

Deep learning models of the discrete component of the Galactic interstellar γ -ray emission

Alexander Shmakov^{1,*}, Mohammadamin Tavakoli^{1,†}, Pierre Baldi^{1,‡}, Christopher M. Karwin^{2,3,§},
Alex Broughton^{3,||} and Simona Murgia^{3,¶}

¹*Department of Computer Science, University of California, Irvine, California 92697, USA*

²*Department of Physics and Astronomy, Clemson University, Clemson, South Carolina 29634, USA*

³*Department of Physics and Astronomy, University of California, Irvine, California 92697, USA*



(Received 3 May 2022; accepted 21 December 2022; published 17 March 2023)

A significant pointlike component from the small-scale (or discrete) structure in the H_2 interstellar gas might be present in the *Fermi*-LAT data, but modeling this emission relies on observations of rare gas tracers only available in limited regions of the sky. Identifying this contribution is important to discriminate γ -ray point sources from interstellar gas, and to better characterize extended γ -ray sources. We design and train convolutional neural networks to predict this emission where observations of these rare tracers do not exist, and discuss the impact of this component on the analysis of the *Fermi*-LAT data. In particular, we evaluate prospects to exploit this methodology in the characterization of the *Fermi*-LAT Galactic center excess through accurate modeling of pointlike structures in the data to help distinguish between a pointlike or smooth nature for the excess. We show that deep learning may be effectively employed to model the γ -ray emission traced by these rare H_2 proxies within statistical significance in data-rich regions, supporting prospects to employ these methods in yet unobserved regions.

DOI: [10.1103/PhysRevD.107.063018](https://doi.org/10.1103/PhysRevD.107.063018)

The Galactic γ -ray interstellar emission (IE) traces interactions of cosmic rays with the interstellar gas and radiation field. In a companion paper [1], we showed that interstellar H_2 gas is more structured and pointlike than current IE models assume, and the related γ -ray emission might be a statistically significant component of the *Fermi*-LAT data. If this structure is not adequately captured by the IE model, it can impact the identification of resolved pointlike sources as well as the characterization of extended components in the γ -ray sky. We demonstrated that unidentified sources in the fourth *Fermi*-LAT catalog [2] could indeed be originating from it. In addition, we have argued that this component could artificially inflate the unidentified and/or unresolved point source component in the data and, depending on its morphology, contribute to confounding the interpretation of the Galactic center (GC) excess observed by *Fermi*-LAT [1,3–12] (see Ref. [13] for a review). Improved modeling of small-scale gas related emission in the *Fermi*-LAT data is therefore necessary to robustly characterize γ -ray sources.

Interstellar H_2 gas is traced indirectly via the emission lines of other molecules that are found concurrently in gas clouds. Carbon monoxide (^{12}CO , or CO hereafter) is used as a proxy; however, its line emission is typically optically thick in the denser cores of molecular clouds, and therefore it underestimates the H_2 column density there. The line emission of rarer isotopologues of CO also present in the clouds, such as ^{13}CO , remains optically thin to larger column densities and are therefore more reliable to probe dense H_2 cloud cores, and therefore the H_2 small-scale structure. We briefly summarize the methodology we developed in [1] to model this emission. We employ observations of the $J = 1-0$ transitions of CO and ^{13}CO from the Mopra Southern Galactic Plane CO Survey [14], which cover a 50 square degree region, spanning Galactic longitudes $l = 300-350^\circ$ and latitudes $|b| \leq 0.5^\circ$. As described in [1], the data reduction process by the Mopra team yields data cubes of the brightness temperature for each spectral line, as a function of Galactic coordinates and local standard of rest velocity. As described in detail in the rest of this paper, we employ these data to train machine learning (ML) models to predict ^{13}CO in regions of the sky where observations are not available. To ensure the highest quality data in the rare isotopologues, only pixels for which the brightness temperature exceeds the 1σ noise provided by Mopra are used. The performance of ML models may be biased by noise, and this conservative 1σ noise threshold allows us to retain the majority of the data for training

