

The diffuse gamma-ray sky of a Milky Way analog: Local diversity and global constraints

Karin Kjellgren^{1,*}, Philipp Girichidis¹, Maria Werhahn², Ralf S. Klessen^{1,3}, Christoph Pfrommer⁴, Juan Soler⁵, Brian Reville⁶, Jim Hinton⁶, Patrick Hennebelle⁷, Noé Brucy⁸, and Simon C. O. Glover¹

¹ Universität Heidelberg, Zentrum für Astronomie, Institut für Theoretische Astrophysik, Albert-Ueberle-Str. 2, 69120 Heidelberg, Germany

² Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85748 Garching, Germany

³ Universität Heidelberg, Interdisziplinäres Zentrum für Wissenschaftliches Rechnen, Im Neuenheimer Feld 225, 69120 Heidelberg, Germany

⁴ Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, D-14482 Potsdam, Germany

⁵ University of Vienna, Department of Astrophysics, Türkenschanzstraße 17, 1180 Vienna, Austria

⁶ Max-Planck-Institut für Kernphysik, Saupfercheckweg 1, Heidelberg 69117, Germany

⁷ Université Paris-Saclay, Université Paris-Cité, CEA, CNRS, AIM, 91191 Gif-sur-Yvette, France

⁸ Université Lyon 1, ENS de Lyon, CNRS, CRAL, UMR 5574, Lyon, France

Received 11 January 2026 / Accepted 16 April 2026

ABSTRACT

Diffuse gamma-ray emission is a key tracer of cosmic rays (CRs) in galaxies, encoding information about their transport, energetics, and interaction with the interstellar medium. Interpreting the Milky Way’s gamma-ray sky, however, remains challenging because the observed emission depends jointly on the three-dimensional CR distribution and gas distribution, as well as the position of the observer within the Galaxy. Using the Rhea suite of CR–magnetohydrodynamic (MHD) simulations of a Milky Way analog, we investigated how pion-decay gamma-ray emission varies with galactic environment, local conditions, and CR transport physics. The emission was computed in post-processing under steady-state CR cooling and interaction assumptions, thus enabling us to analyze luminosities, spectra, full-sky emission maps, and angular power spectra (APS) for many observer positions, including those located inside Local Bubble-like cavities. The simulated galaxy naturally reproduces Milky Way-like gamma-ray luminosities and spectral slopes without any parameter tuning. While the total luminosity remains comparatively stable across the galaxy, the detailed morphology of the gamma-ray sky varies strongly with observer location due to the complex distribution of gas in the nearby environment, consistent with longstanding observational results. Across all observers, the APS closely follows the structure of the gas column density rather than the more diffuse CR energy density, in agreement with previous gamma-ray analyses and CR propagation models. Comparisons with Fermi–LAT data show good agreement for both the all-sky spectrum and the APS. A diffusion coefficient energy-scaling with power-law index $\delta = 0.5$ generally matches the observations best. Our results show that these well-established features of Galactic gamma-ray emission arise naturally in fully self-consistent CR–MHD galaxy simulations. Gas density fluctuations are the primary drivers of the morphology of the pion-decay emission, while CR transport parameters govern its spectral and structural details. The Rhea simulations thus provide a physically grounded framework for interpreting diffuse gamma-ray observations and highlight the importance of understanding the observer’s local surroundings when using gamma rays to trace Galactic CR physics.

Key words. magnetohydrodynamics (MHD) – ISM: bubbles – cosmic rays – ISM: structure – gamma rays: diffuse background

1. Introduction

Cosmic rays (CRs) are ubiquitous in galaxies; they are a crucial energy component of the interstellar medium (ISM) with energy densities that are comparable to the thermal, kinetic, and magnetic components (e.g., [Boulares & Cox 1990](#); [Grenier et al. 2015](#)). As charged particles, the CRs gyrate around and propagate in the direction of the magnetic field lines in what can be approximated as a diffusion and streaming process, allowing the particles to move relative to the thermal gas ([Zweibel 2013](#); [Ruszkowski & Pfrommer 2023](#)). This, combined with their inefficient cooling, makes them excellent sources for driving galactic outflows with a comparatively large mass loading factor, as seen in high-resolution setups of the ISM ([Girichidis et al. 2016](#); [Simpson et al. 2016](#); [Girichidis et al.](#)

[2018](#); [Rathjen et al. 2021](#); [Simpson et al. 2023](#); [Sike et al. 2025](#)) and global galaxy models ([Uhlig et al. 2012](#); [Hanasz et al. 2013](#); [Jacob et al. 2018](#); [Buck et al. 2020](#); [Peschken et al. 2021](#); [Rodríguez Montero et al. 2024](#); [Thomas et al. 2023](#); [Girichidis et al. 2024](#); [Thomas et al. 2025a](#)). CRs can also regulate star formation ([Dashyan & Dubois 2020](#)) and affect the chemistry of gas through ionization (e.g., [Padovani et al. 2020](#)), which has already been quite extensively studied in recent numerical works (e.g., [Hanasz et al. 2021](#)).

We can indirectly observe CRs via the emission that they create across a wide range of wavelengths through a variety of processes. When a CR proton interacts with thermal gas, it can create a neutral pion, which almost immediately decays into two gamma-ray photons. This is the dominant source for the diffuse gamma-ray emission in the Galaxy ([Stecker et al. 1974](#); [Selig et al. 2015](#); [Scheel-Platz et al. 2023](#)), though a non-negligible contribution (around 30% in the 0.1–100 GeV range)

* Corresponding author: karin.kjellgren@stud.uni-heidelberg.de

comes from inverse Compton (IC) and bremsstrahlung emission from CR electrons (e.g., [Strong et al. 2010](#); [Grenier et al. 2015](#); [Werhahn et al. 2021b](#)). Once generated, gamma rays below tens of TeV propagate freely through the ISM, unaffected by magnetic fields or absorption. As a result, the observed emission traces back to its site of origin, providing an indirect probe of the underlying CR distribution, which has motivated recent simulation-based modeling of galactic gamma-ray emission (e.g., [Werhahn et al. 2021b, 2023](#); [Sands et al. 2025](#)).

With instruments such as Fermi-LAT we now have observations of the all-sky diffuse gamma-ray emission in the Milky Way, as well as diffuse emission in other nearby star-forming galaxies. An observed relationship between the infrared (IR) luminosity (which probes star formation, [Kennicutt & Evans 2012](#) and references therein) and gamma-ray luminosity of such galaxies has been revealed ([Ackermann et al. 2012](#); [Rojas-Bravo & Araya 2016](#); [Ajello et al. 2020](#)), which has implications for the efficiency of CR feedback ([Pfrommer et al. 2017b](#); [Kornecki et al. 2020](#); [Werhahn et al. 2021b, 2023](#)). Highly star-forming galaxies are found to be better proton calorimeters at a level of 60–80%, meaning that the CRs lose most of their energy through hadronic interactions within the galaxy, which limits their dynamical impact¹. The Milky Way, however, is found to be only weakly calorimetric, which could be explained by efficient CR escape due to diffusion and streaming ([Thomas et al. 2025a](#)). Ground-based instruments such as LHAASO, Tibet AS_γ, and ARGO have also given us data of the diffuse gamma-ray spectra in different regions of the sky. Observations of our own galaxy in the last years have revealed several interesting features whose origins remain under debate, such as the GeV excess in the inner Galaxy ([Ackermann et al. 2017](#)) and the Fermi bubbles: two massive gamma-ray bubbles emanating from the Galactic Center both above and below the midplane ([Su et al. 2010](#)).

We can observe the diffuse gamma-ray emission from the Milky Way with greater resolution than any other system, and it is the only galaxy in which we can directly detect CR intensities and spectra. However, we can only observe our Galaxy from within it, and only from one single observer’s position, which complicates the derivation of global properties. Our observations are also shaped by our location inside the Local Bubble ([Cox & Reynolds 1987](#); [Zucker et al. 2022](#); [O’Neill et al. 2024](#)), a cavity likely created by a series of clustered supernovae (SNe). Early gamma-ray studies based on data from the satellites SAS-2 and COS-B established the correlation between the distribution of gamma-ray emission and gas structure ([Bignami et al. 1975](#); [Paul et al. 1976](#)), demonstrating that gamma-ray emission from off the Galactic plane serves as valuable tracers of local interstellar gas ([Lebrun & Paul 1978](#); [Bignami et al. 1981](#)). This conclusion has also been re-established with Fermi-LAT data, based on the dominance of the gas distribution in the 8–10 kpc ring ([Casandjian 2015](#); [Acero et al. 2016](#)). However, the role of the local environment in setting the gamma-ray sky has yet to be examined using large-scale galaxy simulations that solve the CR–magnetohydrodynamic (MHD) equations. By solving these equations, the simulations capture CR feedback on the gas and allow the resulting emission to be calculated self-consistently. This is the goal of the present work. Simulations additionally

have the advantage of allowing us to study how the sky changes as observers are put in different locations within the Galactic disk.

Our work is based on the Rhea suite of hydrodynamical simulations of isolated galaxies, aimed at reproducing key characteristics of the Milky Way ([Göller et al. 2025](#)). We used a subset of the suite that includes a magnetic field and CRs ([Kjellgren et al. 2025](#)). The gamma-ray emission was calculated during post-processing using the CRAYON+ code ([Werhahn et al. 2021a,b](#)), which calculates steady-state spectra of CR protons and electrons given the gas properties and CR energy density in every computational cell, as well as all related nonthermal emission processes; relevant for this work is the neutral pion decay calculation from CR proton spectra. This method provides a self-consistent calculation of the emission based on the physical properties of the simulated galaxy with SNe as CR sources that derive from self-consistent star formation modeling and the detailed implementation of relevant cross sections, without fine-tuning parameters and adopting source distributions to match observations.

