https://doi.org/10.3847/1538-4357/abc47c



Inpainting Galactic Foreground Intensity and Polarization Maps Using Convolutional Neural Networks

Giuseppe Puglisi^{1,2,3} and Xiran Bai^{1,2,4}

¹ Department of Physics, Stanford University, Stanford, CA 94305, USA; gpuglisi@lbl.gov

² Kavli Institute for Particle Astrophysics and Cosmology, SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA ⁴ Department of Physics, Yale University, New Haven, CT 06520, USA

Received 2020 March 30; revised 2020 August 17; accepted 2020 October 12; published 2020 December 22

Abstract

The Deep Convolutional Neural Networks (DCNNs) have been a popular tool for image generation and restoration. In this work, we applied DCNNs to the problem of inpainting non-Gaussian astrophysical signal, in the context of Galactic diffuse emissions at the millimetric and submillimetric regimes, specifically Synchrotron and Thermal Dust emissions. Both signals are affected by contamination at small angular scales due to extragalactic radio sources (the former) and dusty star-forming galaxies (the latter). We compare the performance of the standard diffusive inpainting with that of two novel methodologies relying on DCNNs, namely Generative Adversarial Networks and Deep-Prior. We show that the methods based on the DCNNs are able to reproduce the statistical properties of the ground-truth signal more consistently with a higher confidence level. The Python Inpainter for Cosmological and AStrophysical SOurces (PICASSO) is a package encoding a suite of inpainting methods described in this work and has been made publicly available at http://giuspugl.github.io/picasso/.

Unified Astronomy Thesaurus concepts: Cosmic microwave background radiation (322); Extragalactic radio sources (508); Convolutional neural networks (1938); Interstellar synchrotron emission (856); Dust continuum emission (412)

1. Introduction

Over the last few years, the use of machine-learning techniques has become increasingly popular in analyzing scientific data. In particular, the use of the Deep Convolutional Neural Networks (DCNNs) has opened a wide range of interesting applications (Mustafa et al. 2017; Rodríguez et al. 2018; Caldeira et al. 2019; Krachmalnicoff & Tomasi 2019; Perraudin et al. 2019; Aylor et al. 2020; Farsian et al. 2020).

In this work, we investigate how the problem of estimating and reconstructing missing or masked regions of observations can be better solved using DCNNs. These techniques have been widely used for image restoration and face completion and in principle can be similarly applied to generate semantic content for astrophysical signals.

In particular, we focus on the case of the reconstructing polarized signal emitted in the radio and submillimeter regimes: (i) at $\nu \lesssim 60$ GHz, where the emission is mostly dominated by Galactic synchrotron, described by a power law $\beta_{\text{synch}} \sim -3$ (Krachmalnicoff et al. 2018), (ii) at $\nu \gtrsim 150$ GHz where most of the polarization is due to the thermal Galactic dust grains aligning with the Galactic magnetic field, described by a modified blackbody law (Planck Collaboration et al. 2020). Emission coming from molecules (like carbon monoxide) and from anomalous microwave emission are expected to be polarized at this regime of frequencies, although their degree of polarization is expected to be lower (a few percent, for details, see Puglisi et al. 2017; Dickinson et al. 2018) compared to that of dust and synchrotron (about 10%).

At 80 $< \nu < 110$ GHz, the cosmic microwave background (CMB) polarization has a nonnegligible contribution especially at high Galactic latitudes. A reliable assessment of both synchrotron and dust polarized emissions in the two regimes (i) and (ii) is critical to separate the Galactic contamination in CMB

measurements and further detect the divergence-less pattern in the CMB polarization called B-mode. CMB B-modes, at degree angular scales, are directly related to the imprint of a stochastic background of gravitational waves produced during the inflationary phase of our universe, commonly referred as tensorial anisotropies. To date, primordial B-modes have not yet been detected and the latest upper limits have been provided by Sayre et al. (2020), Adachi et al. (2019), and BICEP2 Collaboration et al. (2018). Future experiments aim at better characterizing diffuse polarized emission from our own Galaxy with high-sensitivity measurements (Carlstrom et al. 2019; The Simons Observatory Collaboration et al. 2019).

At the arcminute angular scales, B-modes are sourced by the gravitational lensing of large-scale structures, which deflect the CMB scalar polarization anisotropies into the so-called lensing B-modes (see the latest constraints in Sayre et al. 2020; POLARBEAR Collaboration et al. 2017; Louis et al. 2017). At these scales, extragalactic radio sources and star-forming galaxies are the major polarized contaminants. The majority of these contaminants mostly appears as bright and unresolved point sources in a typical CMB map (the latest measurements can be found in Gupta et al. 2019; Datta et al. 2019). Puglisi et al. (2018) have shown that hundreds of polarized sources will be detected by the forthcoming experiments given the expected nominal sensitivity and the observation sky fraction ($\sim 10\%$ -30%). Hence an aggressive masking may be applied on maps surveyed by the forthcoming CMB experiments, preventing a high-resolution Galactic foreground template as well as a reliable analysis involving high-order estimators beyond the two-point correlation function. Reconstructing signals in the masking area to fill the missing data is done to ameliorate these issues, a procedure sometimes referred to as inpainting (used in, e.g., Starck et al. (2013)).



Figure 1. Sketch of the DP architecture. This U-Net type architecture consists of a series of convolutional blocks: the encoder and the decoder blocks. The inputs first pass the encoder blocks for down-sampling, and then are up-sampled through the decoder blocks for the final outputs. Further details of the architecture and parameters can be found in Table 1.

In this work, three different methodologies are tested to inpaint maps at the locations of extragalactic point sources. Two of the inpainting techniques involve generative DCNNs. We compare the DCNNs inpainting performances with the standard diffusive inpainting approach used in Bucher et al. (2016), which is simply filling the missing pixel with the average value of its nearest-neighbors.

We organize the paper as follows. In Section 2, we present the three inpainting methodologies adopted in this work. Section 3 describes the data used for training and validation purposes. Finally, Section 4 includes the results achieved by inpainting on simulations (Section 4.1) and on more realistic data sets (Section 4.2). Finally, we apply our inpainting method to the map of the *S*-band Polarization All Sky Survey (SPASS; Carretti et al. 2019) at several source locations (Section 4.3) and demonstrate that we robustly recover the background signal.

2. Methods of Inpainting

Inpainting algorithms can be divided into two main groups: (i) diffusive-based methods and (ii) learning-based methods that rely on training DCNNs to fill the missing pixels with the predictions learned from a training data set. We choose three inpainting techniques from both groups: Section 2.1 describes a diffusive-based method from group (i), and Sections 2.2 and 2.3 present methods from group (ii).