*ashmakov@uci.edu

†mohamadt@uci.edu

‡pfbaldi@uci.edu

§ckarwin@clemson.edu

||abrought@uci.edu

¶smurgia@uci.edu




FIG. 1. A graphical overview of our $\text{CO} \rightarrow {}^{13}\text{CO}$ modeling process for a single velocity bin. A CO column density tile of size $1^\circ \times 1^\circ = 131 \times 131$ pixels (left) is split into overlapping smaller patches (center) of size $0.053^\circ = 7$ pixels. Patches are fed to a CNN model to produce the corresponding predicted ${}^{13}\text{CO}$ column density tile (right) by estimating the ${}^{13}\text{CO}$ column density at the center of each CO patch.

purposes while still rejecting very low quality data. We calculate H_2 column densities corresponding to Mopra’s CO and ${}^{13}\text{CO}$, referred to as $N(\text{H}_2)_{\text{CO}}$ and $N(\text{H}_2)_{\text{CO}13}$, respectively. We integrate the column density to produce 17 velocity bins, corresponding to increasing radial distances from the GC. The column density computation and integration procedure is described in [1]. The final data is stored as measurements in $1^\circ \times 1^\circ \times 17$ tiles. With this as input, we build a Modified Map by replacing $N(\text{H}_2)_{\text{CO}12}$ with $N(\text{H}_2)_{\text{CO}13}$ for pixels where $N(\text{H}_2)_{\text{CO}13} > N(\text{H}_2)_{\text{CO}12}$. The “Modified Excess Template,” which corresponds to the small-scale H_2 -related γ -ray emission traced by ${}^{13}\text{CO}$, is determined from the difference between the Modified Map and the baseline CO map.

The central idea of this paper is to harness ML, in particular deep learning [15], to predict the distribution of ${}^{13}\text{CO}$ based on the CO observation, and therefore infer the H_2 small-scale structure in regions where ${}^{13}\text{CO}$ observations do not exist. Since a straightforward and robust analytical mapping between the two distributions is not available, ML would estimate this mapping from data. We train a deep learning model to map between CO and ${}^{13}\text{CO}$ column densities in the Mopra region. To simplify this regression problem and enlarge our effective dataset, we subdivide the tiled Mopra data into small $D_1 \times D_2$ patches of the sky to their respective center points (Fig. 1). The validity of this simplification requires the gas column density to be locally correlated. That is, we assume it is unlikely that column densities outside of our chosen patch will significantly change the ${}^{13}\text{CO}$ estimate. This simplification greatly reduces the model size and provides us with an effective data set of over 50 million patches when $D_1 = D_2 = 0.0581^\circ$, or 7 pixels at the Mopra resolution. We find that larger patch sizes do not improve model accuracy, justifying this locality assumption (see Ref. [16]). We also apply smoothing techniques to eliminate noise from the data (see Ref. [16]). The $\text{CO} \rightarrow {}^{13}\text{CO}$ modeling problem may now be formally written as finding a parametrized function f_θ mapping a source patch $S \in \mathbb{R}^{D_1 \times D_2}$ to an estimate target column density at the center of that

same patch $T \in \mathbb{R}$. The source S_i corresponds to Mopra CO column densities and the targets T_i to the corresponding ${}^{13}\text{CO}$ column densities at the center of the patch (Fig. 1). We evaluate our estimates by splitting the Mopra region into independent subregions which we designate for training and testing. We employ different splitting choices (see Ref. [16]), but we report on one, the *alternate tiled* split, throughout the remainder of this paper. This split consists of alternating $1^\circ \times 1^\circ$ tiles, training on longitudes $\{301\text{--}302, \dots, 349\text{--}350\}$ and testing on longitudes $\{300\text{--}301, \dots, 348\text{--}349\}$. This allows us to evaluate the model’s performance on neighboring data, reducing the variance between training and testing data distributions.