This paper is structured as follows. In Section 2 we describe the setup of our simulations as well as our modeling of the gamma-ray emission from neutral pion decay. In Section 3 we present our results by showing the morphology of the gamma-ray emission, and how well we reproduce global estimates of the luminosity of the Milky Way. Section 4 investigates the gamma-ray sky as seen by different observers, in particular the role of the local environment on the emission. We investigate the correlation between the gamma-ray emission and the gas and CR energy distribution in Section 5, and in Section 6 we compare the simulated gamma-ray sky with observations of the Milky Way. We present our discussion in Section 7 and conclude in Section 8.

2. Methods

2.1. Simulation setup

We simulated isolated galaxies using the moving-mesh code AREPO ([Springel 2010](#); [Pakmor et al. 2016b](#); [Weinberger et al. 2020](#)). Full details of the setup are given in [Kjellgren et al. \(2025\)](#) and [Göller et al. \(2025\)](#); here we summarize the most important points.

We started with a smooth gaseous disk with a gas mass of $1.2 \times 10^{10} M_{\odot}$. To represent the gravitational influence of stars, we used an external flat gravitational potential that results in a flat velocity curve. The gas mass resolution is $3000 M_{\odot}$, and the minimum and maximum allowed cell volumes are 1 pc^3 and 2 kpc^3 , respectively. In the disk ($r_{xy} \leq 30 \text{ kpc}$, $h \leq 1 \text{ kpc}$) there is additional volume refinement that limits the maximum cell volume to $(100 \text{ pc})^3$. Stars are formed using collisionless star particles, each of which represent a stellar population. Star particles are created based on the gas mass of the cell, either deterministically based on its Jeans mass, or stochastically as in the [Springel & Hernquist \(2003\)](#) ISM model, and are populated based on a Kroupa IMF using the algorithm in [Sormani et al. \(2017\)](#). High-mass stars ($8 M_{\odot} < M_{\star} < 120 M_{\odot}$) eventually explode as SNe, injecting momentum and energy into the surrounding 100 pc . We let the simulations evolve for 2 Gyr. During the entire evolution, mass return from SNe is activated, meaning that after the SNe explode, all their mass is returned to the neighboring cells. This has the consequence that no mass is locked up in the star particles, and that we do not deplete the gas. The simulations presented in [Göller et al. \(2025\)](#) disable mass-return

¹ This is in contrast to adiabatic or streaming losses, for example, where the CR losses result in thermal and dynamical impact on the gas. We note that the degree of how much energy can have a dynamical impact also strongly depends on the coupling of CRs with the gas ([Farber et al. 2018](#); [Sike et al. 2025](#); [Thomas et al. 2025b](#)).

after the first two gigayears of evolution; however, due to the computational cost of the CR runs we only focused on this phase for this work.

The Rhea simulations use the NL97 nonequilibrium chemical network from [Glover & Clark \(2012\)](#), which tracks the abundances of atomic, ionized, and molecular hydrogen, as well as a simplified treatment of carbon and CO. These are used as inputs for the atomic and molecular heating and cooling function described in [Clark et al. \(2019\)](#), where a detailed description of the included processes can be found in [Glover et al. \(2010\)](#), with later modifications made by [Glover & Clark \(2012\)](#) and [Mackey et al. \(2019\)](#).

The magnetic field is initialized as a purely toroidal field, with an initial magnetic field strength of either 3 μG (CRMHD) or 3 nG (CRMHD-low). A detailed comparison of these two simulations can be found in [Kjellgren et al. \(2025\)](#), but for the analysis in this paper we purely focus on the CRMHD simulation. To show the robustness of our results we show some of the analysis in this paper applied to CRMHD-low in Appendix C.

Cosmic rays are included in the gray approximation, where we only evolve the total integrated CR-proton energy density ([Pakmor et al. 2016a](#); [Pfrommer et al. 2017a](#)). CR energy is injected for each SN as 10% of the explosion energy, and is deposited into the same cells as the thermal energy. They are transported in the advection-diffusion approximation with a constant diffusion coefficient of $D_0 = 4 \times 10^{28} \text{ cm}^2 \text{ s}^{-1}$ directed along the magnetic field. The following losses of CR energy are accounted for in AREPO: hadronic losses, Coulomb losses, adiabatic losses, and Alfvén losses (which emulate the losses due to streaming; see [Wiener et al. 2017](#)).

2.2. Modeling emission with CRAYON+

Gamma-ray emission from pion decay is calculated in post-processing using CRAYON+, a code that calculates steady-state CR and gamma-ray spectra in every cell of the simulation. Details of how CRAYON+ works can be found in [Werhahn et al. \(2021a,b\)](#); here we summarize the most important points and assumptions.

2.2.1. Steady-state CRp spectra

The steady-state spectrum $f(E) = d^2N/(dE dV)$ is calculated for CR protons by solving the diffusion-loss equation, which means solving

$$\frac{f(E)}{\tau_{\text{esc}}} - \frac{d}{dE}[f(E)b(E)] = q(E), \quad (1)$$

where E is the CR proton energy, and $b(E)$ and $q(E)$ are losses and sources, respectively. The escape of CRs due to diffusion and advection is included in the escape timescale $\tau_{\text{esc}}^{-1} = \tau_{\text{diff}}^{-1} + \tau_{\text{adv}}^{-1}$, where τ_{diff} and τ_{adv} are defined using the CR gradient length $L_{\text{CR}} = e_{\text{CR}}/|\nabla e_{\text{CR}}|$, where e_{CR} denotes the CR energy density, as

$$\tau_{\text{diff}} = \frac{L_{\text{CR}}^2}{D(E)}, \quad \tau_{\text{adv}} = \frac{L_{\text{CR}}}{v_z}. \quad (2)$$

For the advection timescale only the velocity in the z direction, v_z , is used because azimuthal velocities entering and leaving each cell are assumed to cancel each other out (see Fig. 6 of [Werhahn et al. 2021a](#), and surrounding discussion).

Higher-energy CRs diffuse faster (e.g., [Evoli et al. 2020](#)), thereby modifying the resulting CR spectrum as well as the subsequent gamma-ray spectrum. Although we only evolve the total

integrated CR energy in the simulation, and use a constant value for the diffusion coefficient, CRAYON+ is able to account for an energy-dependent diffusion coefficient $D(E) = D_0(E/E_0)^\delta$, where $E_0 = 3 \text{ GeV}$ and D_0 is the diffusion coefficient used in AREPO. We chose a scaling of the CR diffusion coefficient of $\delta = 0.5$ as our fiducial value, though we also tested a shallower scaling of $\delta = 0.3$. Changing δ should in theory also affect the dynamics in the simulation, but to account for that we would need to evolve the full CR spectra in an energy-dependent way, which we do not do here, but leave for future work. For the energy-dependent losses $b(E)$ hadronic and Coulomb losses are accounted for, and the source spectrum $q(E) \propto p(E)^{-\alpha_{\text{inj}}} dp/dE$ is assumed to be a power law in momentum with a spectral index of $\alpha_{\text{inj}} = 2.2$ ([Lacki & Thompson 2013](#)). After the effects of losses and escape have been applied, the resulting CR spectra are re-normalized based on the CR energy density of each cell. The Alfvén losses are implicitly included when we re-normalize the spectra but are not included as an energy-dependent loss, even though this should also be a function of CR energy. However, a detailed modeling of spectrally resolved CRs and a subsequent analysis of the relative energies reveals that the Alfvén losses at above $\sim 100 \text{ GeV}$ do not contribute much to the overall energy budget ([Girichidis et al. 2024](#)).

The code assumes steady-state, meaning that CR losses and sources balance. [Werhahn et al. \(2021a\)](#) investigated the validity of this assumption and found that the cells that predominantly contribute to the gamma-ray emission obey the conditions for steady-state. The assumption breaks down in some regions, however, in particular in the outflow region and close to the CR injection sites. Moreover, [Werhahn et al. \(2023\)](#) compared the gamma-ray emission characteristics derived from a spectral CR simulation of an isolated galaxy ([Girichidis et al. 2020, 2022, 2024](#)) to steady-state gamma-ray spectra derived with CRAYON+ and found excellent agreement after adjusting the CR diffusion coefficient. We also validate this assumption in the context of the Rhea simulations in Appendix A.

2.2.2. Gamma-ray emission

Inelastic collisions between CR protons and other gas protons produce neutral pions, which decay into gamma-ray emission. The source function $q_\gamma(E) = d^3N_\gamma/(dV dt dE)$ of the gamma-ray emission from neutral pion decay is calculated in CRAYON+ using the CR proton distribution described in the section above, and the parameterization of the cross section by [Yang et al. \(2018\)](#) from the pion production threshold ($p > 0.78 \text{ GeV}/c$) up to 10 GeV, and by [Kafexhiu et al. \(2014\)](#) for larger proton kinetic energies. This yields the gamma-ray emissivity $j_{E,\pi^0} = Eq_\gamma$, in every cell. Once again we refer to [Werhahn et al. \(2021b\)](#) for details.

