

CENN: A fully convolutional neural network for CMB recovery in realistic microwave sky simulations

J. M. Casas^{1,2}, L. Bonavera^{1,2}, J. González-Nuevo^{1,2}, C. Baccigalupi^{3,4,5}, M. M. Cueli^{1,2}, D. Crespo^{1,2},
E. Goitia¹, J. D. Santos^{1,2}, M. L. Sánchez^{1,2}, and F. J. de Cos^{2,6}

¹ Departamento de Física, Universidad de Oviedo, C. Federico García Lorca 18, 33007 Oviedo, Spain
e-mail: casasjm@uniovi.es

² Instituto Universitario de Ciencias y Tecnologías Espaciales de Asturias (ICTEA), C. Independencia 13, 33004 Oviedo, Spain

³ SISSA, Via Bonomea 265, 34136 Trieste, Italy

⁴ IFPU – Institute for fundamental physics of the Universe, Via Beirut 2, 34014 Trieste, Italy

⁵ INFN-Sezione di Trieste, via Valerio 2, 34127 Trieste, Italy

⁶ Escuela de Ingeniería de Minas, Energía y Materiales Independencia 13, 33004 Oviedo, Spain

Received 2 March 2022 / Accepted 27 July 2022

ABSTRACT

Context. Component separation is the process with which emission sources in astrophysical maps are generally extracted by taking multi-frequency information into account. It is crucial to develop more reliable methods for component separation for future cosmic microwave background (CMB) experiments such as the Simons Observatory, the CMB-S4, or the LiteBIRD satellite.

Aims. We aim to develop a machine learning method based on fully convolutional neural networks called the CMB extraction neural network (CENN) in order to extract the CMB signal in total intensity by training the network with realistic simulations. The frequencies we used are the *Planck* channels 143, 217, and 353 GHz, and we validated the neural network throughout the sky and at three latitude intervals: $0^\circ < |b| < 5^\circ$, $5^\circ < |b| < 30^\circ$, and $30^\circ < |b| < 90^\circ$. Moreover, we used neither Galactic nor point-source (PS) masks.

Methods. To train the neural network, we produced multi-frequency realistic simulations in the form of patches of 256×256 pixels that contained the CMB signal, the Galactic thermal dust, cosmic infrared background, and PS emissions, the thermal Sunyaev–Zel’dovich effect from galaxy clusters, and instrumental noise. After validating the network, we compared the power spectra from input and output maps. We analysed the power spectrum from the residuals at each latitude interval and throughout the sky, and we studied how our model handled high contamination at small scales.

Results. We obtained a CMB power spectrum with a mean difference between input and output of $13 \pm 113 \mu\text{K}^2$ for multipoles up to above 4000. We computed the residuals, obtaining $700 \pm 60 \mu\text{K}^2$ for $0^\circ < |b| < 5^\circ$, $80 \pm 30 \mu\text{K}^2$ for $5^\circ < |b| < 30^\circ$, and $30 \pm 20 \mu\text{K}^2$ for $30^\circ < |b| < 90^\circ$ for multipoles up to above 4000. For the entire sky, we obtained $30 \pm 10 \mu\text{K}^2$ for $l \leq 1000$ and $20 \pm 10 \mu\text{K}^2$ for $l \leq 4000$. We validated the neural network in a single patch with strong contamination at small scales, obtaining a difference between input and output of $50 \pm 120 \mu\text{K}^2$ and residuals of $40 \pm 10 \mu\text{K}^2$ up to $l \sim 2500$. In all cases, the uncertainty of each measure was taken as the standard deviation.

Conclusions. The results show that fully convolutional neural networks are promising methods for performing component separation in future CMB experiments. Moreover, we show that CENN is reliable against different levels of contamination from Galactic and PS foregrounds at both large and small scales.

Key words. techniques: image processing – cosmic background radiation – submillimeter: general

1. Introduction

The cosmic microwave background (CMB) is the relic emission from the primordial Universe at an age of about 380 000 yr, an epoch called recombination (Schneider 2006). At first approximation, it is homogeneous and isotropic. However, it has small deviations in intensity (temperature) and polarisation from point to point in the sky. These are called anisotropies. They were firstly detected by the Cosmic Background Explorer satellite (COBE; Smoot et al. 1992), and the detection allow fitting the cosmological parameters and cemented the gravitational instability paradigm within a cold dark matter model. After this, the *Wilkinson Microwave Anisotropy Probe* (WMAP; Bennett et al. 2003) showed that the fluctuations are predominantly adiabatic. Finally, *Planck* (Planck Collaboration I 2020) obtained a more precise CMB signal for both intensity and polarisation with its angular resolution, which is higher than that the previous instruments.

The CMB anisotropies are one of the most important fields of research in modern cosmology. They are divided into two types: primary anisotropies, which are temperature fluctuations that originated in the recombination era, and secondary anisotropies, which are due to distortions of photons through their propagation along the Universe. These anisotropies can be studied by estimating the CMB power spectra, which describes the amplitude of the fluctuations $l(l+1)C_l$ on an angular scale $\theta \sim \pi/l = 180^\circ/l$.