2.1. Nearest-neighbors

One of the simplest inpainting methods is the diffusive inpainting described in Bucher et al. (2016), which has been adopted in Ade et al. (2014, 2016). In this method, each masked pixel is iteratively filled with the mean value of its nearest-neighbor pixels, being often referred to as the nearest-neighbors (NN) in the image reconstruction algorithm.

The iterative procedure can be performed in two ways. (i) The Gauss-Seidel method, which computes the average of neighbors at the current iteration. As a consequence, the pixels near the boundary are updated in earlier iterations while pixels near the center of the inpainting regions require several iterations. (ii) The Jacobi method, which estimates the average value from a buffer of pixel values at the previous iteration. Bucher et al. (2016) found that $\sim \mathcal{O}(10^3)$ iterations were needed to inpaint $\sim 10'$ areas on a map with $\sim 2'$ pixel size. Although Bucher et al. (2016) found that both methods did not impact the quality of the inpainted results, the Gauss-Seidel method achieves faster convergence than the Jacobi method. We therefore adopted the former method as suggested by Bucher et al. (2016).

 Table 1

 Architecture for DP Network as Described in Ulyanov et al. (2020)

Deep-Prior	Network
n _d	[16, 32, 64, 128, 128, 128]
k _d	[3, 3, 3, 3, 3, 3]
n _u	[128, 128, 128, 64, 32, 16]
k _u	[5, 5, 5, 5, 5, 5]

Note. n_u , n_d correspond to, respectively, the number of upsampling and downsampling filters, k_d , k_u correspond to the kernel sizes.

2.2. Deep-Prior

Deep-Prior (DP; Ulyanov et al. 2020) is the first methodology encoding DCNNs we use in this work. The fundamental assumption of DP is that the information required to reconstruct an image are essentially encoded in the input image itself (which might be already corrupted, noisy or with missing pixels) and in the network architecture used for the reconstruction. It has the peculiarity of being an untrained network and the inpainting procedure can be summarized as follows: (i) process the image to be inpainted through the convolutional layers for several iterations (usually more than 1000), (ii) fit for the neural network parameters (or weights), and (iii) evaluate a loss function for a given set of parameters. Ulyanov et al. (2020) proposed several loss functions specifically related to the task to be performed (e.g., superresolution, inpainting, and denoising).

Inpainting can be formalized as an image generating procedure, f_{θ} , mapping from the so-called latent space \mathbb{R}^m into the feature space \mathbb{R}^n , with $n = N_{\text{pix}} \times N_{\text{pix}}$ being the size of the input images and with parameters θ . Generative networks take as an input a random vector z, known as the prior and outputs an image $\tilde{x} = f_{\theta}(z)$ given z and the parameterization. Generally, the whole network architecture is symmetric and can be structured into two main parts: an encoder, which represents the input into a lower dimensional space and a decoder meant to do the opposite, upsample from the low dimensional space to the original input one.

We therefore build the DP architecture by following the prescription given in Ulyanov et al. (2020) for the task of inpainting,⁵ and building a *U-Net* type architecture sketched in Figure 1 and summarized in Tables 1 and 2.

The optimal set of parameters θ^* is obtained minimizing the loss function given a ground-truth image x_0 , with missing

⁵ Further details can be found in supplementary material, https://dmitryulyanov.github.io/deep_image_prior.

 Table 2

 Sequence of Layers Encoded in the *i*th Encoder and Decoder Block Adopted for DP

Encoder Block	S	Parameters	Decoder Block	S	Parameters
Conv2D	1	$k_d[i] \times k_d[i] \times n_d[i] \times 1$	Conv2D	1	$k_u[i] \times k_u[i] \times n_u[i] \times 1$
Conv2D	2	$k_d[i] \times k_d[i] \times n_d[i] \times 1$	LeakyRELU		$\alpha_{LR} = 0.1$
LeakyRELU		$\alpha_{LR} = 0.1$	Conv2D	1	$k_u[i] \times k_u[i] \times n_u[i] \times 1$
Conv2D	1	$k_d[i] \times k_d[i] \times n_d[i] \times 1$	LeakyRELU		$\alpha_{LR} = 0.1$
LeakyRELU		$\alpha_{LR} = 0.1$	UpSampling2D		size = (2, 2)

Note. *s* refers to the stride size.

pixels correspondent to a binary mask m (with 0 in the masked region and 1 elsewhere):

$$E(x^*, x_0) = \|(f_{\theta^*}(z) - x_0) \odot m\|^2, \tag{1}$$

with $x^* = f_{\theta^*}(z)$ being the output of the inpainting reconstruction procedure and \odot the Hadamard's product. The minimization process of Equation (1) is performed with the gradient descent algorithm. Notice that this loss function does not depend on the values of the missing pixels, but only on the features outside the mask where the ground-truth is known.

2.2.1. Deep-Prior Architecture

The architecture of DP is sketched in Figure 1, the network read the inputs are 128×128 images of Galactic foreground maps and they are processed through a series of six convolutional blocks aimed at down-sampling the images (with a kernel size of $k_d = 3$) and an increasing number of channels, n_d (ranging from 16 to 128). Each block includes two convolutional layers, respectively, with stride s = 1 and s = 2followed by the application of the Leaky Residual Exponential Linear Unit (LeakyRELU) activation, with a learning rate parameter of $\alpha_{LR} = 0.1$. At the end of each down-sampling block there is another convolutional layer with stride s = 1followed by another LeakyRELU. The decoder blocks follow the exact symmetry of the encoder ones and differ only in the kernel size adopted $k_u = 5$. The structure of each upsampling block include a convolution layer and LeakyRELU repeated twice and terminates with an upsampling layer increasing the size of the output by a factor of 2. In total, the DP network has 4,237,441 parameters. The whole architecture is implemented using the Keras⁶ package with TensorFlow backend⁷

2.3. Generative Adversarial Networks with Contextual Attention

Generative Adversarial Networks (GAN; Goodfellow et al. 2014) is a popular machine-learning algorithm useful in several contexts, especially in image reconstruction. The overall GAN framework can be synthesized as an adversarial interplay between two networks, a generator, G, and a discriminator, D. G is trained to generate fake samples that resemble the training set, and D is responsible for judging whether a given sample is generated by G or is a real image that belongs to the training set. The loss of the network is designed such that G will get penalized for producing images that look fake (statistically deviating from the training set), and D will get penalized for misjudging the originality of the images. The two networks are

then trained until the generator is able to generate good quality images and the discriminator cannot tell if they are generated or not.