We only include gamma-ray emission from the decay of neutral pions generated in collisions between CR protons and gas particles, emission in the form of IC and bremsstrahlung from primary and secondary electrons is not included. Since electrons undergo more rapid losses than the protons, it is crucial to properly model the electron population in a non-steady-state way. In particular, IC emission in the energy range 0.1–100 GeV originates from electrons with different Lorentz factors, depending on the incident radiation field. For example, for IR photons (of energy $\sim 10^{-2} \text{ eV}$), the required electron normalized momenta for IC scattering into the gamma-ray regime (0.1–100 GeV) are $\gtrsim 10^5$, which are affected by rapid losses and deviate from a steady-state ([Werhahn et al. 2025](#)). A proper implementation of the IC emission would also require careful modeling of the




Fig. 1. Face-on and edge-on views of our simulation CRMHD at $t = 1.50$ Gyr. From left to right: Column density, a slice through the midplane of the CR energy density, and the projected gamma-ray emissivity from pion decay. Almost all of the mass is located in the galactic midplane. CR-driven outflows push gas out into the CGM, but of significantly lower column densities than the disk. The CRs diffuse out of the galactic disk and fill the CGM, where they establish pressure gradients capable of launching outflows. The resulting pion-decay gamma-ray emissivity is confined close to the midplane, since the low density in the CGM reduces the amount of possible gas targets for the CRs.

radiation field. Nevertheless, at the energies considered in this work, pion decay likely dominates the gamma-ray budget (e.g., Lacki et al. 2011; Werhahn et al. 2021b). For these reasons we chose to only focus on the gamma-ray emission from pion decay and leave the other components for future work. We discuss the potential impact of the missing electrons in Section 7. For the generation of Mollweide projections of the gamma-ray sky in our simulation, as well as calculation of the APS, we made use of the healpy and HEALpix packages (Gorski et al. 2005; Zonca et al. 2019).

3. Gamma-ray emission in the Rhea simulations

Figure 1 shows face-on and edge-on projections of the gas density and the gamma-ray emissivity from pion decay, and a slice through the midplane of the CR energy density at $t = 1.50$ Gyr. The CR-driven outflows transport gas from the disk into the circumgalactic medium (CGM), which is also as a result infused with CR energy. The pion-decay gamma-ray emission is more confined to the galactic disk, being limited by the low gas densities in the outflow region. The face-on view reveals many small-scale details and local variations in the emission. Explosions from SNe drive expanding low-density bubbles, which become

outlined by an increased amount of CR energy density. These features are also imprinted in the gamma-ray emission.

To reassure ourselves that the post-processed emission is reasonable, we would like to compare our simulations with estimations of the Milky Way gamma-ray luminosity from the literature. This is done in Fig. 2, which depicts the well-known relationship between the gamma-ray luminosity and the star formation rate (SFR) (probed by the IR emission) for a selection of star-forming galaxies detected with Fermi-LAT (Ajello et al. 2020), as well as some Fermi upper limits (Rojas-Bravo & Araya 2016) that could potentially house AGNs. The luminosities (gamma-ray, IR) of the Milky Way, whose properties are what we primarily would like to reproduce, were taken from Ackermann et al. (2012) and are based on a numerical model of CR transport and ISM interaction, operated on static Galactic models (Strong et al. 2010). The SFR of the Milky Way is often quoted as between $1\text{--}2\text{ M}_\odot\text{ yr}^{-1}$ (Chomiuk & Povich 2011; Licquia & Newman 2015; Elia et al. 2022); here we used the value and uncertainty obtained in Elia et al. (2022) based on observational Herschel data from the entire Galactic disk. There are, however, estimates as low as $0.67\text{ M}_\odot\text{ yr}^{-1}$ based on observations of local high-mass stars (Quintana et al. 2025), which is closely in line with our




Fig. 2. Relation between SFR and gamma-ray luminosity in the 0.1–100 GeV energy band for our simulations. The light blue circles are LAT detections from Ajello et al. (2020) and the gray triangles are LAT upper limits from Rojas-Bravo & Araya (2016). The Milky Way gamma-ray luminosity is from Ackermann et al. (2012), while the SFR is from Elia et al. (2022). The shaded purple band shows the 1σ best fit to the LAT detections (Ajello et al. 2020). The square and diamond markers show the simulations CRMHD and CRMHD-low respectively, color-coded by the time of the snapshot. The zoomed-in region in the inset more clearly shows how the simulations change with time.

simulations, but could be underestimating the amount of star formation in the Galaxy by extrapolating from the local ISM. Nonetheless, we acknowledge that the exact value of the SFR of the Milky Way is uncertain.

Our gamma-ray luminosities are consistent with the Milky Way. This particular energy range (0.1–100 GeV) is likely dominated by the pion-decay component (Selig et al. 2015; Scheel-Platz et al. 2023) and should at most have a small contribution from CR electrons, which we do not include. We note that the temporal variations in the simulations are small, indicating that our galaxy is in dynamical equilibrium. The SFR in our simulations is on the lower end compared to the Milky Way, resulting in a degree of calorimetry (i.e., distance to the calorimetric limit, dashed line) higher than expected for the Galaxy. This could suggest that our CR energy injection efficiency of 10% is too high, i.e., that we produce too many CRs considering our SFR. Pais et al. (2018), for example, found that an average efficiency of 5% is more reasonable for SNe expanding in a turbulent magnetic field with varying orientation with respect to the blast wave. However, considering the aforementioned spread in observational values of the SFR, our data remain consistent with current estimates.

In Fig. 3 we further compare the luminosity spectrum of the full galactic disk from the CRMHD simulation with two different estimates for the Milky Way from Strong et al. (2010), which provides one of the most widely accepted models of the Galaxy’s gamma-ray luminosity (e.g., Grenier et al. 2015). The authors used the GALPROP code to calculate the luminosity spectrum from neutral pion decay (along with the other components) based on different CR propagation models, calibrated to match the direct observations of Fermi-LAT. The two models predict slightly different luminosities since the underlying transport parameters differ (the LMS model uses $\delta = 0.33$ and the LMPDS model uses $\delta = 0.5$), but the spectral shapes are similar,




Fig. 3. Estimates of the luminosity spectrum from neutral pion decay of the Milky Way from two different propagation models (dotted and dashed lines) (see Fig. 1 in Strong et al. 2010). The solid colored lines show our simulated gamma-ray spectrum from our closest matching snapshot ($t = 1.92$ Gyr). Orange depicts our fiducial CRAYON+ run with $\delta = 0.5$; the thick gray line shows $\delta = 0.3$ instead. The thin gray lines show the luminosity spectrum for every fifth snapshot ($\Delta t \approx 25$ Myr) between $t = 1.5$ Gyr and $t = 2.0$ Gyr using $\delta = 0.5$.




Fig. 4. Comparison of the vertical scale height in our simulation compared to Milky Way observations. The solid lines show the scale height of the simulation at $t = 1.92$ Gyr, calculated as 75% of the gas mass within a certain radial bin, in the cold ($T < 5050$ K) and warm ($5050 \text{ K} < T < 2 \times 10^4$ K) phase. All other data points are based on observations of the Milky Way: L06 = Levine et al. (2006), K&D08 = Kalberla & Dedes (2008), M17 = Marasco et al. (2017), B19 = Bacchini et al. (2019), and McCG23 = McClure-Griffiths et al. (2023).

in particular the slopes. The thin gray lines show the luminosity spectra for every fifth simulation snapshot ($\Delta t \approx 25$ Myr) between $t = 1.5$ Gyr and $t = 2.0$ Gyr. There are minor variations in the total luminosity with time, as was evident from Fig. 2, but the spectral shape remains constant. The solid orange and gray




Fig. 5. All-sky gamma-ray emission for different observers at the same moment in time ($t = 1.92$ Gyr). The left panel shows a slice through the midplane of the gas density with the positions of the observers marked in red. The right panels show Mollweide projections of the gamma-ray emission from pion decay in the energy range 0.56–1.0 GeV.

lines correspond to a snapshot at $t = 1.92$ Gyr using either a scaling of $\delta = 0.5$ or $\delta = 0.3$ in CRAYON+, respectively. The amplitude of the spectrum at this time matches the z04LMS-model of Strong et al. (2010) particularly well (see Fig. 3). Regardless of the GALPROP model, adopting a scaling of $\delta = 0.3$ in CRAYON+ clearly overestimates the luminosity at higher energies ($\gtrsim 1$ GeV). In contrast, the model with $\delta = 0.5$ reproduces the high-energy slope more accurately, with only a small deviation at $E \gtrsim 10^5$ MeV. We conclude that although it somewhat depends on the model of CR propagation, we manage to reproduce the luminosity spectrum of the Milky Way well for the snapshot at $t = 1.92$ Gyr with $\delta = 0.5$, which we use as our fiducial values.

We note that the more efficient production of gamma-ray emission in our simulations compared to the Milky Way, evident in Fig. 2, may partly reflect the differences between our setup and the GALPROP models of Strong et al. (2010). In particular, those models assume purely isotropic diffusion, enabling more a efficient escape of CRs than in Rhea, which adopts parallel diffusion in a predominantly toroidal magnetic field. Additional discrepancies in assumptions regarding CR injection and the gas density distribution further complicate a direct comparison, which we therefore defer to future work.