Before recombination, photons and baryons are a tightly coupled photon-baryon fluid (Peebles & Yu 1970) that oscillates relativistically, causing a first peak at $l_1 \sim 200$ and additional maxima at approximately integer multiples of l_1 , called acoustic peaks. These peaks provide fundamental information about cosmology and, in particular, about the initial conditions and the energy contents of the Universe. Therefore, it is very important to measure them with the highest precision, implying a very precise recovery of the CMB signal. However, CMB measurements

are contaminated by other signals called foregrounds (Krachmalnicoff et al. 2016, 2018).

The process of cleaning the microwave maps in order to recover the CMB signal is called component separation (Baccigalupi 2003). It usually exploits multi-frequency sky maps to extract cosmological information. Our work focuses on *Planck* central wavelengths, where the foregrounds are divided into two categories: firstly, diffuse emission from our Galaxy, which is mostly dominated by thermal dust emission (Hensley & Draine 2021). This is the main contaminant at large angular scales. Secondly, extragalactic foregrounds such as proto-spheroidal galaxies forming the cosmic infrared background (CIB; Dole et al. 2006), radio galaxies (de Zotti et al. 2010), infrared (IR) dusty galaxies in the process of forming their stellar masses (Toffolatti et al. 1998), and also both thermal and kinetic Sunyaev-Zeldovich effects from galaxy clusters (Carlstrom et al. 2002).

Component separation was one of the central objectives of the *Planck* mission (Leach et al. 2008, 2016, 2016). The mission used four different methods. They were chosen because each of the selected method relies on a different class of algorithms (i.e. blind and non-blind methods, pixel-based, and harmonic-based methods). The first algorithm was NILC (Delabrouille et al. 2009), which makes minimum assumptions for the foregrounds in order to minimise the variance of the CMB using a wavelet on the sphere on the internal linear combination (ILC) algorithm. The second algorithm was SEVEM (Martínez-González et al. 2003), which is a template-based method that removes foreground contamination from the CMB-dominant bands from the lowest and highest frequency channels. The third algorithm was based on a Bayesian parameter estimation technique called Commander (Eriksen et al. 2006), and the fourth method was SMICA (Delabrouille et al. 2003), a parametric approach that allows fitting the amplitude and spectral parameters of the CMB and foregrounds. The main issue of the four methods is that they perform by masking both the Galactic plane and the point sources in order to avoid their contamination. Their final confidence mask in total intensity keeps 78% of the sky. Moreover, the implementation of these methods required years of research and developing in order to reach the current precise CMB measurement, with the introduction of notable improvements from the early to latest releases. It is the not always straightforward application of traditional methods that drives our interest in trying a different approach to component separation.

The amount of data in the past years has been produced a high impact in fields such as astrophysics and cosmology. It is necessary to develop automatic and reliable methods that can perform on the new data. The most promising methods are those based on machine learning (ML) because once they are trained, they can be validated on new data, providing results in a much shorter period of computational time with respect to the whole process of the traditional approaches. Some of the most important ML methods are artificial neural networks, mathematical approaches inspired by neuroscience. They can optimise models with non-linear behaviours. Some recent applications in cosmology were in the PS detection field, both in single-frequency (Bonavera et al. 2021) and in multi-frequency (Casas et al. 2022) approaches. Moreover, they have been used in the statistical study of the CMB for extending foreground models to sub-degree angular scales (Krachmalnicoff & Puglisi 2021) and to perform a foreground model-recognition for B-mode observations (Farsian et al. 2020). They have also been used in works similar to ours for CMB recovery and thermal dust cleaning (Petroff et al. 2020; Wang et al. 2022; an all sky approximation

for multipoles up to $l \sim 900$ in the first case, and in 2D patches for multipoles up to $l \sim 1500$ in the second). Another timely application is the model by Jeffrey et al. (2022), a moment network for polarised foreground removal using a single training image.

The aim of this work is to train a fully convolutional neural network, called the CMB extraction neural network (CENN), for the recovery of the CMB signal with realistic microwave sky simulations in total intensity to determine the potential of the method. We study polarisation data in an upcoming work. The typical application of fully convolutional networks (FCN) is in the Euclidean space, that is, to flat images. This clearly allows a CMB power spectrum estimation only down to angular scales allowed by the patch dimension. This approach might be directly applied to ground-based experiments that will observe limited regions of the sky (e.g. the Simons Observatory and CMB-S4; Ade et al. 2019; Abazajian et al. 2019). In the future, for full-sky coverage experiments (e.g. the LiteBIRD satellite; Sugai et al. 2020), it would be useful to apply FCN to the Hierarchical Equal Area Latitude Pixelization scheme (HEALPix; Górski et al. 2005) using for example the approach by Krachmalnicoff & Tomasi (2019).

The outline of this paper is the following. Section 2 covers the generation of the simulated maps. Section 3 describes our method. The results are explained in Sect. 4, and our conclusions are discussed in Sect. 5.

