of this distribution, and black dashed vertical lines show the mean. Under the hypothesis that our best-fitting model for the real Fermi observations is a true representation of the data, and that disagreements between the model and the data are solely the result of photon counting statistics, the log-likelihood values for our best-fitting model should be drawn from the distributions shown by the blue histograms. For comparison, the red vertical line shows the actual measured log likelihoods for our best fit. The fact that these measured values are well within the range spanned by the Monte Carlo trials indicates that we cannot rule out this hypothesis, indicating that our model is as good a fit to the data as could be expected given the finite number of photons that Fermi has observed.





Extended Data Fig. 3 | Measured photon counts (left), best-fit baseline + Sgr dSph model (middle), and the fractional residuals (Data - Model)/Model (right). The images were constructed by summing the corresponding energy bins over the energy ranges displayed on top of each panel: [0.5, 1.0] GeV, [1.0, 4.0] GeV, [4.0, 15.8] GeV, from top to bottom. The maps have been smoothed with Gaussian filters of radii 1.0*, 0.8*, and 0.5* for each energy range displayed, respectively (where these angular scales are determined by the Fermi-LAT point spread function at the low-edge of the energy interval for the former two, while the latter is determined by the angular resolution of the gas maps). The spectrum of baseline + Sgr dSph model components shown here can be seen in Fig. ??. The 4FGL³⁰ γ-ray point sources included in the baseline model are represented by the red circles.

NATURE ASTRONOMY



Extended Data Fig. 4 | Results from our template mismatch tests. Each of the coloured lines shows the results of a test where we generate synthetic data with one set of templates, and attempt to recover the Sgr dSph in those data using a different set. In the upper two panels, the horizontal axis shows the true, energy-integrated Sgr dSph photon flux in the synthetic data, while the vertical axis shows the value (with 1 σ statistical error bars) retrieved by our pipeline; the black dashed lines indicate perfect recovery of the input, and the vertical bands show the photon flux we measure for the Sgr dSph in the real Fermi data. In the bottom two panels we plot the recovered energy flux in each energy bin (with 1 σ statistical error bars), for the case where the injected photon flux most closely matches the real Sgr dSph flux; the black dashed line again shows perfect recovery of the injected signal. The left panels show experiments where we mismatch the Galactic hadronic and IC templates, while the right panels show experiments where we mismatch the FB templates; see Methods for details.



Extended Data Fig. 5 | Results of our rotation and translation tests. Left: change in TS when repeating the analysis using the default baseline + Sgr dSph model, but with the Sgr dSph rotated about its centre by the indicated angle (blue points); TS values > 0 indicate an improved fit (dashed grey line), with TS = 46.1 corresponding to a 5σ -significant improvement (red dashed line). Centre: same as the left panel, but for tests with the Sgr dSph template rotated about the Milky Way centre, rather than its own centre. Right: tests for translation of the Sgr dSph template. The true position of the Sgr dSph centre is the center of the plot, and the colour in each pixel indicates the change in TS if we displace the Sgr dSph centre to the indicated position; the maximum shown, at a displacement $\Delta b \approx -4^*$, has TS = 40.8, corresponding to 4.5σ significance. For comparison, white contours show the original, unshifted Sgr dSph template, and the green arrow shows the direction anti-parallel to the Sgr dSph's proper motion, back along its past trajectory; red arrows show the projection of the green arrow in the ' and b directions.



Extended Data Fig. 6 | Sgr dSph spectra derived from template analysis using different Galactic diffuse emission models. In all cases the spectrum shown is the flux averaged over the entire ROI, not the flux within the footprint of the Sgr dSph template. The fiducial model is our default choice (first entry in Table 1), while other lines correspond to alternate foregrounds - models 2D A (red), 2D B (black), and 2D C (blue) for the Galactic IC foreground, and models Interpolated (dark green) and GALPROP 3D-gas (light green) for the Galactic hadronic + bremsstrahlung foreground. The error bars display 1 σ statistical errors. See Table 1 and text for details.



Extended Data Fig. 7 | Contribution of each template component to the γ **-ray spectrum averaged over the entire ROI, for our default baseline + Sgr dSph model.** Components shown are as follows: π 0 + brems is the Galactic hadronic plus bremsstrahlung foreground, ICS is the Galactic inverse Compton foreground, 4FGL indicates point sources from the 4th Fermi catalogue, Fermi Bubbles indicates the structured Fermi Bubble template, isotropic is the isotropic γ -ray background, 'other' includes the Sun and Moon, Loop I, and the Galactic Centre Excess, and Sgr stream indicates the Sgr dSph. The error bars display 1 σ statistical errors.