Before proceeding with the analysis, we also compare the vertical scale height in our simulated galaxy with observational estimates of the Milky Way in Fig. 4 as this is important for the resulting gamma-ray sky and the role of the local environment. The scale height is calculated in axisymmetric radial bins as the height that contains 75% of the gas mass, and is separated into the cold ($T < 5050$ K) and warm ($5050 \text{ K} < T < 2 \times 10^4$ K) phase (Kim & Ostriker 2018). Estimations of the Milky Way gas scale height from the literature are shown for comparison, pri-




Fig. 6. Histogram of the circularized bubble radii of 18 bubbles selected from one snapshot in our simulation, all from within 2.5 kpc of the solar circle. The dashed line shows the radius of the Local Bubble.

marily based on HI data, and using a variety of different methods. Though we do not see the same amount of flaring at large radii as suggested by HI observations (e.g., Kalberla & Dedes 2008), the scale height in our simulation is well in line with observational estimates.

4. Variations in the gamma-ray sky and the role of local emission

The left panel of Fig. 5 depicts a density slice through the midplane of our simulation at $t = 1.92$ Gyr, the time which provided the best match to the luminosity spectrum of the Milky




Fig. 7. Left panel: For bubble 6 in Fig. 5, distance from which up to 90% of the emission in each line of sight originates. Right panel: Corresponding gamma-ray flux with the contours of the left panel overlaid. At higher galactic latitudes the gamma-ray sky is dominated by local (≤ 2 kpc) emission.




Fig. 8. Top: Distance within which 90% of the gamma-ray emission originates, as a function of galactic latitude, for a selection of 18 bubbles in one snapshot of our simulation. The black solid line shows the median evolution, and the shaded region marks the 25th and 75th percentiles. Bottom: Sky covering fraction of regions dominated by emission from different distances, in 1 kpc bins.

Way in Fig. 3; the numbered dots represent the different locations of observers. The locations were all selected to be in low-density bubbles, mimicking the fact that the Sun is inside the Local Bubble (Zucker et al. 2022). The six panels on the right show what the gamma-ray sky (from pion decay of CR protons, in the energy range 0.56–1.0 GeV) looks like for different observers at these locations. Once gamma rays are created they do not interact, meaning it is straightforward to sum up the contributions of the individual cells to create all-sky maps. The sky projections have been rotated so they all look at the galactic center head-on, and the dynamic range and unit of the flux have been chosen to match the reconstructed diffuse dust-

correlated gamma-ray sky based on ten-year Fermi data, presented in Fig. 14 of Scheel-Platz et al. (2023).

The bubbles in the simulation were chosen based on visual inspection and not in an attempt to find the one bubble that most closely resembles the Local Bubble. The Local Bubble that we ourselves reside in is estimated to have a radius of around 165 pc, however with large variations in different directions, in particular toward the top and bottom (O’Neill et al. 2024). In Fig. 6 we show a histogram of the circularized radii of the 18 bubbles selected for the analysis in this work; all have galactocentric radii between 6.5 kpc and 10.8 kpc, so as to be comparable to the location of the Sun at 8.25 kpc. The radii of the bubbles from the simulation are all 2–4 times larger than the average radius of the Local Bubble. Although our bubbles are larger, we do not believe this compromises our analysis. For some comparisons of the gamma-ray emission in bubbles versus randomly selected positions in the disk, we refer to Appendix B.

Visually, there are striking variations between the projections. Filaments and loops dominate the sky away from the galactic midplane, resulting in widely different skies even between bubbles 2 and 3, which are relatively close to each other. This suggests that the local environment plays an important role in determining the distribution of the out-of-plane emission. To center our discussion, we focus on a single bubble for much of the remaining analysis presented in this paper; specifically, we chose bubble 6 from Fig. 5. We refer to Appendix B for some additional comparisons between the bubbles.

In order to further understand the importance of the local environment, it is useful to quantify the range of distances from which the majority of the observed emission originates. Figure 7 shows, in the left panel, the maximum distance that one needs to travel along each line of sight in order to account for 90% of the observed emission along that line of sight. For example, the orange regions in this plot denote areas of the sky for which this maximum distance is only 2 kpc. The right panel shows the flux map with the contours of the distance map overlaid. Most of the emission naturally comes from the galactic midplane, and we note that we need to include emission from distances of up to $d(0.9\text{Flux}) = 5\text{--}10$ kpc from the observer to capture 90% of the flux in these regions, and as far as 15 kpc in the center. With increasing galactic latitude there is less overall flux, and the majority of the emission originates from gas that is much closer, in many cases less than 2 kpc away. A comparison of the left and right panels shows that there is clearly visible gamma-ray emission in the flux map that overlaps with the 1–2 kpc distance-bin,




Fig. 9. APS of the gamma-ray emission, column density, and projected CR energy as a function of multipoles ℓ . The solid lines are medians of several bubbles, the shaded regions are the 25th and 75th percentiles. The saw-tooth pattern, visible in the density and gamma-ray emission, are an imprint of the galactic disk, since it is a large-scale anisotropic feature that is more dominant in the even multipoles due to its symmetry. The gamma-ray emission is mainly set by the gas density distribution.

implying that the emission is a result of the local environment where the observer is found.

Figure 8 shows how $d(0.9 \text{ Flux})$ varies with galactic latitude for a larger sample of 18 bubbles from the same snapshot, again selected based on visual inspection. We see, for example, that at a galactic longitude of 20° , 90% of the emission comes from within 2 kpc of the observer. There is not a large overall variation between different the observers; statistically, the bubbles are similar and show that with increasing latitude, local emission becomes more important. This trend is a direct imprint of the vertical scale height of the gas, as portrayed in Fig. 4, which limits the path length through dense material for lines of sight away from the midplane. Because the gamma-ray emission from pion decay traces dense structures, it is likewise confined close to the midplane (see Fig. 1, bottom right panel). The picture could conceivably change if we were to include the leptonic IC emission, which would affect the sky at high galactic latitudes and higher gamma-ray energies as it traces the hot outflowing gas (Selig et al. 2015; Scheel-Platz et al. 2023). However, at these energies, pion decay gamma rays still dominate. The bottom panel shows how much of the sky is dominated by emission from specific distance ranges. For example, over the 18 selected bubbles, the median sky fraction dominated by emission from 2–3 kpc away is $\sim 19\%$. Emission from 1–2 kpc away dominates the sky more than any other distance bin, though the exact fractions vary between bubbles. In some cases, a nonnegligible portion of the sky is dominated by emission arising from within the nearest kiloparsec.

5. Correlation with gas density and CR energy density

To zeroth order, the pion-decay gamma-ray emission is proportional to the product of the gas density and the CR energy




Fig. 10. Comparing the all-sky gamma-ray spectra in simulations and observations. The colored lines are spectra from the perspective of different bubbles in the same snapshot; the dashed orange line is bubble 6 in Fig. 5, which we selected as our best candidate. The light and dark blue data points are observational results from Selig et al. (2015) and Ackermann et al. (2017), respectively: both are based on 6.5 year Fermi-LAT data and show the total gamma-ray flux, with the point source contribution subtracted. The yellow data points are specifically the diffuse gamma-ray emission from pion-decay and bremsstrahlung, based on GALPROP modeling (Ackermann et al. 2017).

density. To investigate whether the structure seen in the emission is preferentially determined by one or the other, we computed the angular power spectrum (APS) of the gamma-ray sky as well as the gas mass and CR energy projected onto the sky for a large number of bubbles, the results of which are shown in Fig. 9. The APS tells us about the strength of fluctuations, i.e., the amount of structure, at different angular scales, represented by the multipole moment ℓ , which scales inversely with angular size. Because the three fields differ substantially in their absolute normalization and dynamic range, we standardize each individual sky map prior to computing the APS. Specifically, for a given realization i and field $X \in \{\text{gamma, gas, CR}\}$ we compute $F_{X,\text{stand}}^{(i)} = (F_X^{(i)} - \mu_{F_X^{(i)}}) / \sigma_{F_X^{(i)}}$, where $\mu_{F_X^{(i)}}$ and $\sigma_{F_X^{(i)}}$ are the mean and standard deviation of the given map. The APS $C_\ell^{(i)}$ is then computed from $F_{X,\text{stand}}^{(i)}$. The solid lines in Fig. 9 show the average APS for each of the three fields, averaged over 21 bubbles, with the shaded regions showing the 20th and 80th percentiles. The narrow width of these bands indicates that there is little statistical difference between the bubbles, and that the structure is consistent between different observer locations.

The most eye-catching feature at first glance is the saw-tooth pattern, most evident in the gas density APS. This is, however, not a numerical issue, but an imprint of the galactic disk. So much of the total mass lies in a relatively thin strip around the galactic midplane that its structure is better captured in the even multipoles, and the power in the odd multipoles is suppressed. The CRs, in contrast, diffuse out of dense regions and tend to smooth out any sharp gradients, resulting in a much smoother APS with a steep decline in angular power toward larger multipoles.