2. Simulations

To train CENN, we used realistic simulations of the microwave sky. Each simulation was formed by sky patches at the central channels of the *Planck* mission: 143, 217, and 353 GHz. Each patch has an area of 256×256 pixels in order to have good statistics with the shortest possible computation time. The pixel size was 90 arcsec, a number close to the 1.72 arcmin used in the *Planck* mission, which correspond to $N_{\text{side}} = 2048$ in the HEALPix all-sky pixelisation schema. The maps were randomly simulated at any position in the sky, without the use of PS or Galactic masks, and they were selected in such a way that the probability of having two exact sky maps was negligible. We simulated different datasets: one formed by 60 000 simulations used for training, another one with 6000 simulations to evaluate the model during that training, and validation datasets with simulations apart from the train one in order to determine how CENN performs on new data. The first dataset consisted of 6000 simulations of patches centred at any position in the sky (all-sky case). Other three datasets of 2000 simulations each were built with patches centred at three different latitude intervals: $0^\circ < |b| < 5^\circ$ (inner Galactic region), $5^\circ < |b| < 30^\circ$ (intermediate Galactic region), and $30^\circ < |b| < 90^\circ$ (extragalactic region).

First, the CMB signal was projected from all-sky simulations available in the *Planck* Legacy Archive (PLA¹), which were produced using the *Planck* Sky Model software (PSM; Delabrouille et al. 2013). The CMB signal used in the train and test datasets was taken from a different PLA simulation than we used in the validation dataset.

As Galactic foregrounds, we assumed that only thermal dust emission contaminate the CMB, and we considered it by adding FFP10 PLA simulations from the third *Planck* release. Although it is a good modelling at first approximation, some studies predict significant variations of the two modified blackbody model of the simulations used in this work (Boudet et al. 2005; Meny et al. 2007; Paradis et al. 2011), which directly affect the behaviour of

¹ <http://pla.esac.esa.int/pla/#home>

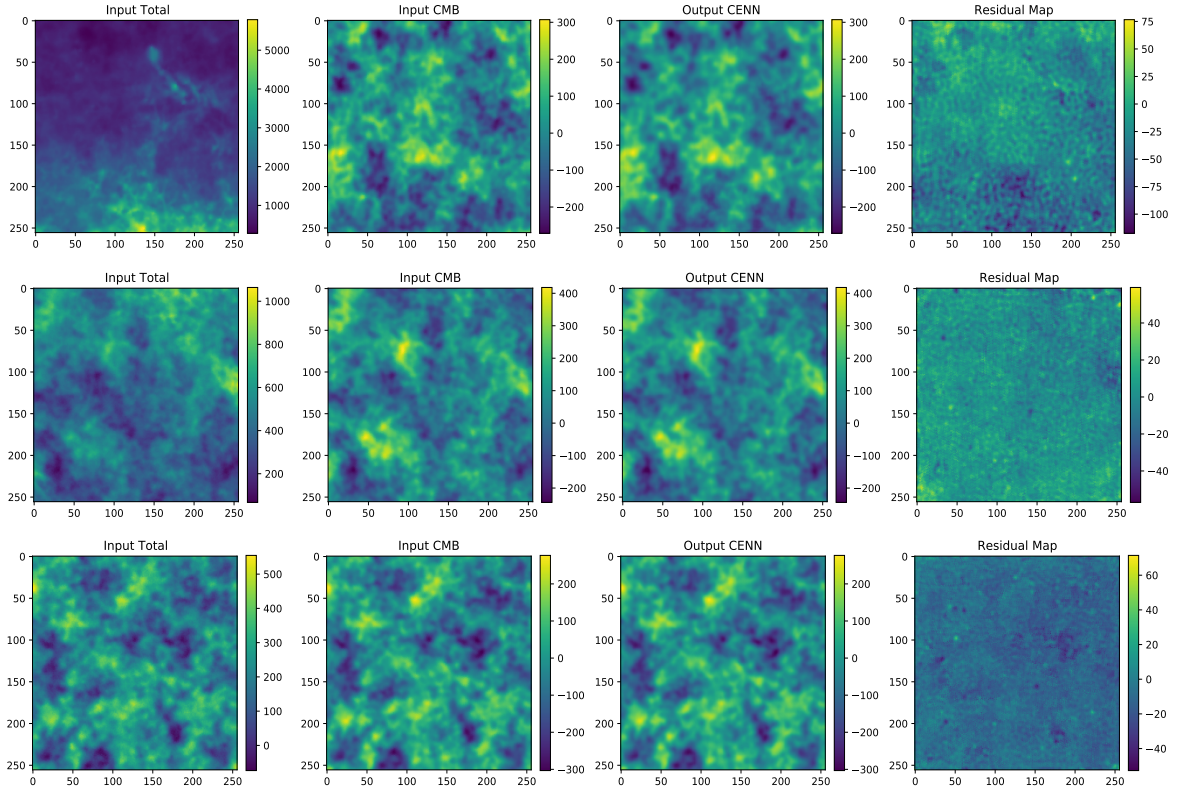


Fig. 1. Patches at $0^\circ < |b| < 5^\circ$ (top row), $5^\circ < |b| < 30^\circ$ (middle row), and $30^\circ < |b| < 90^\circ$ (bottom row) latitude intervals. They represent the maps with all the emissions, the CMB label, the neural network output, and the residual map computed as the difference between the second and the third column patches from left to right. The frequency is 217 GHz, and their temperature (in μK_{CMB}) is shown in the right vertical bars. The x and y labels are the pixel coordinates for each 256×256 patch.

dust at small scales, neglecting the real non-Gaussian structures of this foreground.