We employ convolutional neural networks (CNNs) to model and predict the ${}^{13}\text{CO}$ column density in a translation-invariant manner. CNNs use neurons with shared connection parameters in order to implement convolution operations that provide the basis for building translation invariant architectures [15]. Each convolution operation is associated with a *kernel*, or filter, corresponding to a set of connection weights that are shared by all the neurons in the corresponding layer. For each of the 17 velocity bins, we apply K learnable convolution filters of size $D_1 \times D_2$, operating on an entire patch and independently for each velocity bin, producing 17 K -dimensional latent vectors for each patch. Afterwards, we apply parametric rectified linear units [17–19] with a learnable slope, a batch normalization layer [20], and a random dropout layer [21,22]. The latent vectors are then processed by several fully-connected layers, each with their own parametric rectified linear units, batch-normalization, and dropout. Unlike the convolution layer, these hidden layers are shared between velocity bins, learning identical weights for every bin. This design allows spatial components (CNN) of the network to be specialized for each velocity bin while allowing latent higher level components to be shared between bins. The resulting latent vectors are fed through a final fully connected layer which produces the ${}^{13}\text{CO}$ concentration point estimates for each patch and velocity bin. A diagram of the network architecture is presented in Fig. 2.

The Mopra dataset contains widely varying column densities from bright to very dim. To effectively learn this high-spread distribution, we model the incidence of photons on the Mopra detector as a Poisson process and use a Poisson log-likelihood loss. We further reweight this loss based on the ${}^{13}\text{CO}$ density to elevate the importance of bright pixels, prioritizing accuracy in hot spots over a slight degradation in the background. This increases the importance of accurate measurement in the bright regions, encouraging the network to focus on the accuracy of these regions first. As the network trains, we anneal this weighting back to uniform (with all target values having equal weighting) in order to minimize any bias introduced by this loss. The design of this loss is guided by the overall goal of finding small angular scale features while limiting




FIG. 2. A diagram of the internal layers of the network. We maintain independent weights for the 17 convolution layers, one for each velocity bin. However, the hidden and output fully connected layers share a single set of weights and the same layers are applied to every velocity bin.

the amount of overestimation throughout the Mopra region. A formal description of the loss function is available in [16].

We tune the CNN’s hyperparameters using the Sherpa hyperparameter optimization library [23], testing 2000 network variations, scoring each parametrization with Poisson likelihood, and using Gaussian process optimization to suggest parameters. Our final parametrization is presented in [16], which is evaluated after training the network for 200 iterations. We generate the predicted ^{13}CO density by stitching the network output for every patch in a source image. When performed on a single Nvidia Titan XP GPU, and batching 8192 patches at once, this inference process requires approximately 11 seconds to cover the entire Mopra region.

Following the same procedure as in [1], we employ the cosmic ray propagation code GALPROP (v56) [24–34] to calculate γ -ray sky maps in 17 galactocentric radial bins for the H_2 -related emission, where the latter is traced by CO and ^{13}CO . The CO and ^{13}CO related H_2 column densities, $N(\text{H}_2)_{\text{CO13}}$ and $N(\text{H}_2)_{\text{CO}}$, respectively, are used to determine the H_2 that is missed in dense regions when

only CO is used as a tracer. We follow the procedure from [1] to determine the Modified Map and “Modified Excess Template” (Fig. 3, top panel) to construct their ML analogs. The “Smart Map” is determined similarly to the Modified Map, but instead of using the true Mopra ^{13}CO data, we use the ^{13}CO estimates from the CNN. We use this to determine the “Smart Excess Template,” defined as the difference between the Smart Map and the baseline CO map. The Smart Excess Template, also integrated over all annuli and energies, is shown in Fig. 3, bottom panel. The map covers the full Mopra region, which includes alternating training and testing tiles for the CNN.

To assess how closely the CNN predicts the γ -ray emission inferred by Mopra observations, we compare the excess templates with three different metrics. We quantify the differences with the following three metrics which are determined by using the predicted photon counts per pixel for each of the templates and scaled to ~ 12 years of *Fermi*-LAT data (we assume the same *Fermi*-LAT observation parameters and event selection as for the simulations described in [1] and in this work): absolute difference (Smart Modified), fractional difference $[(\text{Smart}-\text{Modified})/\text{Modified}]$, difference in units of the standard deviation, σ , for the Modified counts $[(\text{Smart}-\text{Modified})/\sqrt{\text{Modified}}]$. The latter metric allows us to compare the magnitude of the differences due to the CNN performance in light of the statistical power of the data. All results are shown in Fig. 4 as a function of longitude and latitude for the Mopra region (see [16] for distributions of residuals integrated over the region). We observe that for the vast majority of the region (83.6% of the pixels), the predicted counts for the Smart Excess Template are within $\pm 1\sigma$ of the Modified Excess Template counts, and therefore the difference is generally within the statistical uncertainty of the data. This result indicates that the CNN performance is adequate for modeling the small-scale structure in H_2 -related γ -ray emission traced by Mopra, for the statistics achieved by *Fermi*-LAT.