Fig. 11. Comparison of gamma-ray fluxes to observations in different regions of the sky, for bubble 6 in Fig. 5. The numbered boxes on the sky projection show which region is being counted, and the corresponding gamma-ray spectrum is shown in the smaller panels, marked by the same number. In orange is our fiducial run with $\delta = 0.5$, while the gray line shows the spectrum with $\delta = 0.3$. In the regions that are not symmetric in longitude (regions 2 and 5) the spectrum from the same region mirrored in longitude is shown. All other points and lines are from observations. The instruments used are given in the respective legends, and the data are from [Selig et al. \(2015\)](#) (region 1), [Bartoli et al. \(2015\)](#), [Amenomori et al. \(2021\)](#) (region 2), [Cao et al. \(2023\)](#), [Zhang et al. \(2023\)](#) (regions 3 and 5), and [Ackermann et al. \(2017\)](#) (region 4). We note that differences at the factor of ~ 2 level may arise from both modeling details in the simulations and from analysis choices in the observations, such as point-source removal (see, e.g., Fig. A.1 of [Zhang et al. 2023](#) for an illustration in region 5). Source confusion with diffuse emission can also affect the inferred spectral slope, particularly at low energies where confusion is more severe. We further note that, with the exception of region 4, the observations shown correspond to total gamma-ray emission rather than gas-correlated components. Regardless, we find remarkably good agreement between the simulations and observations.

The APS of the gamma-ray emission primarily follows the column density rather than the projected CR energy. As in the column density APS we see the imprint of the saw-tooth pattern associated with the midplane where most of the emission originates, although it is not as pronounced. Since the CR energy field is relatively diffuse, the structure seen in emission is set by the gas mass distribution. This is consistent with early GALPROP modeling, which found that in order to reproduce the gamma-ray emission large CR halo heights of several kiloparsec are needed ([Strong & Moskalenko 1998](#)), and that varying the size of the halo between 2 and 10 kpc had little effect on the total luminosity.

6. Comparison to observations

While we found that our luminosities are consistent with the Milky Way (see Fig. 2), it is of interest to see whether this also applies to the gamma-ray sky as seen by an observer, which we can directly compare to observed fluxes of the Galaxy. The all-sky flux spectra for 18 simulated bubbles are shown together in Fig. 10, colored by the galactocentric radius r_{xy} of the observer, together with the observed spectrum of dif-

fuse gamma rays based on 6.5-year Fermi-LAT data ([Selig et al. 2015](#); [Ackermann et al. 2017](#)). In addition to the total measured spectrum, [Ackermann et al. \(2017\)](#) decomposed the diffuse emission into individual components using GALPROP modeling, and we therefore also show their inferred contribution from pion decay and bremsstrahlung. Our simulated all-sky flux lies a factor of ~ 2 – 3 below the total observed diffuse gamma-ray emission, but is fully consistent with the expected hadronic component in both normalization and spectral slope. To capture the full observed flux we would also have to include the leptonic gamma-ray emission as well as the isotropic background flux. The normalization of the all-sky spectrum does vary slightly between different observers, which roughly correlates with galactocentric radius.

In addition to the all-sky spectrum, we computed the flux spectrum for different regions of the sky that have been studied observationally. These regions are shown and labeled in Fig. 11: $|b| < 1.5^\circ$, $|l| < 40^\circ$ (region 1) based on observations by Fermi-LAT ([Selig et al. 2015](#)); $|b| < 5^\circ$, $l = [25^\circ, 100^\circ]$ (region 2) based on observations by ARGO ([Bartoli et al. 2015](#)) and TibetAS ([Amenomori et al. 2021](#)); $|b| < 5^\circ$, $l = [15^\circ, 125^\circ]$ (region 5) and $l = [125^\circ, 235^\circ]$ (region 3) based on observations by LHAASO

(Cao et al. 2023) and Fermi-LAT (Zhang et al. 2023); and $|b| < 10^\circ$, $|l| < 10^\circ$ (region 4) based on observations by Fermi-LAT (Ackermann et al. 2017). In the case of regions 2 and 5, which are not symmetric in longitude, we additionally show the spectra from the corresponding mirrored regions, since this is an arbitrary choice in the simulation.

For most regions, our fiducial model with $\delta = 0.5$ reproduces the observed spectra well. In particular, regions 3 and 5 show a striking agreement between the simulated flux and the Fermi-LAT measurements. However, because the observational data include both hadronic and leptonic contributions, whereas our simulations only model hadronic emission, such a close agreement could indicate a mild overestimation of the hadronic flux in these regions. Region 4 provides a useful benchmark since Ackermann et al. (2017) provide estimates of the pion-decay and bremsstrahlung component, which lies a factor of $\lesssim 2$ below the total observed diffuse emission. Ideally, our simulated flux in the other regions should therefore also fall below the observed total emission by a similar amount. However, it should not be over-interpreted as a discrepancy in the regions where it does not. In regions 2 and 5 we can get a comparable variation in the normalization by simply considering the mirrored region across the galactic center. We also need to be mindful of the possibility of point source confusion in the observational data. For example, for the data shown in regions 3 and 5, the same point source masking was used by both Zhang et al. (2023) and Cao et al. (2023), based on LHAASO sources. Zhang et al. (2023) note, however, that this likely removes a nonnegligible fraction of the diffuse emission, as the flux is noticeably reduced in the inner region (region 5) compared to when subtracting point sources using the Fermi catalog.

The Galactic Center is a complicated region with puzzling gamma-ray measurements that could possibly hint at unresolved millisecond pulsars (Abazajian 2011; Abazajian & Kaplinghat 2012; Yuan & Zhang 2014) or dark matter annihilation (Huang et al. 2016; Daylan et al. 2016; Ackermann et al. 2017). In the simulations we are also missing certain features of the Galactic Center, such as a galactic bar and the Fermi bubbles. However, despite this, our simulation reproduces the observed flux from this region well. In region 1 the simulated flux falls below the Fermi-LAT measurements from Selig et al. (2015), as expected given that these data include diffuse emission from channels beyond pion decay. In region 4, which targets a similar area, the data also mostly lie below the total observed emission, but it seems we have overestimate the amount of flux near the spectral peak when we make a comparison only to the hadronic+bremsstrahlung component.

Assuming a scaling of $\delta = 0.5$ with energy for the CR diffusion coefficient changes the high-energy tail of the spectra, and generally fits the observational data better. There are regions, however, where the slope seems better captured by a weaker scaling (region 4) or where it is unclear which scaling is better. We discuss the choice of δ in Section 7.

Figure 12 compares the diffuse gamma-ray emission in our simulation with the reconstructions of Scheel-Platz et al. (2023), who used ten-year Fermi-LAT data and a Bayesian inference framework to decompose the diffuse gamma-ray emission into two components. The one shown here (bottom row) is the dust-correlated diffuse emission, i.e., the diffuse emission whose target population distribution traces the gas density. This should therefore predominantly target the hadronic component of the emission, which is what we simulate, as well as bremsstrahlung emission. It is important to note, however, that the relationship between dust emission and gas column density is not straight-

forward (e.g., Ren et al. 2025). Opacity steeply rises in denser clouds as you move into the molecular phase (Remy et al. 2017), meaning that some caution is required when equating dust-correlated emission with CR proton emission.

Quantitatively there is good agreement between the maps in all three energy bins, and the dynamic ranges are similar. There is small-scale structure evident in all maps, though it is more prominent in the observations. In the two higher-energy bins it looks as if we are missing some flux at higher latitudes and toward the galactic center, in particular in the 178–316 GeV bin. This is somewhat puzzling considering the simulated all-sky flux spectra in Fig. 10 and the GALPROP predictions (yellow data points), which show especially good agreement in the energy range 0.1–100 GeV. Because the relationship between dust opacity and column density becomes nonlinear at high densities, correlations between dust emission and column density may overestimate the gas content, which might bias the diffuse emission in the reconstruction toward higher values, in particular in the inner Galaxy. Scheel-Platz et al. (2023) also acknowledge that there is a mismatch in the point spread function (PSF) modeling, resulting in possible point-source contamination (which typically have harder spectra than diffuse emission) as well as an artificial sharpening of the emission from the disk. Together, these factors could partially explain the differences we see between our simulation and the reconstruction in the higher-energy bins.

Figure 13 compares the APS of the diffuse gamma-ray skies shown in Fig. 12, as well as the simulation with $\delta = 0.3$. The top row shows the APS applied on the full gamma-ray sky, while the bottom row shows the APS for the sky with $b > 5^\circ$, i.e., omitting the galactic midplane.

In the lowest energy bin (0.56–1 GeV), the all-sky APS shows good agreement between simulations and observations on large angular scales (low ℓ), but the two diverge at higher multipoles, where the observations exhibit a greater amount of small-scale structure. In the Scheel-Platz et al. (2023) reconstruction there is an artificial amount of power at the smallest scales due to the PSF modeling mismatch, for this reason we limit the APS to $\theta > 2^\circ$ in the lowest-energy bin and $\theta > 1^\circ$ in the two higher-energy bins. Nonetheless, our simulation is likely underestimating the small-scale power due to a combination of limited simulation resolution and simplifications in our CR modeling. Finite resolution restricts the smallest spatial scales that can be represented and can therefore suppress power at large ℓ , particularly for emission originating at nearby distances. This is illustrated in Fig. 14, which shows a histogram of cell sizes within 2 kpc of the simulated observer, averaged over 21 realizations and weighted by gamma-ray luminosity. As seen in Fig. 8, for galactic latitudes $|b| > 20^\circ$ the sky is dominated by emission from cells ≤ 2 kpc away, which have sizes of ~ 75 –100 pc (light blue histogram). As an example, a cell of transverse size $x = 75$ pc at a distance $d = 2$ kpc away, approximately corresponds to an ℓ -mode of $\ell = \pi/(x/d) \approx 84$. Considering this, it is not unreasonable that we could be suppressing power for large multipoles, which emphasizes the need for higher-resolution simulations.