As extragalactic foregrounds, firstly, the thermal Sunyaev–Zel’dovich effect produced by galaxy clusters was also taken into account using FFP10 PLA simulations as well. Furthermore, point sources (PS) were added as two different populations: the first population consisted of radio galaxies, mainly radio quasars and BL Lacertae objects, commonly known as blazars. These were simulated in the central channel (i.e. 217 GHz) following the model by Tucci et al. (2011) with the software CORRSKY (González-Nuevo et al. 2005). The second population contained IR late-type galaxies, mainly starburst and local spiral galaxies (IRLT), which were simulated with CORRSKY in the central channel using the model by Cai et al. (2013) and normalised following the update by Negrello et al. (2013). Their spectral behaviour was considered following Casas et al. (2022) by assuming that their emission at 143 and 353 GHz varied as

$$S = S_0 \left(\frac{\nu}{\nu_0} \right)^\alpha, \quad (1)$$

where S is the flux density of the sources at each frequency ν , S_0 is the flux density in the central channel with frequency ν_0 , and α is the spectral index, which was randomly selected for each population using the distributions given in Planck Collaboration XXVI (2016).

Proto-spheroidal galaxies were also considered as the elements forming the cosmic IR background (CIB; Hauser & Dwek 2001), a main contaminant at high frequencies. They were simulated using the source number counts given by Cai et al. (2013),

the angular power spectrum obtained by Lapi et al. (2011), and the software CORRSKY. Their spectral behaviour was also considered following Eq. (1), assuming that they have the same spectral index distribution as the IRLT population.

Finally, instrumental noise was added to the total input maps in order to simulate the *Planck* HFI instrument. To do this, we used 0.55, 0.78, and $2.56 \mu\text{K}_{\text{CMB}} \text{ deg}$ values for 143, 217, and 353 GHz, respectively (Planck Collaboration IV 2020).

Each simulation was formed by two different sets of maps called the input total images and the labels. The input total images are multi-frequency realistic simulations of the microwave sky at 143, 217, and 353 GHz formed by the CMB added to other signals acting as contaminants. The labels are maps with only the CMB signal at 217 GHz. To project the maps into flat patches, the `gnomview` function of HEALPix framework was used. An example of simulations at different latitudes is shown in the first two columns of Fig. 1 for 217 GHz. The first column represents the total input images with both the CMB signal and the other contaminants, and the second column shows the CMB label maps.

3. Method

ML is the ability of a computer program to learn some class of tasks and performance measures from experience. Neural networks are one of the most popular ML methods. They are inspired by neuroscience and are generally used to learn non-linear parameters in a model. They are usually composed of connected layers with basic computing units called neurons, which adjust their weights and bias values in each training step.

One of the most popular applications of neural networks is to learn patterns in images to perform a detection using a special type of architecture called convolutional neural networks (LeCun et al. 1989). Image segmentation is an evolution of this task because it consists of classifying each pixel instead of the whole image. This allows separating one class from the others. Fully convolutional neural networks (FCN; Long et al. 2015) are used to perform image segmentation. They learn parameters by extracting the main features of an image through a set of convolutional blocks, while they make predictions through a set of deconvolutional blocks. Each convolutional block is formed by a layer that reads information from a matrix of input data, and after this, it performs convolutions in parallel. A set of activation functions linearize the convolutions, and a pooling layer aggregates information by grouping neighbouring pixels using their local statistics. Deconvolutional blocks have a similar architecture as convolutional blocks, but in this case, the linearized information is deconvolved, resulting in a segmentation of the class that we wish the neural network to learn.

Like other neural networks, FCN learn by performing optimisation through minimising a loss function located at the end of their architecture, a process called forward propagation. To drive the loss function to a minimum, a gradient-based optimiser is used on a small set of samples from the training dataset. This is called a minibatch. When the gradient is updated, the model parameters, which adjust some components of the architecture known as the filters, the weights, or the bias, are also updated. This process is called back propagation (Rumelhart et al. 1986). It results in a theoretically better knowledge of the information read in each epoch, which is the interval within which the entire train dataset flows forward and backward in the neural network.

3.1. Model

On the one hand, our model is based on the physical assumptions explained in Eriksen et al. (2006, 2008). On the other hand, it follows the mathematical behaviours of CNN and FCN, which are extensively revised in Rumelhart et al. (1986); Lecun et al. (1998) and Goodfellow et al. (2016). In this subsection, we explain the physical concepts of our model. The mathematical concepts are summarised in Appendix A.