In the context of image reconstruction, GAN seeks for coherency between generated and existing pixels by adopting a convolutional encoder-decoder generator network (similar to the one described in Section 2.2). As a result, one of the biggest advantages of using GAN based methods is that they are less affected by problems observed in simple convolutional networks such as boundary artifacts and blurry textures, making the reconstructed images inconsistent with the surrounding regions (Iizuka et al. 2017; Yu and Koltun 2015).

Recently, Yu et al. (2018) presented a novel generative inpainting procedure based on GAN with an extra-branch in the architecture aiming at providing more coherence in the reconstructed area given features in the uncorrupted regions of the image, referred to here as the contextual attention branch. The proposed generator network can be then structured into two stages: a coarse reconstruction stage, which employs a generator built with an encoder-decoder architecture trained to inpaint the missing region, and a refinement stage aimed at improving the generator with local and global features estimated with contextual attention. The first stage, the coarse stage, is based on an encoder-decoder network to roughly generate the content in the missing region. The second stage, the refinement stage, is organized into two parallel convolutional pathways both fed with the output of the coarse stage, i.e., the full image with an approximated content in the missing region. One of the two branches aims at hallucinating novel contents in the missing region in order to better refine the content inside the mask by injecting smaller scale features. The other branch encodes the contextual attention to enhance spatial coherency of the local features inside the masked area with the global features. To better visualize how the contextual attention works, we report in Figure 3 an example of inpainting with GAN and the attention map related to this case. The attention map helps to identify which regions of the input image the contextual attention focuses on in order to refine the corrupted image.

The output of the refinement stage is then combined and fed to a single decoder for the final inpainting output. Both the coarse and refinement stages of the generator network are trained end-to-end with reconstruction losses. We follow the prescriptions in Yu et al. (2018) and adopt a weighted sum of pixel-wise ℓ_1 loss to train explicitly the coarse and the refinement networks. For the discriminator, we adopted two Wasserstein GAN (WGAN) adversarial losses (Arjovsky et al. 2017) aimed at evaluating separately the global and local features of the generated images.

⁶ https://keras.io/api/

⁷ https://www.tensorflow.org/



Figure 2. Sketch of the GAN architecture with contextual attention implemented in this work. The Generator network consists of a coarse and a refinement stage, which, respectively, produce a rough and refined reconstruction of the corrupted image. The input of the Discriminator coincides with the output of the refinement stage. Both Generator and Discriminator are trained end-to-end together. Further details of the network architecture and parameters can be found in Table 3.

Table 3		
Architecture and Parameters Used for GAN as Described in Yu et al. ((2018)	

Inpainting Network		Contextual-attention Branch
n _d	[32,64,64,128,128,128,128	L1:[32,64,64,128,128,128,128]
	128,128,128,128,128]	L2:[128,128]
k _d	[5,3,3,3,3,3,3,3,3,3,3]	L1:[5,3,3,3,3,3], L2:[3,3]
S_d	[1,2,1,2,1,1,1,1,1,1,1]	L1:[1,2,1,2,1,1], L2:[1,1]
D	[1,1,1,1,2,4,8,16,1,1]	
n _u	[128,64,64,32,16]	[128,64,64,32,16]
k _u	[5,5,5,5]	[5,5,5,5]
S _u	[1, 1, 1, 1, 1]	[1,1,1,1]
Local WGAN		Global WGAN
n	[64,128,256,512]	[64,128,256,256]
k	[5,5,5]	[5,5,5,5]
S	[2,2,2,2]	[2,2,2,2]

Note. For each network, we indicate with n, k, s, respectively, the number of channels, the kernel size, and the stride size. In particular, for the inpainting and contextual attention networks we label with d(u) the sizes for the down-sampling (upsampling) of filters. D refers to the dilation rate. L1 and L2 in the contextual attention column refer to the two parallel encoders in the refinement stage.

2.3.1. GAN Architecture

The overall architecture used in generative inpainting is shown in Figure 2 (details for convolutional layers can be found in Table 3). Moreover, we choose to adopt the same hyperparameters in setting the network as the one provided in Yu et al. (2018). Each convolutional layer is implemented by using mirror padding, without batch normalization and with Exponential Linear Units (ELUs; Clevert et al. 2015) activation functions. We summarize into four subnetworks the GAN implementation in this work:

- 1. *Inpainting network* performs the coarse stage of the Generator network, and presents an encoder-decoder architecture made of 16 blocks of convolutional layers. For the down-sampling part, we choose k = 3 for all the kernels, except for the first one (k = 5), and set the stride to s = 2 only for the second and fourth convolutional block (for the rest s = 1). Starting from the sixth to the ninth block dilation is also applied to the convolution with an increasing dilation rate from 2 to 16. The application of dilations aims at capturing more global features from the input without increasing the size of parameters. The upsampling part encodes two upsampling layers after the twelfth and the fourteenth block, for all the convolutions we choose $k_u = 5$ and $s_u = 1$. The number of channels are reported in Table 3.
- 2. Contextual attention branch is aimed at refining the coarsely inpainted image output from the previous network. This branch has a very similar architecture as the inpainting one with the main difference that it is made of two parallel encoders concatenated to a single decoder. One of the encoders estimates the contextual attention. In this branch, the attention map (Figure 3) is estimated. The architecture details are listed in Table 3 and we refer to L1 and L2 for the two encoders.
- 3. Local and global critics in the Discriminator network aim at labeling whether the inpainted image is coherent inside and outside the masked area. Both critics are made of four blocks with k = 5 and s = 2 and with an increasing number of channels *n* from 64 to 512 (256) for the local (global) network. The last convolution is linked onto a fully connected layer and to the WGAN loss function.

3. Data and Simulations

In this study, we mainly focus on inpainting maps of two emissions in the microwave regimes: i.e., synchrotron and thermal dust. These two foreground components contaminate the CMB polarization measurements and need to be modeled both at large and small angular scales for foreground removal. Since the statistical properties of Galactic foregrounds are highly



Figure 3. Visualization of an inpainting contextual-attention map. The colors in the color coding indicate the portion of the whole image the neural network has focused on in order to inpaint a given pixel location in the masked area. In this particular example, most of the pixels are inpainted by GAN by looking at the upper and lower left region of the image.

non-Gaussian, an interesting application of DCNNs is to reconstruct images with complex, non-Gaussian features. In contrast, traditional inpainting methods used in CMB studies such as the Gaussian Constrained Inpainting methods (Hoffman & Ribak 1991; Bucher & Louis 2012) assume that the background is a Gaussian field, making it incapable of capturing the non-Gaussianity in the foregrounds.

Both Galactic foregrounds, unpolarized and polarized emissions data (respectively, encoded in the brightness temperature, T, and Stokes parameters Q and U maps) are simulated using the PySM package (Thorne et al. 2017) and will be described below in Sections 3.1 and 3.2.