In addition, with the steady-state approximation applied to the simulated gray CRs, there is little spectral variation across the sky. In reality, low-energy CRs lose their energy faster (e.g., Evoli et al. 2020) and therefore trace the regions where they are injected, while the higher-energy CRs propagate farther and produce a smoother distribution. This energy dependence could enhance small-scale structure at lower energies, an effect that is not fully captured in our modeling.




Fig. 12. Comparison of the diffuse gamma-ray emission between our simulations (upper panels) and observations (lower panels) by Fermi-LAT (Scheel-Platz et al. 2023), in three different energy bands: 0.56–1.0 GeV (left), 10–17.8 GeV (middle), and 178–316 GeV (right). The top row shows the gamma-ray sky for the same bubble as in Fig. 7. The bottom row shows the dust-correlated component of the observed diffuse gamma-ray emission.




Fig. 13. APS of the maps shown in Fig. 12. Top row: APS of the full gamma-ray sky in different energy bands. Bottom row: APS of the sky, omitting all pixels with $|b| < 5^\circ$. In pink are the observations, where we limit the APS to $\theta > 2^\circ$ in the low-energy bin and $\theta > 1^\circ$ in the two higher-energy bins, due to the PSF mismodeling artifacts present in the reconstruction. In blue and yellow are the APS from our simulated gamma-ray sky with $\delta = 0.5$ and $\delta = 0.3$ respectively.

In the two higher-energy bins the shape of the full-sky APS is well reproduced in the simulation, but the overall normalization is systematically lower than observed. For the $\delta = 0.5$ model, this discrepancy increases with energy. Using a weaker scaling of the diffusion coefficient in this case does naturally increase the intensity at higher energy, but we still do not capture the same amount of power as in the observations. By masking away the midplane, shown in the bottom row of Fig. 13,

we ignore any potentially unresolved point sources that could be present in the reconstructed sky-maps, and by doing this we achieve a close match between the observations and the $\delta = 0.3$ model. This could suggest that a weaker energy dependence of the diffusion coefficient is more appropriate at higher latitudes, and once again emphasizes the importance of accurately modeling CR transport when interpreting the details of the gamma-ray emission.




Fig. 14. Histogram of the cell sizes of all the cells within 2 kpc of an observer, averaged over 21 different bubbles and weighted by the gamma-ray luminosity. In dark blue are cells with latitudes $|b| < 20^\circ$, and in light blue cells with $|b| > 20^\circ$.

7. Discussion

7.1. CR transport details

The choice of CR transport parameters has a clear impact on the resulting gamma-ray spectrum. Our simulations assume one constant diffusion coefficient throughout the Galaxy, while our steady-state modeling of the CR proton spectrum assumes one constant value of the energy dependence δ of the diffusion coefficient. In the case of external confinement of CRs, where their scattering is dominated by externally driven turbulence, a scaling of $\delta = 0.3$ is expected for isotropic Kolmogorov turbulence (Ruszkowski & Pfrommer 2023). For this work we used $\delta = 0.5$ based on work by Evoli et al. (2020) who find a slope of $\delta = 0.54$ by fitting experimental data of secondary to primary ratios by AMS-02. In other works an even steeper slope is found (Maurin et al. 2010). Our results favor $\delta = 0.5$ rather than $\delta = 0.3$, but this should not be overinterpreted. Most modeling of the energy dependence of CR diffusion assumes a spatially constant diffusion coefficient that is only a function of particle momentum, but the diffusive properties of CRs likely change in different regions of the Galaxy. The nature of turbulence is different in the halo compared to the disk, which causes a latitudinal change in δ in the case where CRs scatter off externally driven turbulence (Tomassetti 2012).

Previous simulation work has shown how the effective diffusion coefficient (for GeV CRs) varies in different phases of the ISM if CR streaming is accounted for (Armiglotta et al. 2022; Thomas et al. 2023, 2025b). In this case the CRs excite Alfvén waves through gyroresonance, which are damped by processes such as ion-neutral or nonlinear Landau damping (Xu & Lazarian 2022; Thomas et al. 2025a). Differences in the effectiveness of the damping processes, or in which process dominates, affects the diffusion coefficient. This would also affect the resulting gamma-ray spectrum. Considering how well our results fit with the observational data, we expect this effect to be minor.

We only accounted for diffusion parallel to the magnetic field lines, when in reality a small amount of perpendicular diffusion is expected. The degree of anisotropy is not known, but is

generally estimated to be $D_\perp/D_\parallel = 10^{-4}-10^{-1}$ (Shalchi et al. 2010), and could depend on energy (Dundovic et al. 2020). Dörner et al. (2025) tested the effect of varying the degree of anisotropy on the gamma-ray sky, and found only small changes in the flux and the overall tomography of the sky.

7.2. Leptonic emission

For this work, we modeled the gamma-ray emission from neutral pion decay, and found good agreement between our simulations and the observations. However, we did not include the contribution of CR electrons to the emission, which warrants discussion. In the Milky Way, the gamma-ray emission from neutral pion decay is believed to dominate in the 0.1–100 GeV range, with leptonic processes contributing 20–40% for typical values of the electron-to-proton ratio of ~ 0.01 at GeV energies (Strong et al. 2010; Martin 2014; Werhahn et al. 2021b). Roughly 15–30% of the total luminosity is contributed by IC, and 7–20% from bremsstrahlung. IC dominates the emission at $\lesssim 10$ MeV, and is also important at the highest energies, depending on the interstellar radiation field (ISRF). Bremsstrahlung likely becomes more important at lower energies ($\lesssim 10$ GeV), though this depends on the gas density (Martin 2014). However, the exact contribution of IC emission, in particular, is highly dependent on the model of electron injection and propagation, and is limited by the uncertainties involved in the spatially varying ISRF; it is often quoted as being one of the largest sources of uncertainty in modeling of the diffuse gamma-ray emission in the Milky Way (e.g., Ackermann et al. 2015). This further justifies our choice of neglecting this component for this particular work.

Like the hadronic emission, bremsstrahlung is correlated with the gas, so its inclusion is unlikely to affect our conclusions regarding the role of local emission at higher Galactic latitudes. In contrast, IC emission might be more diffuse and will subsequently have a much smoother distribution, and is therefore more important above and below the mid-plane (Selig et al. 2015; Werhahn et al. 2021b; Scheel-Platz et al. 2023), while the density-correlated components always dominate near the Galactic plane at gigaelectronvolt energies. Though we believe our findings to be robust, it will be interesting to see how the inclusion of IC emission will impact our results.

8. Conclusions

In this work we simulated and analyzed gamma-ray emission from pion decay of simulated, isolated, Milky Way-like galaxies from the Rhea suite (Göller et al. 2025; Kjellgren et al. 2025). The simulations are run using the moving-mesh code AREPO, and evolve the total integrated CR proton energy density in a self-consistent manner. The gamma-ray emission from neutral pion decay is calculated in post-processing using CRAYON+ (Werhahn et al. 2021a,b), which accounts for the various energy-dependent losses of CR protons to calculate a steady-state spectrum in each computational cell. This is done without fine-tuning the source spectrum to match observations, and provides a self-consistent calculation of the emission based on the parameters of our galaxy simulation and a few parameter assumptions, such as the energy-dependence of the diffusion coefficient. We used this to investigate the impact of local emission on the full gamma-ray sky seen by different observers, and succeeded in reproducing several well-established observational trends of Galactic gamma-ray emission.

Our main conclusions are as follows:

1. We successfully reproduced the gamma-ray luminosity expected for the Milky Way in our simulations. The total luminosity varies moderately with time due to the time variability of the SFR, meaning that it is possible to find a simulation time that closely matches the modeled luminosity spectrum of the Milky Way. The energy scaling of the diffusion coefficient is important, with models using $\delta = 0.3$ overestimating the luminosity at higher energies (>10 GeV). A scaling of $\delta = 0.5$ reproduces the correct luminosity spectrum significantly better, showing how sensitive our results are to the details of CR transport.
2. The gamma-ray sky an observer sees can vary drastically for different positions in the galaxy, thus showing that the local environment is important. In particular, at high latitudes the local arms and spurs differ perceptibly between bubbles.
3. For a given gamma-ray sky seen by an observer, it is common for lines of sight with prominent features and/or filaments in the gamma-ray emission to come from the local ($\lesssim 2$ kpc) environment. Local emission becomes increasingly more important with increasing galactic latitude.
4. The structure of the gamma-ray emission in the sky is primarily set by the gas density distribution rather than the CR energy density distribution, which is much more diffuse.
5. We reproduced the all-sky flux spectrum from pion-decay, which varies slightly depending on which bubble the observer is placed in. The observed gamma-ray flux from different regions of the sky is consistent with observations. As is the luminosity, the emission spectrum is sensitive to the choice of CR transport parameters, and in most cases $\delta = 0.5$ is favored.