The CNN performance varies with respect to longitude. Most tiles display either an overall underprediction by the CNN across the entire tile or an overprediction, rather than




FIG. 3. *Top*: the Modified Excess Template (from Ref. [1]). *Bottom*: the Smart Excess Template. The color scale (logarithmic) indicates the γ -ray intensity per 0.03125^2 degree pixel. The first tile to the left (closest to the GC, and covering longitudes 349–350) is testing data, and the neighboring one is training data for the CNN. The interleaving pattern of testing and training tiles is adopted for the entire Mopra region.




FIG. 4. The pixel-wise residuals for the full region between the predicted counts from the Smart Excess Template and the predicted counts from the Modified Excess Template. *Top*: residual in counts (Smart Modified), *Middle*: fractional residuals $[(\text{Smart}-\text{Modified})/\text{Modified}]$, *Bottom*: residual in units of 1 standard deviation for the modified excess template $[(\text{Smart}-\text{Modified})/\sqrt{\text{Modified}}]$. The color scale is per 0.08^2 degree pixel.

a comparable mixture of under/over predictions across the same tile. This may be explained by the design of the loss function, placing increased importance towards brighter pixels while also biasing the network to prefer under-prediction via Poisson regression. Prominent examples include the training tile $343^\circ < l < 344^\circ$ and testing tile $348^\circ < l < 349^\circ$, the brightest in each dataset, where the CNN underpredicts the gas column density, as shown in Fig. 4. The dependence of the residuals as a function of pixel brightness (in counts/pixel), is shown in Fig. 5 for the full Mopra region, and for the training and testing subregions separately. Generally, better agreement between Mopra (Modified Excess Template) and the CNN prediction (Smart Excess Template) is found for the brighter pixels. There is a broad spread in the fractional residuals, but it is confined to be between $\pm 50\%$ for the vast majority of the pixels. We overlay contours to indicate the $\pm 1\sigma$ statistical fluctuations in the Modified Excess Template, which encompasses 83.6% of the pixels. We also overlay

the median and median absolute deviation of the fractional difference in bins of the Modified Excess Template counts to guide the eye, but emphasize that the scatter of the points within each interval cannot be simply characterized by these quantities and does not always indicate a most probable outcome. As expected, we find that the CNN performs (somewhat) better in the training tiles. Approximately 58% of pixels where the CNN prediction is beyond the $\pm 1\sigma$ level are in the testing subregion. The CNN is more likely to overpredict the emission for the dimmer pixels (60.8% of pixels across the full region are below the 25% flux percentile), consistently with the design. The CNN overprediction in the dim pixels essentially spreads out the brightness of hot spots over a larger area. We also note that small statistics causes the fractional difference to increase dramatically for the faintest pixels. Above the 90% flux percentile, the distribution of the fractional residuals bifurcates into two separate distributions. The underpredicted pixels are in the $343^\circ < l < 344^\circ$ tile,




FIG. 5. The fractional residuals between the nominal Smart Excess Template (*Smart*) and Modified Excess Template (*Modified*) predicted counts per pixel as a function of Modified Excess Template pixel counts for all tiles (*left*), testing tiles (*center*), and training tiles (*right*). The points and error bars represent the median and median absolute deviation of the fractional difference for the pixels in 8 domain intervals (we note that the scatter of the points within each interval is not simply characterized by these quantities and does not always indicate a most probable outcome.) We also plot the $\pm 1\sigma$ boundaries of the Poisson noise for the Modified Excess Template. We overlay the 25%, 50%, 95% flux percentiles to indicate the flux fraction encompassed by pixels below that threshold.