3.1. Galactic Synchrotron

For the synchrotron data, we consider SPASS (Carretti et al. 2019) which observed the Southern sky ($\delta < -1^{\circ}$) at 2.3 GHz with an 8.9 arcmin full width at half maximum (FWHM). The methodology used to generate the intensity and polarization maps are described in Carretti et al. (2019). 98.6% of the pixels in the Q and U SPASS maps have signal-to-noise ratio (SNR) >3, making these maps a promising synchrotron polarization template (Krachmalnicoff et al. 2018). Lamee et al. (2016) cross-matched the extragalactic radio quasars (mostly steep spectrum sources) detected by SPASS with the ones detected by the NRAO/VLA Sky Survey (NVSS; Condon et al. 1998), at 1.4 GHz and released a polarization catalog with 533 bright sources in the overlapping area of the two surveys. However, because the SPASS T, O, and U maps are filled with radio sources, an assessment on the map level to the smallest angular scales is essentially compromised by the point-source bias. Therefore, masking and inpainting these sources provide an overall benefit in fully exploiting the angular scales probed by SPASS.

In order to inpaint the SPASS map, we first create a simulated training data set, which will be used for training the GAN (training set) and for evaluating the quality of the reconstructions (testing set). We therefore simulate SPASS synchrotron-only TQU maps at the SPASS frequency with s1 PySM model. This model is one of the most representative since it parameterizes the synchrotron power-law emission with a spatially varying spectral index. Maps are pixelized on a HEALPix⁸ (Górski et al. 2005; Zonca et al. 2019) nside=2048 grid convolved with an 8/9 FWHM beam.

3.2. Thermal Dust

We use the thermal dust maps at 353 and 857 GHz from the third Planck public release.⁹ Both frequency maps are dominated by the thermal dust emission emitted by our own Galaxy and encode contribution from cosmic infrared background (CIB). We choose the 353 GHz frequency channel in order to test the inpainting techniques on both dust temperature and polarization maps. Because the 857 GHz channel is not polarization sensitive, we use this channel to assess how the different SNR and different CIB contribution affect inpainting reconstructions in total intensity. At these frequencies, the emission is dominated by star-forming galaxies, and blazars are expected to have minor contribution to the total intensity.

We build the training set by simulating TQU Thermal dust maps at 353 GHz with the d1 PySM model, which describes the modified blackbody emission law with a spectral index and a temperature, both spatially varying. Maps are simulated on a nside = 2048 grid and convolved with a 5' FWHM beam, similarly to the one in the Planck 353 GHz observations.

Furthermore, since the pixel values of the maps are rescaled during the training and inpainting processes, the GAN reconstruction is not affected by the overall amplitude of a signal. We will show in Section 4.2 that inpainting performances do not change as a function of frequency as long as the brightest signal in the map coincides with the one used in the training. This independence of frequency is the reason why we trained the GAN for thermal dust emission using PySM signal-only simulations with single frequency at 353 GHz.

3.3. The Training Data Set

Both the training and testing data sets are made from the PySM simulated maps. We forecast with PS4C package¹⁰ (Puglisi et al. 2018), the number of sources, $N_{\rm src}$ whose density fluxes will be detected at 5σ significance above the sensitivity flux. For a generic large aperture forthcoming CMB experiment (e.g., The Simons Observatory Collaboration et al. 2019), the forecasted detection is about $N_{\rm src} \sim 30,000$. We then generate a point-source mask by randomly extracting $N_{\rm src}$ locations following a Poisson distribution. We mask the sources with circular holes centered at the source locations and with a radius three times larger than the beam FWHM size (namely 26.7 and 15', respectively, for synchrotron and dust maps).

Both masked and unmasked maps are then split into $3 \times 3 \text{ deg}^2$ square tiles, composed of 128×128 pixels with resolution closer to the HEALPix one (i.e., $\sim 1'.5$ at nside=2048). We finally build the training set for the GAN network by combining 45, 000 images from square patches extracted equally from T, Q, and U maps. The remaining 5000 images are used for validation and 500 for testing.

4. Results

Figures 4, 5, and 6 show examples of maps extracted from the test set and reconstructed with the three methods outlined in Section 2. We estimate the minimum and maximum values of each ground-truth image to rescale it (together with the respective inpainted ones) to [0, 1] with the MinMax normalization. This rescaling forces the generated maps to have the same range as the test ones, so the differences between the

⁸ https://healpix.sourceforge.io

⁹ https://pla.esac.esa.int

¹⁰ https://gitlab.com/giuse.puglisi/PS4C



(a)



(b)

Figure 4. Thumbnail $1.5 \times 1.5 \text{ deg}^2$ crops for (a) Q and (c) U map predictions of thermal dust. The radius of the reconstructed area (black circle) is 15'. Columns from left to right show ground-truth maps from the test set, predictions obtained with DP, NN, and GAN, respectively. The colorbar is set to be the same in each row by MinMax rescaling all the images with the same min and max values of the ground-truth ones. Notice that we further zoom in on the dust maps to better inspect of the inpainted region (maps are originally $3 \times 3 \text{ deg}^2$). Temperature maps are shown in Figure 6 of the Appendix.

ground-truth and reconstructed map can be spotted more easily. Notice that we further zoom in on dust maps $(1.5 \times 1.5 \text{ deg}^2 \text{ crops})$ to better inspect the inpainted region.

As expected, the inpainting performed with the NN algorithm is smooth and lacks finer details, making them distinguishable from the original map.





(b)

Figure 5. Thumbnail $3 \times 3 \deg^2$ crops for (a) Q and (c) U map predictions of synchrotron emission. The arrangement and the colorbar setting are the same as those in Figure 4. The radius of the reconstructed area (black circle) is 30'. Temperature maps are shown in Figure 6.

On the contrary, for the inpainting performed with DP and GAN, it is harder to point out which one is the ground-truth and which is the reconstructed maps. Moreover, we note that DP

and GAN are able to reproduce the large-scale features, the most correlated angular scales, as well as the typical Q/U pattern for the polarization maps.



(b)

Figure 6. Thumbnail images for (a) dust and (b) synchrotron T maps predictions from the testing set simulated with PySM. The radius of the reconstructed area is shown as a gray circle.