Acknowledgements. We thank the anonymous referee for a careful reading of the manuscript and the many insightful and valuable comments that helped improve the quality of the paper. The team in Heidelberg acknowledges financial support from the European Research Council via the ERC Synergy Grant “ECOGAL” (project ID 855130), from the German Excellence Strategy via the Heidelberg Cluster of Excellence (EXC 2181 – 390900948) “STRUCTURES”, and from the German Ministry for Economic Affairs and Climate Action in project “MAINN” (funding ID 50002206).00 The authors gratefully acknowledge the scientific support and HPC resources provided by the Erlangen National High Performance Computing Center (NHR@FAU) of the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU) under the NHR project a104bc. NHR funding is provided by federal and Bavarian state authorities. NHR@FAU hardware is partially funded by the German Research Foundation (DFG) – 440719683. They also thank for computing resources provided by the Ministry of Science, Research and the Arts (MWK) of the State of Baden-Württemberg through bwHPC and the German Science Foundation (DFG) through grants INST 35/1134-1 FUGG and 35/1597-1 FUGG, and for data storage at SDS@hd funded through grants INST 35/1314-1 FUGG and INST 35/1503-1 FUGG. NB acknowledges support from the ANR BRIDGES grant (ANR-23-CE31-0005). KK is a fellow of the International Max Planck Research School for Astronomy and Cosmic Physics at the University of Heidelberg (IMPRS-HD). CP acknowledges support from the European Research Council via the ERC Advanced Grant “PICOGAL” (project ID 101019746).

References

- Abazajian, K. N. 2011, *JCAP*, **2011**, 010
- Abazajian, K. N., & Kaplinghat, M. 2012, *Phys. Rev. D*, **86**, 083511
- Ajero, F., Ackermann, M., Ajello, M., et al. 2016, *ApJS*, **223**, 26
- Ackermann, M., Ajello, M., Allafort, A., et al. 2012, *ApJ*, **755**, 164
- Ackermann, M., Ajello, M., Albert, A., et al. 2015, *ApJ*, **799**, 86
- Ackermann, M., Ajello, M., Albert, A., et al. 2017, *ApJ*, **840**, 43
- Ajello, M., Di Mauro, M., Paliya, V. S., & Garrappa, S. 2020, *ApJ*, **894**, 88
- Amenomori, M., Bao, Y., Bi, X., et al. 2021, *Phys. Rev. Lett.*, **126**, 141101
- Armstrong, L., Ostriker, E. C., & Jiang, Y.-F. 2022, *ApJ*, **929**, 170
- Bacchini, C., Fraternali, F., Pezzulli, G., et al. 2019, *A&A*, **632**, A127
- Bartoli, B., Bernardini, P., Bi, X. J., et al. 2015, *ApJ*, **806**, 20
- Bignami, G. F., Fichtel, C. E., Kniffen, D. A., & Thompson, D. J. 1975, *ApJ*, **199**, 54
- Bignami, G. F., Barbareschi, L., Caraveo, P. A., et al. 1981, *Int. Cosmic Ray Conf.*, **1**, 182
- Boulares, A., & Cox, D. P. 1990, *ApJ*, **365**, 544
- Buck, T., Pfrommer, C., Pakmor, R., Grand, R. J. J., & Springel, V. 2020, *MNRAS*, **497**, 1712
- Cao, Z., Aharonian, F., An, Q., et al. 2023, *Phys. Rev. Lett.*, **131**, 151001
- Casandjian, J.-M. 2015, *ApJ*, **806**, 240
- Chiu, H.-H. S., Ruszkowski, M., Werhahn, M., & Pfrommer, C. 2024, *ApJ*, **976**, 136
- Chomiuk, L., & Povich, M. S. 2011, *AJ*, **142**, 197
- Clark, P. C., Glover, S. C. O., Ragan, S. E., & Duarte-Cabral, A. 2019, *MNRAS*, **486**, 4622
- Cox, D. P., & Reynolds, R. J. 1987, *ARA&A*, **25**, 303
- Dashyan, G., & Dubois, Y. 2020, *A&A*, **638**, A123
- Daylan, T., Finkbeiner, D. P., Hooper, D., et al. 2016, *Phys. Dark Univ.*, **12**, 1
- Dörner, J., Hellrung, J., Becker Tjus, J., & Fichtner, H. 2025, *39th International Cosmic Ray Conference*, 634
- Dundovic, A., Pezzi, O., Blasi, P., Evoli, C., & Mattheaus, W. H. 2020, *Phys. Rev. D*, **102**, 103016
- Elia, D., Molinari, S., Schisano, E., et al. 2022, *ApJ*, **941**, 162
- Evoli, C., Morlino, G., Blasi, P., & Aloisio, R. 2020, *Phys. Rev. D*, **101**, 023013
- Farber, R., Ruszkowski, M., Yang, H.-Y. K., & Zweibel, E. G. 2018, *ApJ*, **856**, 112
- Girichidis, P., Naab, T., Walch, S., et al. 2016, *ApJ*, **816**, L19
- Girichidis, P., Naab, T., Hanasz, M., & Walch, S. 2018, *MNRAS*, **479**, 3042
- Girichidis, P., Pfrommer, C., Hanasz, M., & Naab, T. 2020, *MNRAS*, **491**, 993
- Girichidis, P., Pfrommer, C., Pakmor, R., & Springel, V. 2022, *MNRAS*, **510**, 3917
- Girichidis, P., Werhahn, M., Pfrommer, C., Pakmor, R., & Springel, V. 2024, *MNRAS*, **527**, 10897
- Glover, S. C. O., & Clark, P. C. 2012, *MNRAS*, **421**, 116
- Glover, S. C. O., Federrath, C., Mac Low, M.-M., & Klessen, R. S. 2010, *MNRAS*, **404**, 2
- Göller, J., Girichidis, P., Brucy, N., et al. 2025, *A&A*, **704**, A331
- Gorski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Grénier, I. A., Black, J. H., & Strong, A. W. 2015, *ARA&A*, **53**, 199
- Hanasz, M., Lesch, H., Naab, T., et al. 2013, *ApJ*, **777**, L38
- Hanasz, M., Strong, A. W., & Girichidis, P. 2021, *Liv. Rev. Comput. Astrophys.*, **7**, 2
- Huang, X., Enßlin, T., & Selig, M. 2016, *JCAP*, **2016**, 030
- Hunter, G. H., Sormani, M. C., Beckmann, J. P., et al. 2024, *A&A*, **692**, A216
- Jacob, S., Pakmor, R., Simpson, C. M., Springel, V., & Pfrommer, C. 2018, *MNRAS*, **475**, 570
- Kafexhiu, E., Aharonian, F., Taylor, A. M., & Vila, G. S. 2014, *Phys. Rev. D*, **90**, 123014
- Kalberla, P. M. W., & Dedes, L. 2008, *A&A*, **487**, 951
- Kennicutt, R. C., & Evans, N. J. 2012, *ARA&A*, **50**, 531
- Kim, C.-G., & Ostriker, E. C. 2018, *ApJ*, **853**, 173
- Kjellgren, K., Girichidis, P., Göller, J., et al. 2025, *A&A*, **700**, A124
- Kornecki, P., Pellizza, L. J., del Palacio, S., et al. 2020, *A&A*, **641**, A147
- Lacki, B. C., & Thompson, T. A. 2013, *ApJ*, **762**, 29
- Lacki, B. C., Thompson, T. A., Quataert, E., Loeb, A., & Waxman, E. 2011, *ApJ*, **734**, 107
- Lebrun, F., & Paul, J. A. 1978, *A&A*, **65**, 187
- Levine, E. S., Blitz, L., & Heiles, C. 2006, *ApJ*, **643**, 881
- Licquia, T. C., & Newman, J. A. 2015, *ApJ*, **806**, 96
- Mackey, J., Walch, S., Seifried, D., et al. 2019, *MNRAS*, **486**, 1094
- Marasco, A., Fraternali, F., Van Der Hulst, J. M., & Oosterloo, T. 2017, *A&A*, **607**, A106
- Martin, P. 2014, *A&A*, **564**, A61
- Maurin, D., Putze, A., & Derome, L. 2010, *A&A*, **516**, A67
- McClure-Griffiths, N. M., Stanimirović, S., & Rybarczyk, D. R. 2023, *ARA&A*, **61**, 19
- O’Neill, T. J., Zucker, C., Goodman, A. A., & Edenhofer, G. 2024, *ApJ*, **973**, 136
- Padovani, M., Ivlev, A. V., Galli, D., et al. 2020, *Space Sci. Rev.*, **216**, 29
- Pais, M., Pfrommer, C., Ehlert, K., & Pakmor, R. 2018, *MNRAS*, **478**, 5278
- Pakmor, R., Pfrommer, C., Simpson, C. M., Kannan, R., & Springel, V. 2016a, *MNRAS*, **462**, 2603
- Pakmor, R., Springel, V., Bauer, A., et al. 2016b, *MNRAS*, **455**, 1134
- Paul, J., Casse, M., & Cesarsky, C. J. 1976, *ApJ*, **207**, 62
- Peschken, N., Hanasz, M., Naab, T., Wóltański, D., & Gawryszczak, A. 2021, *MNRAS*, **508**, 4269
- Pfrommer, C., Pakmor, R., Schaal, K., Simpson, C. M., & Springel, V. 2017a, *MNRAS*, **465**, 4500
- Pfrommer, C., Pakmor, R., Simpson, C. M., & Springel, V. 2017b, *ApJ*, **847**, L13