FIG. 6. Statistical significance of the Smart Excess Template for the full Mopra region (solid lines), and for the testing subregion (dashed lines). Each contains 1000 realizations of 12 years of *Fermi*-LAT data. A fit with a Gaussian distribution is overlaid to each distribution. These distributions are compared to the results for the Modified Excess Template from [1].

and the overpredicted pixels are in $333^\circ < l < 334^\circ$ tile, both in the training subregion. This is likely caused by our choice to combine independent convolution layers for each velocity bin with a shared hidden layer. Since certain bins have higher overall brightness than others, the shared hidden layer will attempt to average the error between the two extremes.

We determine the significance of the Smart Excess Template in the *Fermi*-LAT data, similarly to [1] for the Modified Excess Template. The simulations cover the same observations and event selection (~ 12 years, 1–100 GeV, P8R3 CLEAN FRONT + BACK). We only simulate the H_2 -related γ -ray emission in the Mopra region, excluding all other components, since our goal is to establish the performance of the Smart Excess Template in the optimistic

scenario where all other components are known. The simulated events trace the H_2 -related γ -ray emission modeled with the Modified Map from Mopra. The simulated data are fit based on a binned maximum likelihood method to a model that includes the baseline CO map (as observed by Mopra) and the Smart Excess Template. The normalization of the Smart Excess Template is free to vary in the fit. The energy spectrum, calculated by GALPROP, is assumed, and held fixed during the likelihood fit. The normalization and spectral index of the CO baseline contribution is also free to vary. The 17 radial bins from GALPROP are combined into 4 radial annuli, which we refer to as A1, A2, A3, and A4, with the same partitioning as in [1]. Also consistently to [1], the normalization of A4 is held constant to a normalization of 1.0, contributing a flux of $4.98 \times 10^{-9} \text{ ph cm}^{-2} \text{ s}^{-1}$.

We simulate 1000 realizations and calculate the test statistics (TS) for the nested models $[-2 \log(L_0/L)]$, where L_0 is the null hypothesis (CO baseline), and L is the alternative hypothesis (CO baseline and Smart Excess Template.) The statistical significance is approximated by $\sigma \approx \sqrt{TS}$. The distribution of the $\sqrt{TS}$ for the simulations is shown in Fig. 6 for the full Mopra region (solid line), and for the testing subregion only (dashed line). The Smart Excess Template corresponds to $\sqrt{TS} = 27.5 \pm 1.1$ (mean and standard deviation) in the full Mopra region and 12.9 ± 1.0 in the testing subregion. The significance in the testing subregion is lower, not only because of the smaller statistics but also because the full region contains the training tiles where the CNN more closely matches the Mopra data. For comparison, the Modified Excess Template has a $\sqrt{TS} = 48.3 \pm 1.0$ and $\sqrt{TS} = 31.5 \pm 1.0$ for the full region and testing subregion, respectively. In Fig. 7, we show the distributions for the best-fit flux of the Smart Excess Template overlaid with those for the Modified Excess Template, as well as the flux for the CO baseline emission, separated into annuli. The integrated flux for the Smart Excess Template is approximately 83.7%




FIG. 7. Distribution of fluxes for each component of the best-fit model using the Smart Excess Template maps over the full region for 1000 simulations of *Fermi*-LAT data. The results for the Modified Excess Template maps from [1] are overlaid. Note that one of the annuli (A4), not shown, has negligible contribution and the normalization was fixed in the fitting process.




FIG. 8. Fractional count residuals for the fit with the Smart Excess Template. The green and yellow bands show the 68% and 95% confidence regions from 1000 simulations, respectively. We also plot the results for a single simulation as an example, which is shown with black data points.

of the flux of the Modified Excess Template, indicating an overall underprediction by the CNN. The fractional count residuals $[(\text{data}-\text{model})/\text{model}]$ as a function of energy for the best fit model are shown in Fig. 8 for the fits that include Smart Excess Template, in addition to the other CO components. They are consistent with zero. The best fit spectra of the CO components agree with the GALPROP prediction. Finally, Fig. 9 shows the fractional count residuals in latitude, longitude. The residuals, which incorporate differences between the (best fit) Smart Excess Template and the (simulated) Modified Excess Template, are smaller than those shown in Fig. 4 (middle panel), indicating that the other CO model components partially compensate for the discrepancies between the excess templates. We have performed this analysis also using Gaussian processes and find that the performance is worse compared to the CNN [16].