4.1. Evaluation of Fidelity

 $\Delta \ell$ [deg]

We employ several methods used in the literature to assess the quality of the reconstructed maps quantitatively. First, we follow the approach from Mustafa et al. (2017) and Aylor et al. (2020), focusing on evaluating the ability to replicate the summary statistics of the underlying signal that needs to be

Figure 7. Top: thermal dust pixel intensity distribution of 500 generated maps with NN (diamonds), DP (stars), GAN (circles) compared to 500 maps from test set (black squares). Bottom: synchrotron pixel intensity distribution. The KS test statistics and the *p*-values listed in Table 4 indicate that the distribution of inpainted images are likely to be drawn from the distribution of the test set.

reconstructed. Those statistics are based on (i) the pixel intensity distribution, (ii) the angular power spectra of the two-point correlation function, and (iii) the first three Minkowski functionals. We, therefore, consider the inpainting as successful if it passes these three statistical tests.

The distribution of pixel intensity provides information about whether the range of pixel values of the ground-truth maps are reproduced in the generated maps. Figure 7 shows the histogram of the pixel intensities of 500 generated maps compared with the corresponding ground-truth maps from the test set. For the types of foregrounds we consider in this analysis, the amplitude is strongly directionally dependent. Therefore, we scale the sample maps with MinMax rescaling to [-1, 1] to reduce the patch-to-patch variation. Although the differences between the pixel distributions are nearly negligible, we run a Kolmogorov-Smirnov (KS) two-sample test on each case and assess how likely the distribution of inpainted images is drawn from the ground-truth distribution. We thus estimate the empirical distribution function on the pixel samples from the ground-truth images and from images inpainted with three methods introduced in Section 2, and then derive the KS two-sample statistics to test the null hypothesis. From the KS p-values summarized in Table 4, we find that p > 0.86 (0.997) for synchrotron (thermal dust) maps so that we cannot reject the null hypothesis.

In order to check whether the Fourier modes in the groundtruth are reproduced in the inpainted maps, we additionally evaluate the power spectra of intensity (TT), and polarization (EE and BB) maps¹¹ in each flat square map from the test set and the ones inpainted with the three methods.

Each spectra shown in Figure 8, is estimated with NAMAS-TER¹² (Alonso et al. 2019), binned into equally spaced multipoles with $\Delta \ell = 450$. The maximum multipole is chosen accordingly to the beam FWHM with whom the signal is convolved, i.e., $\ell_{max} = 4000$ for dust and $\ell_{max} = 2000$ for synchrotron. The median of the binned power spectra is plotted at each multipole estimated from the test set including 500 ground-truth maps and corresponding inpainted maps.

 Table 4

 p-Value of KS Test Performed on the Pixel Intensity Distribution of Q Maps

 Shown in Figure 7

	Synchrotron	Thermal Dust
NN	>0.997	>0.999
DP	>0.963	>0.999
GAN	>0.861	>0.997

The shaded-gray area represents 95% of the power spectra estimated from the test set and its vertical width indicates how much the amplitude of the signal can vary at different locations of the sky (as much as 2 orders of magnitude). The median power spectra estimated from inpainting the test set are shown as points, and the area within the dashed lines corresponds to the 95% of the power spectra. Notice that DP power spectra for thermal dust tend to systematically depart from the ones estimated with the test set spectra at $\ell > 2500$ scales. On the contrary, GAN and NN are overall consistent with the spectra from the ground-truth maps.

To assess more quantitatively that the power spectrum at a given ℓ bin is correctly reproduced, we consider three different multipole bins. We bootstrap resample 5000 times the distribution in each bin of 500 spectra and perform the KS test on the resampled distribution.

Figure 9 shows the distribution of EE spectra for thermal dust (top) and synchrotron (bottom) and favors GAN as the method that better resembles the ground-truth distribution, with KS *p*-value >0.978 (>0.808) for dust (synchrotron) spectra. On the other hand, NN spectra present the lowest *p*-values of >0.808 for dust and >0.538 for synchrotron.

In Table 5, we summarize the KS test *p*-values estimated on each multipole bin and after having bootstrapped resampled the TT, EE, and BB spectra 5000 times in each bin. In total, we account for (3×9) 27 KS *p*-values for dust and (3×5) 15 KS *p*-values for synchrotron spectra. Although there are few bins where the *p*-value is as low as 0.153, the KS statistics for those bins is small enough that the null hypothesis cannot be rejected at a significance level of $\alpha > 5\%$. We therefore conclude that all three methodologies are able to reproduce angular correlations of the underlying signal coherently.

¹¹ Following the decomposition of Q and U maps proposed in Seljak & Zaldarriaga (1997) and Hu & White (1997).

¹² https://github.com/LSSTDESC/NaMaster

Figure 8. From left to right, TT, EE, and BB power spectra estimated from a sample of 500 thermal dust (top) and synchrotron (bottom) maps. Median power spectra are shown as black dotted lines for the test set, blue diamonds for NN inpainting, orange stars for DP inpainting, and green circles for GAN inpainting. The gray shaded area corresponds to 95% of the test set spectra to be compared with 95% of the set of power spectra inpainted with NN, DP, and GAN, respectively, shown as dashed blue, dashed orange, and dashed green. The spectra are uniformly binned with $\Delta \ell = 450$.

Figure 9. Distributions of EE power spectra of thermal dust (top panel) and synchrotron (bottom panel) at three different multipole bins resampled from the power spectra of 500 images inpainted with NN (diamonds), DP (stars), GAN (circles). We compare them with the resampled EE spectra of test maps (black squares). Error bars are estimated as the squared root of number of elements in each bin after bootstrap resampling.

THE ASTROPHYSICAL JOURNAL, 905:143 (19pp), 2020 December 20

 Table 5

 p-Value of KS Test Performed on Each Multipole Bin

	Synchrotron	Thermal Dust
NN	>0.306 (2 bins)	>0.537 (4 bins)
	>0.792 (13 bins)	>0.801 (23 bins)
DP	>0.538 (5 bins)	>0.153 (4 bins)
	>0.792 (10 bins)	>0.788 (23 bins)
GAN	>0.538 (5 bins)	>0.153 (3 bins)
	>0.792 (10 bins)	>0.788 (24 bins)

Note. We combined together TT, EE, and BB spectra, 27 (15) bins in total for dust (synchrotron).

Since the Galactic emission is highly non-Gaussian, it is essential to evaluate how well the methodologies described here are able to reproduce the non-Gaussian features. We thus evaluate the three Minkowski functionals V_0 , V_1 , V_2 for each rescaled map (ranging in [-1, 1]). The first three Minkowski functionals are related, respectively, to the area, the perimeter and the connectivity in an image as a function of a threshold ρ . The dotted black lines in Figures 10 and 11 show the median for the three Minkowski functionals (calculated using Mantz et al. 2008) estimated from 500 samples and evaluated at 10 equally spaced thresholds. Minkowski functionals for inpainted images are shown in Figures 10 and 11 with the same color scheme as in Figure 8. We notice that both GAN and NN are able to fully reproduce the non-Gaussianity of both the unpolarized and the polarized Galactic emission.