The model consists of patches x_ν of the microwave sky at frequency ν (143, 217, and 353 GHz in this work), each one composed of a linear combination of c astrophysical components with e_ν emission maps convolved with the instrumental beam of the experiment i_ν , plus instrumental noise n_ν , that is,

$$x_\nu = i_\nu * e_\nu + n_\nu, \quad (2)$$

where the asterisk denotes convolution. Moreover, the emission maps are a superposition of different astrophysical components, each one with a particular non-linear spectral behaviour. Their superposition can be written as a linear combination of each foreground,

$$e_\nu = \sum_c A_{c,\nu} s_c, \quad (3)$$

where A is called the mixing matrix, and s is the vector of components. Therefore, Eq. (2) becomes

$$x_\nu = i_\nu * \left(\sum_c A_{c,\nu} s_c \right) + n_\nu. \quad (4)$$

In the first epoch, the parameters of the model are randomly initialised, and the neural network reads a tensor formed by a matrix of $N_{\text{pix}} \times N_{\text{pix}} \times N_\nu$ in its first block, where $N_{\text{pix}} = 256$ is the number of pixels, and N_ν is the number of frequency channels. Then, the spectral dimension of the input data is considered by estimating a multidimensional stride convolution (Goodfellow 2010),

$$H_{\nu,i,j} = \sum_{k,m} W_{k,m,i} V_{s \circ \nu + k,m,j}, \quad (5)$$

where H are the hidden units (i.e. related to the resulting outputs in the first convolutional block), which index each spectral dimension ν within the feature map i for sample j . W is the kernel, which connects the weights between i and j for each spectral channel. V are the visible units (i.e. related to the input images), which have the same format as H . The subindex s represents a vector of strides². The open circle represents the elementwise product. The next convolutional blocks have similar operations, but they consider the information computed in their previous convolutional block as inputs.

In the deconvolutional blocks, the spectral dimensions of the data are considered by computing a multidimensional transpose of strided convolution

$$R_{q,m,j} = \sum_{\nu,k|s \circ \nu + k = q} \sum_i W_{k,m,i} H_{\nu,i,j}, \quad (6)$$

where R are the hidden units in the deconvolutional block, and H are the visible units (i.e. the information from the previous block), which index each spectral dimension q within the feature map m for the sample j . The vertical line denotes the modulus operator³. The other terms are those defined in Eq. (5). The kernel values W in all convolutional and deconvolutional blocks are updated on each epoch after minimising the mean-squared error loss function,

$$\text{MSE} = \frac{1}{2} |y - y'|^2, \quad (7)$$

where y is the value predicted by the model, and y' is the true value, which in our case is the map with the true CMB emission. This function is located at the end of the architecture, and after computing it, a gradient is then estimated. When the gradient is computed, the parameters for each layer become updated as explained in Appendix A. In the last deconvolutional block, all the spectral information of its previous block is deconvolved into a single channel in order to predict the CMB signal at $\nu_0 = 217$ GHz, that is, the network computes the quantity $R_{q=\nu_0,m,j}$ by indexing $H_{\nu,i,j}$ in Eq. (6) over all the filters on the previous deconvolutional block, producing the quantity

$$\tilde{x}_{\nu_0} = i_{\nu_0} * e_{\text{CMB},\nu_0} + \tilde{n}, \quad (8)$$

that is, we obtain a map \tilde{x}_{ν_0} at ν_0 frequency composed of the e_{CMB,ν_0} CMB emission convolved with i_{ν_0} (i.e. the instrumental beam of the experiment) plus \tilde{n} , which is generalised noise due to deconvolution.

² They are the number of pixels shifts over the input matrix. Because in our case, $s = 2$ in every layer, the network move the filters to two pixels at a time in each layer.

³ It is used to sum over the correct set of values for ν and k .