The ML methodology presented here must be refined to be extended to other regions of the sky. The available all sky CO $J = 1-0$ observations have significantly poorer spatial resolution ($8'$ from [35]) compared to Mopra ($0.6'$). The CNN model will therefore require additional transfer learning to operate on lower resolution maps. The resolution of the *Fermi*-LAT data is worse (for most energies and event types), suggesting that the impact of the poorer CO resolution on the CNN predictions for the γ -rays might ultimately be less pronounced. Another consideration is that the available ^{13}CO training data for the CNN (from

Mopra, and other observations [36–40] that could also be included), are confined to the Galactic plane. The H_2 scale height is low, but its small-scale structure contribution at higher latitudes might not be negligible because of the contribution of more local H_2 . However, because of the lack of adequate training data the CNN prediction for this component could be more uncertain and this is relevant for the characterization of the GC excess, which extends to high latitudes. Finally, this analysis inherits the limitations and uncertainties in modeling the H_2 component with traditional methods, including the loss of kinematic resolution for the gas in the direction of the GC [41]. And it adds another: the H_2 column densities from CO and ^{13}CO have been treated independently in this work. Their estimates can be related analytically, but rely on assumptions pertaining to the optical depth, beam filling factor, spatial variation, etc. (e.g., [42,43]). More ambitiously, ML may be used to constrain some of these uncertainties by using the γ -ray data. Finally, small-scale structure in the γ -ray data could also arise from other components of the IE, e.g., related to HI , and shall be included in a more comprehensive study.

We have presented a methodology that harnesses ML to predict the small-scale component of the interstellar H_2 gas and its related γ -ray emission for the first time. Our work is motivated by the fact that while CO is widely employed as a tracer of H_2 , its emission alone does not reliably capture the more highly structured component of H_2 . Incorporating this small-scale component in γ -ray IE models is crucial for determining γ -ray point sources and could affect the characterization of extended γ -ray sources, e.g., the *Fermi*-LAT GC excess. Here, we have improved H_2 modeling by employing observations from the Mopra survey which includes isotopologues of CO such as ^{13}CO that do not suffer from the same limitations and can, therefore, more reliably trace the H_2 small-scale structure (the survey spans Galactic longitudes $l = 300\text{--}350^\circ$ and latitudes $|b| \leq 0.5^\circ$). In particular, we designed and trained a CNN to predict the concentration of ^{13}CO in the Mopra region. We have tested the performance of this methodology in predicting the contribution of the H_2 related γ -ray small-scale structure to the *Fermi*-LAT data. We find that the accuracy of the predicted γ -ray emission, which incorporates the CNN-predicted contribution, is within the statistical uncertainty of the *Fermi*-LAT data for the vast majority of the region. Moreover, the predicted γ -ray contribution can have a large statistical significance, $\sqrt{TS} = 28$ if we assume other components of the γ -ray IE are known. We designed a




FIG. 9. Fractional count residuals $[(\text{data}-\text{model})/\text{model}]$ in latitude and longitude for the fits with the Smart Excess Template, in addition to the other CO components. The residuals are calculated using the mean over 1000 simulations.

dynamically weighted loss function for the CNN's training to prioritize emission accuracy from brighter spots in the Mopra data while avoiding overpredicting the ^{13}CO concentration in these regions. This is a conservative choice aimed at minimizing the injection of spurious γ -ray sources in the IE model. As a consequence of this design, overprediction is more likely for the dimmer spots, which we show to be less important. Based on these results, we conclude that deep learning can improve the modeling of γ -ray emissions in data-rich regions, supporting prospects to extend this methodology to other regions of the sky. In this work, we restrict our analysis to the Mopra observations of CO and ^{13}CO which are confined to the Galactic plane.