- Quintana, A. L., Wright, N. J., & Martínez García, J. 2025, [MNRAS](#), **538**, 1367
- Rathjen, T.-E., Naab, T., Girichidis, P., et al. 2021, [MNRAS](#), **504**, 1039
- Remy, Q., Grenier, I. A., Marshall, D. J., & Casandjian, J. M. 2017, [A&A](#), **601**, A78
- Ren, H. X., Remy, Q., Ravikularaman, S., et al. 2025, [A&A](#), **697**, A42
- Rodríguez Montero, F., Martín-Alvarez, S., Slyz, A., et al. 2024, [MNRAS](#), **530**, 3617
- Rojas-Bravo, C., & Araya, M. 2016, [MNRAS](#), **463**, 1068
- Ruszkowski, M., & Pfrommer, C. 2023, [A&ARv](#), **31**, 4
- Sands, I. S., Hopkins, P. F., Ponnada, S. B., et al. 2025, *Phys. Rev. D*, submitted [arXiv:2509.18351]
- Scheel-Platz, L. I., Knollmüller, J., Arras, P., et al. 2023, [A&A](#), **680**, A2
- Selig, M., Vacca, V., Oppermann, N., & Enßlin, T. A. 2015, [A&A](#), **581**, A126
- Shalchi, A., Büsching, I., Lazarian, A., & Schlickeiser, R. 2010, [ApJ](#), **725**, 2117
- Sike, B., Thomas, T., Ruszkowski, M., Pfrommer, C., & Weber, M. 2025, [ApJ](#), **987**, 204
- Simpson, C. M., Pakmor, R., Marinacci, F., et al. 2016, [ApJ](#), **827**, L29
- Simpson, C. M., Pakmor, R., Pfrommer, C., Glover, S. C. O., & Smith, R. 2023, [MNRAS](#), **520**, 4621
- Sormani, M. C., Treß, R. G., Klessen, R. S., & Glover, S. C. O. 2017, [MNRAS](#), **466**, 407
- Springel, V. 2010, [MNRAS](#), **401**, 791
- Springel, V., & Hernquist, L. 2003, [MNRAS](#), **339**, 289
- Stecker, F. W., Puget, J. L., Strong, A. W., & Bredekamp, J. H. 1974, [ApJ](#), **188**, L59
- Strong, A. W., & Moskalenko, I. V. 1998, [ApJ](#), **509**, 212
- Strong, A. W., Porter, T. A., Digel, S. W., et al. 2010, [ApJ](#), **722**, L58
- Su, M., Slatyer, T. R., & Finkbeiner, D. P. 2010, [ApJ](#), **724**, 1044
- Thomas, T., Pfrommer, C., & Pakmor, R. 2023, [MNRAS](#), **521**, 3023
- Thomas, T., Pfrommer, C., & Pakmor, R. 2025a, [A&A](#), **698**, A104
- Thomas, T., Pfrommer, C., Pakmor, R., Lemmerz, R., & Shalaby, M. 2025b, *ArXiv e-prints* [arXiv:2510.16125]
- Tomassetti, N. 2012, [ApJ](#), **752**, L13
- Uhlig, M., Pfrommer, C., Sharma, M., et al. 2012, [MNRAS](#), **423**, 2374
- Weinberger, R., Springel, V., & Pakmor, R. 2020, [ApJS](#), **248**, 32
- Werhahn, M., Pfrommer, C., Girichidis, P., Puchwein, E., & Pakmor, R. 2021a, [MNRAS](#), **505**, 3273
- Werhahn, M., Pfrommer, C., Girichidis, P., & Winner, G. 2021b, [MNRAS](#), **505**, 3295
- Werhahn, M., Girichidis, P., Pfrommer, C., & Whittingham, J. 2023, [MNRAS](#), **525**, 4437
- Werhahn, M., Pfrommer, C., Whittingham, J., et al. 2025, *MNRAS*, submitted [arXiv:2511.13811]
- Wiener, J., Pfrommer, C., & Oh, S. P. 2017, [MNRAS](#), **467**, 906
- Xu, S., & Lazarian, A. 2022, [ApJ](#), **927**, 94
- Yang, R.-Z., Kafexhiu, E., & Aharonian, F. 2018, [A&A](#), **615**, A108
- Yuan, Q., & Zhang, B. 2014, [J. High Energy Astrophys.](#), **3**, 1
- Zhang, R., Huang, X., Xu, Z.-H., Zhao, S., & Yuan, Q. 2023, [ApJ](#), **957**, 43
- Zonca, A., Singer, L., Lenz, D., et al. 2019, [J. Open Source Softw.](#), **4**, 1298
- Zucker, C., Goodman, A. A., Alves, J., et al. 2022, [Nature](#), **601**, 334
- Zweibel, E. G. 2013, [Phys. Plasmas](#), **20**, 055501




Fig. A.1. Top: face-on and edge-on slice through the midplane showing the ratio of the timescale at which the CR energy density changes τ_{CR} , to the timescale of losses and escape τ_{all} , evaluated in every computational cell. Bottom: mass-weighted histogram of the same ratio in each cell.

Appendix A: Validity of steady-state assumption

The steady-state approximation relies on the assumption that the timescale at which CR losses and escape occur, defined as $\tau_{\text{all}}^{-1} = \tau_{\text{esc}}^{-1} + \tau_{\text{cool}}^{-1}$, is shorter than the timescale over which the CR energy density changes, τ_{CR} . We evaluate this assumption in Fig. A.1 where we show in the left panel the ratio of these two timescales $\tau_{\text{CR}} / \tau_{\text{all}}$ in a face-on and edge-on slice through the center of our galaxy, and in the right panel where we show a mass-weighted histogram of each cell. The escape timescale τ_{esc} is defined as a combination of diffusion and advection, $\tau_{\text{esc}}^{-1} = \tau_{\text{diff}}^{-1} + \tau_{\text{adv}}^{-1}$, where the advection and diffusion timescales are defined in Eq. (2). The cooling timescale is defined as the total CR energy divided by the sum of the hadronic and Coulomb energy losses, $\tau_{\text{cool}} = e_{\text{CR}} / |b_{\text{had}} + b_{\text{Coul}}|$. The timescale at which the CR energy density in each cell changes is calculated as $\tau_{\text{CR}} = e_{\text{CR}} / (\Delta e_{\text{CR}} / \Delta t)$, where Δe_{CR} is change in CR energy density between two consecutive snapshots separated by Δt . Though there are a few regions where the steady-state assumption breaks down, the vast majority of cells obey that $\tau_{\text{CR}} \gtrsim \tau_{\text{all}}$, owing to fast hadronic and diffusive losses. This result agrees with the steady-state analysis using a smooth pressurized ISM (figure 9 of Werhahn et al. 2021a) and an analysis of the simulated radio-synchrotron emission in an edge-on galaxy (figure A.1 of Chiu et al. 2024), adopting the structured multi-phase ISM model CRISP (Thomas et al. 2025a), similar to the present study.




Fig. B.1. Left: average all-sky flux spectrum from 18 bubbles compared to 18 randomly selected positions at similar galactocentric radii. Shaded regions are 1σ uncertainty intervals. Right: Average flux spectrum of 18 bubbles together with the average spectra if there was a $\pm 10\%$ uncertainty in the distances to the flux-emitting cells.

Appendix B: Comparison between bubbles

For completeness's sake, and to show the variation between different observer locations, we include some analysis on other bubbles. All bubbles were selected by taking a density-slice through the midplane and picking out locations in low-density regions, with galactocentric radii r_{xy} within 2 kpc of the solar circle. Though we found that there are only minor variations between different bubbles, we can ask ourselves how much of an impact the specific bubble location has. The left panel of B.1 shows the average all-sky spectrum for a selection of 18 hand-picked bubbles versus the same number of random locations in the disk. We see that while the flux from random locations is slightly higher, the variations are minor. The right panel shows how much the flux spectrum changes if we assume the distances to the emitting cells have an uncertainty of $\pm 10\%$ (Hunter et al. 2024). Using standard error propagation, one can analytically show that this corresponds to an uncertainty of $\pm 20\%$ in the flux, in agreement with our results.

In Fig. B.2 we show the same sky projections as in Fig. 7 but for different observer positions. In all cases there is visible contribution from local emission, showing that our conclusions regarding the latitude-dependent integration-length are robust. In particular there are cases where the immediate surroundings (0–1 kpc from the observer) set the emission in the sky. Our conclusions also hold even in the case where an observer is placed in a nonbubble location, two examples of which are shown in Fig. B.3.

Appendix C: CRMHD versus CRMHD-low

For the analysis we chose to focus on one of our simulations, CRMHD. For the sake of completion and to show the robustness of our results we also show some results from our other simulation, CRMHD-low, which was initialized with 0.1% of the magnetic field strength of CRMHD, but is in other regards identical. The main difference between the two simulations is that CRMHD-low has a larger star-forming disk as a result of the weaker magnetic support, as seen in the face-on panels on the left in Fig. C.1. This is also reflected in the $d(0.9\text{Flux})$ -vs- $|b|$ plot in the right panel. Both simulations follow the same relationship, except at small galactic latitudes where the emission originates larger distances in CRMHD-low, due to the larger disk.




Fig. B.2. Same as Fig. 7, but for different bubbles. The top right, bottom left, and bottom right panels correspond to the bubbles 1, 2, and 5. The top left bubble has coordinates $x = -5$ kpc, $y = 6.4$ kpc, but is not included in Fig. 7. As before, we see that local emission dominates at higher galactic latitudes.




Fig. B.3. Same as Fig. 7, but for two random positions in the galaxy, that do not correspond to bubbles.




Fig. C.1. Left: Face-on and edge-on projections of the gamma-ray emissivity for our simulations CRMHD (left column) and CRMHD-low (right column). Right: Same as Fig. 8, but comparing CRMHD and CRMHD-low. Due to the initially weaker magnetic field, CRMHD-low has a larger disk, meaning that for small galactic latitudes most of the emission is coming from farther away than in CRMHD. At higher latitudes the two simulations behave the same.