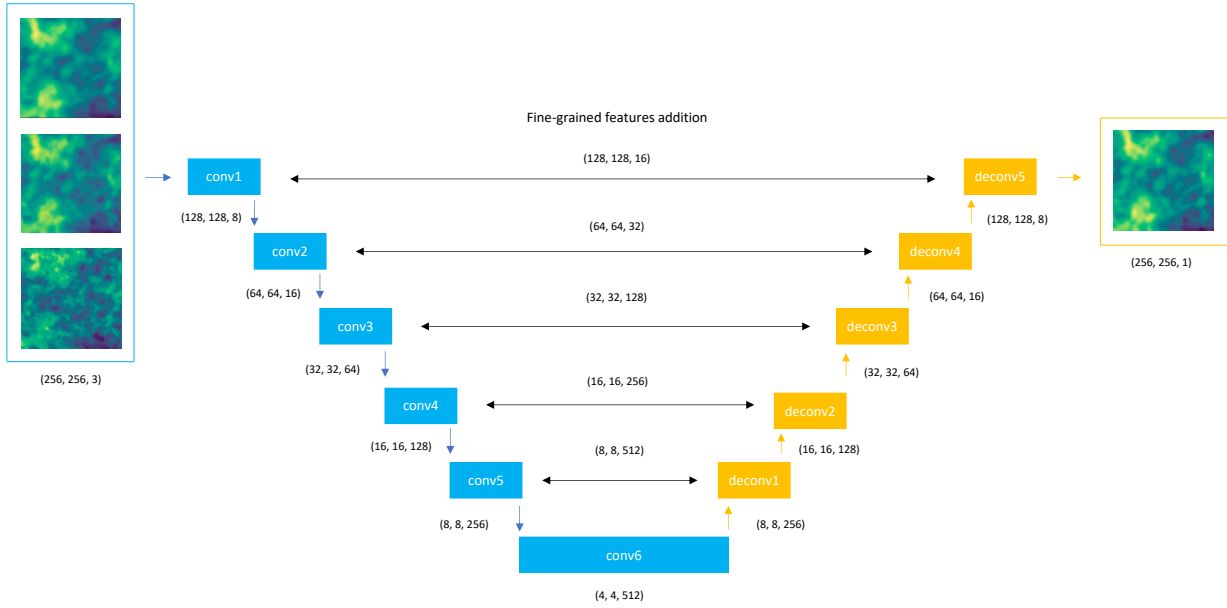


Fig. 2. Architecture of CENN. It has a convolutional block that produces 8 feature maps. After this, the space dimensionality increases to 512 feature maps through five more convolutional blocks. These layers are connected to deconvolutional blocks that decrease the space dimensionality to one feature map in the last deconvolutional block. Fine-grained features are added from each convolution to its corresponding deconvolution.

3.2. Architecture

Our topology is an FCN based on the U-Net architecture (Ronneberger et al. 2015). The inputs of the network are three patches (256×256 pixels) with the total sky emission, one for each frequency at the same position. The output is a single patch, with the same size as the input patches, containing the recovered CMB signal. The label for minimising the loss function is the true CMB signal on each simulation. The network was trained during 500 epochs with 60 000 simulations divided into minibatches of 32.

The architecture of CENN is shown in Fig. 2 and is detailed as follows: firstly, it has a set of six convolutional blocks, all of them formed by convolutional and pooling layers with 8, 2, 4, 2, 2, and 2 kernels of sizes of 9, 9, 7, 7, 5, and 3, respectively. The number of filters is 8, 16, 64, 128, 256, and 512, respectively. They have a subsampling factor of 2. The padding type “Same” is added in all the layers in order to add some space around the input data or the feature map, to deal with possible loss in width and/or height dimension in the feature maps after having applied the filters. The activation function is leaky ReLU in all the layers. Secondly, the FCN has a set of six deconvolutional blocks that are formed by deconvolutional and pooling layers with 2, 2, 2, 4, 2, and 8 kernels of sizes of 3, 5, 7, 7, 9, and 9, respectively. The number of filters is 256, 128, 64, 16, 8, and 1. Their subsampling factor is 2. The padding type “Same” is also added in all the layers, and the activation function is the leaky ReLU.

In order to take fine-grained features in the image into account (Wei et al. 2021), we have added layers to the architecture that connect the convolutional and deconvolutional blocks. This addition doubles the space of the feature maps before each deconvolutional block. The main goal for these layers is to help the network to predict low-level features with the deconvolutional blocks by taking high-level features into account that are inferred by the convolutional blocks. After each fine-grained layer, a deconvolutional block is connected in order to predict patterns that are dependent on the high and low-level features. In our model, the addition of these layers is related to the task

of predicting small-scale regions of the CMB signal by taking already inferred large-scale structures into account. An example of CENN output patches is shown in the third column of Fig. 1.

4. Results

After CENN was trained, we validated it using a dataset of 6 000 simulations of random patches in the whole sky and using three different validation datasets of 2000 simulations each in different latitude ranges: the first dataset at $0^\circ < |b| < 5^\circ$, the second dataset at $5^\circ < |b| < 30^\circ$, and the third dataset at $30^\circ < |b| < 90^\circ$. All the validation datasets consisted of simulations that were different from the training dataset and, in particular, the CMB signal was from another PLA simulation. The analysis of the results is split into three parts: firstly, we study the CMB power spectrum in the whole sky in Sect. 4.1. Secondly, we analyse both the mean residuals and the residuals of the mean patch in the all sky case and in the three latitude intervals in Sect. 4.2. Thirdly, we estimate the levels of contamination at small scales in Sect. 4.3.

4.1. Power spectra

We compare in the top panel of Fig. 3 the power spectrum of the true CMB maps with the power spectrum from the outputs of CENN, validated over the whole sky. More precisely, we compare the average of the power spectra computed on the input (blue line) and output (red line) sky patches. The power spectrum was rebinned using a step of 50 up to $l = 1000$ and a step of 200 above $l > 1000$. The standard deviation of each bin for the input and output patches is also shown (blue and red areas, respectively), which is considered the uncertainty of each quantity. Furthermore, the bottom panel shows the difference between the input and the output average power spectra (black line) and its uncertainty (grey area).