Similar to what is stated above, the Minkowski functionals estimated on the maps inpainted with DP clearly depart from the ground-truth especially at intermediate thresholds, $(-0.5 < \rho < 0.5)$ for the V_1 and V_2 functionals. This can point to further investigations for the DP inpaintings especially for better characterizing the non-Gaussian features and the small angular scales $\ell > 2500$ of both dust and synchrotron that are not fully captured by the DP network.

4.2. Validation on Real Data

To further test and validate our methodologies, we run tests on real data, cropping 500 images at random locations from the Planck maps at 353 and 857 GHz.

Figure 12 shows the $1.5 \times 1.5 \text{ deg}^2$ Planck maps at (a) 353 and (b) 857 GHz. Notice that we deliberately choose two locations to highlight reconstruction performances at two different SNR regimes. Closer to the Galactic midplane, the dust emission is stronger so that SNR is high at both 353 and 857 GHz, e.g., see top panels of Figures 12(a) and (b). On the contrary, at high Galactic latitudes and at 353 GHz, the SNR can be lower and noise contribution is clearly noticeable in the maps, e.g., as in the bottom panel of Figure 12(a). In this case, the inpainting performances with DP and GAN can be affected by noise, and the reconstruction area can be visually distinguished in the square patch.

For a more quantitative assessment, we thus estimate the summary statistics shown in Section 4.1 and we show the results in Figure 13.

A clear indication that the thermal dust emission data is contaminated by instrumental noise can be inferred by comparing the shapes of Minkowski functionals estimated for the maps at 353 GHz, with the ones from signal-only simulations (Figure 4) or from signal-dominated maps (Figure 13(b)). The morphology of the Minkowski functionals of the former largely resembles the functionals estimated from a Gaussian signal. Two possible candidates for the Gaussian component are (i) Planck instrumental noise, which can be approximated as white within $\sim 9 \text{ deg}^2$ patches, and (ii) CMB residual emission.

DP inpaintings are the most affected ones in the presence of the noise at 353 GHz, e.g., notice how the pixel distribution (top left panel) of Figure 13(a) noticeably departs from the ground-truth one. However, the KS test *p*-value performed on the DP and ground-truth samples is p > 0.536, i.e., not significantly low enough to reject the null hypothesis that the two samples are different. On the other hand, inpainting with GAN and NN does not show any dependence with SNR, as the performances with these methodologies are essentially similar to the ones observed with the signal-only simulated maps (e.g., Figures 7, 8, and 10).

Finally, we would like to point out that for the case of inpainting with GAN, we use the weights derived from the training set composed of signal-only dust TQU simulated maps. Looking at Figure 13 we notice that GAN is able to statistically reproduce at 353 GHz the features composed by signal and noise, indicating that the network has correctly learned the features related to the intrinsic signal in the presence of noise, and injects signal plus noise features statistically coherent with the ones outside the masked area. However, when the noise is highly dominating in the patch, we can clearly distinguish smooth artifacts in some cases inpainted with GAN (see Figure 12(a, bottom)). This is somewhat expected because GAN is trained on signal-only simulated images. Further investigations on training GAN with noisy data are needed, and we will address it in a future work.

4.3. Inpainting Maps with Point Sources

In this section, we aim at showing real world applications of the three reconstruction methodologies by inpainting areas in real maps with extragalactic radio sources.

We consider the TQU SPASS synchrotron map at 2.3 GHz¹³, and we run a Matched Filter (Marriage et al. 2011) to detect the brightest polarized sources in the map. We consider 7σ as a threshold for a point-source detection. In particular, we focus on unresolved point sources (i.e., sources whose projected solid angle is smaller than the SPASS beam solid angle) detected at intermediate Galactic latitudes $|b| > 20^{\circ}$. As a result, 45 polarized sources are detected in the SPASS map, which is very close to 60, the number forecasted by PS4C with the adopted SPASS specifications.

Figures 14 and 15 show a selection of images extracted from the SPASS TQU maps and centered at the coordinates of the detected sources. The shape and the size of the region to be inpainted are chosen proportionally to the flux of each source (see the black circle in Figure 14). However, we expect the inpainting not to be affected by different shapes and/or sizes of the masked area as demonstrated in Yu et al. (2018) and Ulyanov et al. (2020). Moreover, we set the color scale of the input SPASS map to be the same as the one in the inpainted maps to highlight the consistency with reconstructed maps. As expected, images inpainted with GAN are visually injecting

¹³ Maps are available online at https://sites.google.com/inaf.it/spass.

Figure 10. From left to right, V_0 , V_1 , V_2 Minkowski functionals estimated from the test set of 500 thermal dust T, Q, and U maps, respectively, in (a), (b), and (c). We use the same coloring scheme as in Figure 8: (black dotted) median of the functionals estimated from the test set, 95% of the functionals is shown as a gray shaded area. Points and dashed lines refer to medians and the 95% interval of the functionals estimated from the sets of inpainted maps with NN (blue diamonds and blue dashed), DP (orange stars and orange dashed), and GAN (green circles and green dashed).

more coherent features and less artifacts with respect to DP and NN.

Given the presence of the point source at the center of the patch biasing the evaluation of fidelity with the pixel distribution and the Minkowski functionals, we estimate the power spectra from all the sets of maps in order to assess more quantitatively the quality of the generated maps. We masked the source with a circular mask with 30' radius and estimate the spectra in the area outside the mask. On the other hand, we did not apply the point-source mask for the power spectra estimated from the inpainted maps.

We estimate the power spectra as described in the previous sections. In Figure 16, we show the TT, EE, and BB power spectra, which indicates all methodologies essentially are able to reproduce consistently the power spectrum at all angular scales of the input SPASS masked maps. Moreover, on smaller

Figure 11. From left to right, V_0 , V_1 , V_2 Minkowski functionals estimated from the test set of 500 synchrotron TQU maps, respectively, in (a), (b), and (c). We use the same coloring scheme as in Figure 8: (black dotted) median of the functionals estimated from the test set, 95% of the functionals is shown as a gray shaded area. Points and dashed lines refer to medians and 95% interval of the functionals estimated from the sets of inpainted maps with NN (blue diamonds and blue dashed), DP (orange stars and orange dashed), and GAN (green circles and green dashed).

angular scales, the power spectra from maps inpainted with GAN show a lower amplitude tail, possibly implying that Poissonian bias from undetected point sources is further reduced (a factor of \sim 4 for EE and BB spectra).