Extending this methodology to other regions will require the CNN model to model ^{13}CO from the poorer resolution CO maps available for the majority of the sky, and will require extrapolating predictions to higher latitudes.

We thank Troy Porter for many helpful discussions and insights. The work of A. S., M. T., and P. B. is in part supported by Grants No. NSF NRT 1633631 and No. ARO 76649-CS to P. B. The work of A. B. was supported in part by the UCI Physical Sciences Machine Learning NEXUS Fellowship. This work was performed in part at the Aspen Center for Physics, which is supported by National Science Foundation Grant No. PHY-1607611.

-
- [1] C. Karwin, A. Broughton, S. Murgia, A. Shmakov, M. Tavakoli, and P. Baldi, [arXiv:2206.02809](#).
 - [2] J. Ballet, T. H. Burnett, S. W. Digel, and B. Lott, [arXiv:2005.11208](#).
 - [3] L. Goodenough and D. Hooper, [arXiv:0910.2998](#).
 - [4] D. Hooper and L. Goodenough, *Phys. Lett. B* **697**, 412 (2011).
 - [5] D. Hooper and T. Linden, *Phys. Rev. D* **84**, 123005 (2011).
 - [6] K. N. Abazajian, *J. Cosmol. Astropart. Phys.* **03** (2011) 010.
 - [7] K. N. Abazajian and M. Kaplinghat, *Phys. Rev. D* **86**, 083511 (2012).
 - [8] K. Abazajian, N. Canac, S. Horiuchi, and M. Kaplinghat, *Phys. Rev. D* **90**, 023526 (2014).
 - [9] C. Gordon and O. Macías, *Phys. Rev. D* **88**, 083521 (2013).
 - [10] T. Daylan, D. P. Finkbeiner, D. Hooper, T. Linden, S. K. Portillo, N. L. Rodd, and T. R. Slatyer, *Phys. Dark Universe* **12**, 1 (2016).
 - [11] F. Calore, I. Cholis, and C. Weniger, *J. Cosmol. Astropart. Phys.* **03** (2015) 038.
 - [12] M. Ajello, A. Albert, W. Atwood, G. Barbiellini, D. Bastieri, K. Bechtol, R. Bellazzini, E. Bissaldi, R. Blandford, E. Bloom *et al.*, *Astrophys. J.* **819**, 44 (2016).
 - [13] S. Murgia, *Annu. Rev. Nucl. Part. Sci.* **70**, 455 (2020).
 - [14] C. Braiding, G. F. Wong, N. I. Maxted, D. Romano, M. G. Burton, R. Blackwell, M. D. Filipović, M. S. R. Freeman, B. Indermuehle, J. Lau *et al.*, *Pub. Astron. Soc. Aust.* **35**, E029 (2018).
 - [15] P. Baldi, *Deep Learning in Science* (Cambridge University Press, Cambridge, United Kingdom, 2021).
 - [16] See the Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevD.107.063018> for details on the data cleaning of the Mopra maps, choice of the training and testing regions, description of the Gaussian processes approach, and additional comparisons between Smart and Modified Maps.
 - [17] F. Agostinelli, M. Hoffman, P. Sadowski, and P. Baldi, [arXiv:1412.6830](#).
 - [18] K. He, X. Zhang, S. Ren, and J. Sun, [arXiv:1502.01852](#).
 - [19] M. Tavakoli, F. Agostinelli, and P. Baldi, *Neural Netw.* **140**, 1 (2021).
 - [20] S. Ioffe and C. Szegedy, in *Proceedings of the 32nd International Conference on International Conference on Machine Learning—Volume 37*, ICML'15 (JMLR.org, 2015), pp. 448–456, [10.5555/3045118.3045167](#).
 - [21] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, *J. Mach. Learn. Res.* **15**, 1929 (2014).
 - [22] P. Baldi and P. Sadowski, *Artif. Intell.* **210C**, 78 (2014).
 - [23] L. Hertel, J. Collado, P. Sadowski, J. Ott, and P. Baldi, [arXiv:2005.04048](#).
 - [24] I. V. Moskalenko and A. W. Strong, *Astrophys. J.* **493**, 694 (1998).
 - [25] I. V. Moskalenko and A. W. Strong, *Astrophys. J.* **528**, 357 (2000).
 - [26] A. W. Strong and I. V. Moskalenko, *Astrophys. J.* **509**, 212 (1998).
 - [27] A. W. Strong, I. V. Moskalenko, and O. Reimer, *Astrophys. J.* **537**, 763 (2000); **541**, 1109(E) (2000).
 - [28] V. S. Ptuskin, I. V. Moskalenko, F. C. Jones, A. W. Strong, and V. N. Zirakashvili, *Astrophys. J.* **642**, 902 (2006).
 - [29] A. W. Strong, I. V. Moskalenko, and V. S. Ptuskin, *Annu. Rev. Nucl. Part. Sci.* **57**, 285 (2007).
 - [30] A. E. Vladimirov, S. W. Digel, G. Jóhannesson, P. F. Michelson, I. V. Moskalenko, P. L. Nolan, E. Orlando, T. A. Porter, and A. W. Strong, *Comput. Phys. Commun.* **182**, 1156 (2011).
 - [31] G. Jóhannesson *et al.*, *Astrophys. J.* **824**, 16 (2016).
 - [32] T. A. Porter, G. Jóhannesson, and I. V. Moskalenko, *Astrophys. J.* **846**, 67 (2017).
 - [33] G. Jóhannesson, T. A. Porter, and I. V. Moskalenko, *Astrophys. J.* **856**, 45 (2018).
 - [34] Y. Génolini, D. Maurin, I. V. Moskalenko, and M. Unger, *Phys. Rev. C* **98**, 034611 (2018).
 - [35] T. M. Dame, D. Hartmann, and P. Thaddeus, *Astrophys. J.* **547**, 792 (2001).
 - [36] J. S. Urquhart, C. König, A. Giannetti, S. Leurini, T. J. T. Moore, D. J. Eden, T. Pillai, M. A. Thompson, C. Braiding, M. G. Burton, T. Csengeri, J. T. Dempsey, C. Figura, D. Froebrich, K. M. Menten, F. Schuller, M. D. Smith, and F. Wyrowski, *Mon. Not. R. Astron. Soc.* **473**, 1059 (2018).