We recover the CMB power spectrum with a mean difference between input and output of $13 \pm 113 \mu\text{K}^2$ for all multipoles

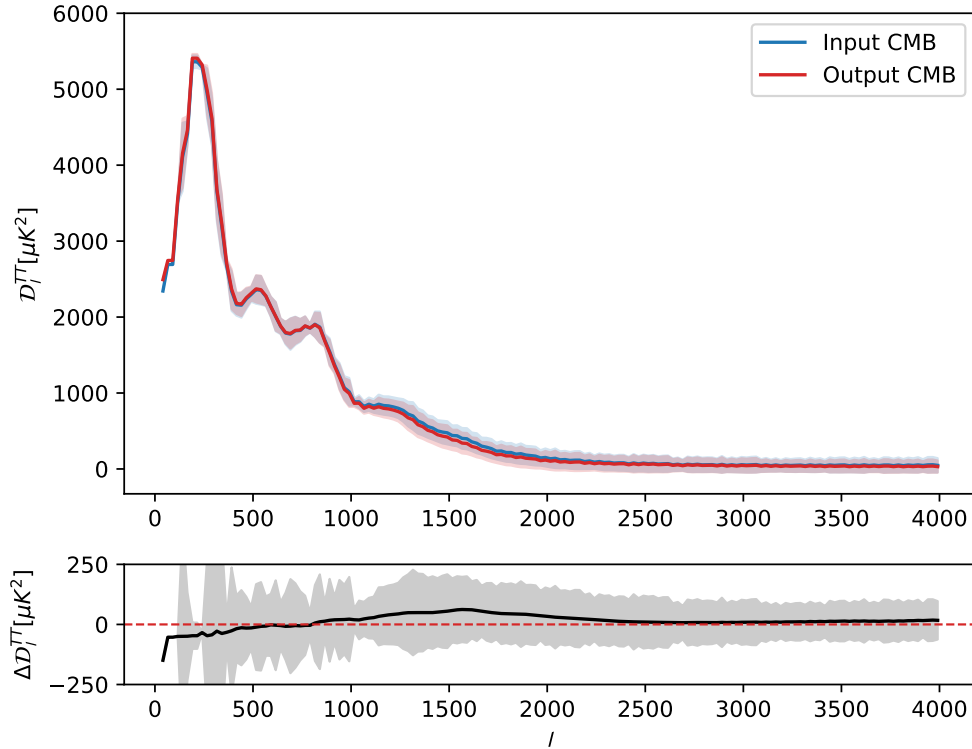


Fig. 3. Mean CMB power spectrum comparison computed over the entire validation dataset. The true CMB is represented as a blue line, and the output from CENN is shown as a red line. The corresponding uncertainties, computed as the standard deviation of each bin, are the blue and red areas, respectively. The difference between input and output is plotted in the bottom panel as a black line, and its uncertainty is the grey area.

and $-18 \pm 150 \mu\text{K}^2$ for multipoles up to $l \sim 1000$. Petroff et al. (2020) recovered the CMB with a difference of $3.8 \mu\text{K}^2$ up to the same multipoles after masking map pixels with a predicted standard deviation of $>50 \mu\text{K}^2$. This was a per-pixel error comprising both statistical and model uncertainties. This was not the case in this work. However, in our model, we tested CENN without the use of any pixel mask. At the first multipoles, the CMB cannot be properly recovered because we work with patches of the sky. We therefore have poor statistics at the largest angular scales. Moreover, between $l \sim 1000$ and $l \sim 2500$, we obtain a CMB power spectrum with a difference of $20 \pm 100 \mu\text{K}^2$ with respect to the input data, while Petroff et al. (2020) obtained a mean value above $100 \mu\text{K}^2$. This result is interesting mostly because the traditional component separation methods applied to *Planck* data (Planck Collaboration IV 2020) also failed at these multipoles, and they needed to mask the Galactic region and the brightest PS to avoid this contamination. Our method can be used to recover the CMB in the entire sky, without applying any sky mask at all. This means that on the one hand, we successfully recover the CMB signal even in the region that is mostly affected by galactic emission and generally avoided by traditional component separation methods. On the other hand, because we do not need to mask the PS, we successfully recover the signal also at large angular scales. Therefore, it is quite likely that our network will recover the CMB signal with satisfying accuracy and reliability in future CMB experiments with better angular resolution than *Planck*.

4.2. Residuals

In this section, we study the mean power spectrum of the residual maps (i.e. the difference between input CMB and CENN output patches) and the power spectrum of the mean residual patch.

Firstly, we computed the mean and the root mean-squared (RMS) of the residual patch for the whole sky and for the three different latitude intervals. Table 1 lists the mean value (top row) and the RMS value (bottom row) for $0^\circ < |b| < 5^\circ$ (second column), $5^\circ < |b| < 30^\circ$ (third column), $30^\circ < |b| < 90^\circ$ (fourth column), and the whole sky (fifth column).