AStrophysical SOurces (PICASSO github.com/giuspugl/ picasso §). It has been made publicly available together with a documentation web page.¹⁴ We trained GAN separately on dust and synchrotron map

sets.¹⁵ The training process for each set took \sim 12 hr on a GPU

5. Code Release

The three inpainting methods have been collected into a python package, the Python Inpainter for Cosmological and

¹⁴ http://giuspugl.github.io/picasso/

¹⁵ Training weights can be downloaded from https://bit.ly/2TI6x4o.

(b)

Figure 12. Thumbnail images of Planck temperature maps at (a) 353 and (b) 857 GHz. The radius of the reconstructed area (black circle) is 15'. Two locations are chosen to highlight different inpainting performances with high and low SNRs of Planck 353 map, respectively, at high Galactic latitude, i.e., $(l, b) = (65^\circ, 7, -7^\circ)$ and at low Galactic latitude, i.e., $(l, b) = (73^\circ, 3, -37^\circ)$.

Figure 13. Summary statistics estimated on Planck dust T maps at (a) 353 and (b) 857 GHz.

(b)

Figure 14. From left to right, input SPASS, DP, NN, and GAN polarization maps inpainted in the 30' region surrounding the detected point source (black circle). The range of the input map is chosen to be the same as that of the inpainted ones. Temperature maps are shown in Figure 15.

node of SHERLOCK Cluster of Stanford Supercomputing Center with four interconnected NVIDIA Tesla P40 GPUs.¹⁶

Finally, we measure the inpainting speed for each of the three methods. We performed this benchmark within a GPU

node at NERSC equipped with four interconnected NVIDIA Tesla –V100 GPUs.¹⁷ GAN is the fastest method since fast-forwarding the trained weights is very quick and it takes $\sim 4 \div 6$ s. Although NN does not involve any DCNN in the

¹⁶ For more details, see https://www.sherlock.stanford.edu.

¹⁷ https://docs-dev.nersc.gov/cgpu

Figure 15. SPASS temperature maps inpainted in the 30' region surrounding the detected point source (black circle). The range of the input map is chosen to be the same as that of the inpainted ones.

Figure 16. TT, EE, and BB power spectra estimated on the SPASS map. In this case, the input map encodes a point source located at the center and is masked out before computing the power spectra. Vice versa, spectra estimated on inpainted maps do not have the mask applied.

map reconstruction, it iterates over the pixels in the missing region and it takes ~ 15 s per image. A single inpainting with DP takes ~ 30 s, since this is the time spent to minimize the loss function in Equation (1) over 3000–5000 epochs with gradient descent.

6. Discussion

In the near future, several high-resolution experiments from the ground, e.g., Simons Observatory, the South Pole Observatory, and CMB-Stage IV (The Simons Observatory Collaboration et al. 2019; Carlstrom et al. 2019; Abazajian et al. 2019), are expected to detect more than 10,000 sources in total density flux representing a critical contaminant especially at small angular scales (Puglisi et al. 2018; Naess 2019; Lagache et al. 2020). This is mostly due to an improvement in sensitivity and an increase in the footprint area (\sim 40% of the sky) with respect to the ones that have been surveyed so far. To mitigate the contamination issue, detected point sources are usually masked out from the maps. However, a mask encoding more than 10,000 can result in biasing the two-point (as well as the higher-order) correlation functions used to estimate cosmological parameters. Several methodologies have been proposed in the literature to clean the maps from point sources, e.g., if the source is unresolved by fitting the beam profile at the map level (Naess 2019; Datta et al. 2019) as well as inpainting with several methodologies, not necessarily involving DCNNs (Bucher & Louis 2012).

On the other hand, point sources contaminate the Galactic foreground emission especially far from the Galactic midplane (i.e., at high Galactic latitudes). A mitigation of point-source contamination in polarized foreground maps at small angular scales is thus needed in order to de-lens the gravitational lensing B-mode and assess the amplitude of the primordial B-mode.

In this work, we inpaint thermal-dust and synchrotron intensity and polarized emissions in order to deliver pointsource-free foreground template maps. We show that not only the pixel distribution and two-point correlation function of the inpainted maps, but also their higher statistic moments commonly adopted for non-Gaussian maps, i.e., the Minkowski functionals, are statistically consistent with the ones from the real data set by means of the KS test. Given the fact that applications related to the foreground maps mostly involve maps and spherical harmonic expansions, the statistical tests, implemented in this work, ensure that the inpainted maps do not have any bias when compared with the authentic ones outside the masked area. However, we plan to address in a future work what is the eventual bias introduced by inpainting methodology by estimating the bi- and tri-spectrum of angular correlation.

Finally, both the two DCNN methodologies implemented in this work are deterministic, meaning that they generate one single inpainting result for a given corrupted image. This makes it very hard to perform a χ^2 analysis on the inpainted images as the pixel–pixel covariance matrix cannot be defined with one sole inpainting case. However, a novel methodology based on DCNN has been proposed by Cai & Wei (2019), the Pluralistic Image Inpainting GAN (PiiGAN¹⁸). The generator in PiiGAN is designed such that it can generate multiple results from the contextual semantics of one image allowing us to perform the χ^2 analysis on inpainted maps. We devote a future work to apply PiiGAN for the specific case of Galactic Foreground inpainting.

7. Summary and Conclusions

In this work, we demonstrate three inpainting methodologies that can reproduce an underlying non-Gaussian signal without modifying the overall summary statistics of the signal itself.

The first method (NN) has already been used in the literature and it is based on diffusing the pixels in the masked area with the average of nearest-neighbor pixels. We further adopted two novel techniques relying on DCNNs, namely DP and GAN. They are first introduced as a tool to inpaint natural images by the deep-learning community. We validate the three techniques on simulated data, and test them on a data set with a wide range of SNRs. In addition, we show a real-life application by inpainting a map in regions where bright point sources are detected. To evaluate the quality of inpainted results, we adopted three summary statistics based on the pixel distribution, including the angular power spectra, and the first three Minkowski functionals. We find that all techniques are able to reproduce the overall summary statistics when applied to signal-only data.

Inside the masked regions inpainted with NN, the results are smooth but lack finer details. Generally, because NN averages the neighboring pixels, a map inpainted with NN sharply transitions from the area outside the mask with substructures and noisy pixels to a very smooth one encoding only long and smooth modes inside the masked area. However, we did not notice any clear effect or bias due to NN inpainting on the statistical tests we adopted in this work.

On the other hand, DP reconstructions on images extracted from signal and noise maps present different Minkowski functionals with respect to the ground-truth ones. As pointed out at the end of Section 4.1, this failure case of DP needs to be further investigated by means of a better tuning of hyperparameters.