- [37] D.J. Eden *et al.*, *Mon. Not. R. Astron. Soc.* **498**, 5936 (2020).
- [38] P. M. W. Kalberla, J. Kerp, and U. Haud, *Astron. Astrophys.* **639**, A26 (2020).
- [39] R. Blackwell, M. Burton, and G. Rowell, in *The Multi-Messenger Astrophysics of the Galactic Centre*, Vol. 322, edited by R. M. Crocker, S. N. Longmore, and G. V. Bicknell (Cambridge University Press, Cambridge, 2017), pp. 164–165.
- [40] E. M. Di Teodoro, N. M. McClure-Griffiths, F. J. Lockman, S. R. Denbo, R. Endsley, H. A. Ford, and K. Harrington, *Astrophys. J.* **855**, 33 (2018).
- [41] M. Ackermann, M. Ajello, W. Atwood, L. Baldini, J. Ballet, G. Barbiellini, D. Bastieri, K. Bechtol, R. Bellazzini, B. Berenji *et al.*, *Astrophys. J.* **750**, 3 (2012).
- [42] M. G. Burton, C. Braiding, C. Glueck, P. Goldsmith, J. Hawkes, D. J. Hollenbach, C. Kulesa, C. L. Martin, J. L. Pineda, G. Rowell *et al.*, *Pub. Astron. Soc. Aust.* **30**, E044 (2013).
- [43] D. Cormier, F. Bigiel, M. J. Jiménez-Donaire, A. K. Leroy, M. Gallagher, A. Usero, K. Sandstrom, A. Bolatto, A. Hughes, C. Kramer, M. R. Krumholz, D. S. Meier, E. J. Murphy, J. Pety, E. Rosolowsky, E. Schinnerer, A. Schrubba, K. Sliwa, and F. Walter, *Mon. Not. R. Astron. Soc.* **475**, 3909 (2018).