GAN has been demonstrated to be promising as it is able to statistically reproduce signal and noise maps, and it generates images visually very similar to the ground-truth on signaldominated maps. However, we have identified cases where it fails to produce high fidelity images in noise dominated maps. We plan to further investigated this in a future work by training GAN with a more realistic data set including several levels of SNR.

To our knowledge, this is the first time that GAN has been used to successfully generate high-resolution intensity and polarization maps of Galactic foreground polarization maps. This approach opens up many possibilities of generating foreground maps using adversarial networks, which overcome the limitations of existing templates.

In conclusion, we focused this work on Galactic foreground emission motivated by the challenges in inpainting the non-Gaussian signal and in dealing with Galactic (and extragalactic) foregrounds in CMB B-mode polarization studies. For future work, we plan to apply similar techniques to different non-Gaussian signals spanning from galaxy weak lensing to HI data, which are highly affected by foreground emission as well.

The authors are very thankful to Yuuki Omori and Ben Thorne for carefully reading the manuscript. They also thank Alexandre Refregier, Nicoletta Krachmalnicoff, Ioannis Liodakis, and Warren Morningstar for the fruitful discussions and useful comments shared during the development of this work. G.P. acknowledges support from the TCAN grant "Modeling Polarized Galactic Foregrounds For Cosmic Microwave Background Missions," No. 80NSSC18K1487. Some of the results in this paper have been derived using the HEALPIX (Górski et al. 2005; Zonca et al. 2019) package.

References

- Abazajian, K., Addison, G., Adshead, P., et al. 2019, arXiv:1907.04473
- Adachi, S., Aguilar Faúndez, M. A. O., Arnold, K., et al. 2019, arXiv:1910. 02608
- Ade, P. A. R., Aghanim, N., Armitage-Caplan, C., et al. 2014, A&A, 571, A24
- Ade, P. A. R., Aghanim, N., Arnaud, M., et al. 2016, A&A, 594, A17
- Alonso, D., Sanchez, J., Slosar, A., Dark, L. S. S. T. & Energy Science Collaboration 2019, MNRAS, 484, 4127
- Arjovsky, M., Chintala, S., & Bottou, L. 2017, arXiv:1701.07875
- Aylor, K., Haq, M., Knox, L., Hezaveh, Y., & Perreault-Levasseur, L. 2020, MNRAS, 500, 3889
- BICEP2 Collaboration, Keck Array Collaboration, Ade, P. A. R., et al. 2018, PhRvL, 121, 221301
- Bucher, M., & Louis, T. 2012, MNRAS, 424, 1694
- Bucher, M., Racine, B., & van Tent, B. 2016, JCAP, 2016, 055
- Cai, W., & Wei, Z. 2019, IEEE Access, 8, 48451
- Caldeira, J., Wu, W., Nord, B., et al. 2019, A&C, 28, 100307
- Carlstrom, J., Abazajian, K., Addison, G., et al. 2019, BAAS, 51, 209
- Carretti, E., Haverkorn, M., Staveley-Smith, L., et al. 2019, MNRAS, 489, 2330
- Clevert, D.-A., Unterthiner, T., & Hochreiter, S. 2015, arXiv:1511.07289
- Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, AJ, 115, 1693
- Datta, R., Aiola, S., Choi, S. K., et al. 2019, MNRAS, 486, 5239
- Dickinson, C., Ali-Haïmoud, Y., Barr, A., et al. 2018, NewAR, 80, 1
- Farsian, F., Krachmalnicoff, N., & Baccigalupi, C. 2020, JCAP, 2020, 017 Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., et al. 2014, arXiv:1406.2661
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
- Gupta, N., Reichardt, C. L., Ade, P. A. R., et al. 2019, MNRAS, 490, 5712

Hoffman, Y., & Ribak, E. 1991, ApJL, 380, L5

Hu, W., & White, M. 1997, PhRvD, 56, 596

Iizuka, S., Simo-Serra, E., & Ishikawa, H. 2017, ACM Trans. Graphics, 36, 107

¹⁸ https://github.com/vivitsai/PiiGAN

- Krachmalnicoff, N., Carretti, E., Baccigalupi, C., et al. 2018, A&A, 618, A166 Krachmalnicoff, N., & Tomasi, M. 2019, A&A, 628, 129
- Lagache, G., Bethermin, M., Montier, L., Serra, P., & Tucci, M. 2020, A&A, 642, A232
- Lamee, M., Rudnick, L., Farnes, J. S., et al. 2016, ApJ, 829, 5
- Louis, T., Grace, E., Hasselfield, M., et al. 2017, JCAP, 2017, 031
- Mantz, H., Jacobs, K., & Mecke, K. 2008, JSMTE, 2008, P12015
- Marriage, T. A., Baptiste Juin, J., & Lin, Y.-T. 2011, ApJ, 731, 100
- Mustafa, M., Bard, D., Bhimji, W., et al. 2017, ComAC, 6, 1
- Naess, S. K. 2019, JCAP, 2019, 060
- Perraudin, N., Srivastava, A., Lucchi, A., et al. 2019, arXiv:1908.05519
- Planck Collaboration, Akrami, Y., Ashdown, M., et al. 2020, A&A, 641, A4 POLARBEAR Collaboration, Ade, P. A. R., Aguilar, M., et al. 2017, ApJ, 848, 121
- Puglisi, G., Fabbian, G., & Baccigalupi, C. 2017, MNRAS, 469, 2982

- Puglisi, G., Galluzzi, V., Bonavera, L., et al. 2018, ApJ, 858, 85
- Rodríguez, A. C., Kacprzak, T., Lucchi, A., et al. 2018, ComAC, 5, 4
- Sayre, J. T., Reichardt, C. L., Henning, J. W., et al. 2020, PhRvD, 101, 122003 Seljak, U., & Zaldarriaga, M. 1997, PhRvL, 78, 2054
- Serjak, U., & Zaldamaga, M. 1997, PirkvL, 78, 2034
- Starck, J.-L., Fadili, M. J., & Rassat, A. 2013, A&A, 550, A15
- The Simons Observatory Collaboration, Ade, P., Aguirre, J., et al. 2019, JCAP, 2019, 056
- The Simons Observatory Collaboration, Ade, P., Aguirre, J., Ahmed, Z., et al. 2019, JCAP, 2019, 056
- Thorne, B., Dunkley, J., Alonso, D., & Naess, S. 2017, MNRAS, 469, 2821
 Ulyanov, D., Vedaldi, A., & Lempitsky, V. 2020, Int. J. Comput. Vision, 128, 1867
- Yu, F., & Koltun, V. 2015, arXiv:1511.07122
- Yu, J., Lin, Z., Yang, J., et al. 2018, arXiv:1801.07892
- Zonca, A., Singer, L., Lenz, D., et al. 2019, JOSS, 4, 1298