The mean value is $5.072 \mu\text{K}$ for the inner Galactic region. This is close to the value for the extragalactic region, $5.238 \mu\text{K}$. At intermediate latitudes, the mean value is $-0.026 \mu\text{K}$. Finally, when the whole sky is taken into account, we obtain a mean value of $-0.83 \mu\text{K}$. For the RMS, we obtain $7.732 \mu\text{K}$ for the Galactic region and $7.536 \mu\text{K}$ for the intermediate region, which is a lower number than was reported in Leach et al. (2008) for some of the component separation methods used in *Planck*. More precisely, they presented an RMS value of about 15 and $13 \mu\text{K}$ for SMICA and 18 and $16 \mu\text{K}$ for SEVEM for the latitude intervals, both obtained using realistic simulations.

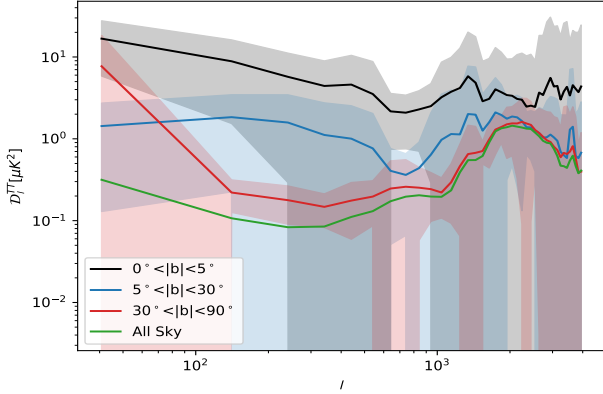
In Fig. 4, we compare the power spectrum of the mean residual patch, that is, firstly we computed the mean residual patch for each validation dataset, and after this, we estimated its power spectrum. This type of study is not only useful for a comparison against other segmentation methods, but also allow us to improve our statistical knowledge of the method, for example by identifying possible systematic biases such as artefacts near the borders, which are common for methods working with patches.

The residuals are lower than $15 \mu\text{K}^2$ for multipoles up to $l \sim 4000$. In the inner Galactic region, the strong contamination implies $7 \pm 25 \mu\text{K}^2$, that is, only a higher residual for all multipoles with respect to the other cases. This is an important achievement because this region is usually avoided when more traditional component separation methods are used.

Although it is still a region with high Galactic contamination, our method performs well in the intermediate case, obtaining

Table 1. Mean value of the residual patch (top row) and its root mean-squared value (bottom row) for each sky latitude and for the whole sky.

	$0^\circ < b < 5^\circ$	$5^\circ < b < 30^\circ$	$30^\circ < b < 90^\circ$	All sky
Mean residual patch (μK)	5.072	-0.026	5.238	-0.83
RMS residual patch (μK)	7.732	7.536	6.85	7.372


Fig. 4. Power spectrum of the mean residual patches for the whole sky (green continuous line) and for $0^\circ < |b| < 5^\circ$ (continuous black line), $5^\circ < |b| < 30^\circ$ (continuous blue line), $30^\circ < |b| < 90^\circ$ (continuous red line). The corresponding coloured areas are the standard deviation of each bin.

residuals of about $2 \pm 10 \mu\text{K}^2$ for multipoles up to $l \sim 4000$. As in the inner Galactic case, this region was usually partially masked by the component separation methods used in *Planck*. Finally, in the extragalactic region, the power spectrum of the mean residual patch shows the expected behaviour: it has lower values overall because in this region, the Galactic contamination is significantly lower than in the other cases.

The residual behaviour remains almost constant even at high multipoles. More precisely, for $l > 1000$, we obtain values of about $3 \pm 11 \mu\text{K}^2$ for $0^\circ < |b| < 5^\circ$, $1 \pm 2 \mu\text{K}^2$ for $5^\circ < |b| < 30^\circ$, $0.5 \pm 0.6 \mu\text{K}^2$ for $30^\circ < |b| < 90^\circ$, and $0.8 \pm 1.5 \mu\text{K}^2$ for the whole sky, which are still a low values for the residuals. This suggests that in principle, masking might not be needed with our network. This gives us the advantage that we do not have to deal with the holes left by PS masks in CMB reconstructed maps. The performance at these multipoles is probably due to the fact that the network is trained to segmentate diffuse signals, which have completely different characteristics than the point-like signals. Then, the PS contamination is relatively irrelevant while reconstructing the CMB signal with CENN.

In Fig. 5, we compare the mean power spectrum (black line) of each residual patch in the validation datasets with the input and recovered CMB signal and the level of contamination of each foreground in the simulations (the thermal Sunyaev-Zel'dovich effect appears in the 143 and 353 GHz input simulations during validation, but because the output is at 217 GHz, the contribution is neglected). The corresponding uncertainties for each quantity (i.e. the standard deviation, as in the above subsection) are represented as coloured areas.

The residuals depend strongly on contamination levels. In the Galactic plane region, we obtain residuals of about $700 \pm 60 \mu\text{K}^2$. In the other Galactic region, the residuals decrease to $80 \pm 30 \mu\text{K}^2$. However, the contamination levels in these regions are about three and two orders of magnitude higher,

respectively. In the extragalactic region, the impact of contaminants is not as high as in the other cases, and this is also true for the residuals, which are above $30 \pm 20 \mu\text{K}^2$. Finally, with random simulations in the whole sky, we obtain residuals of about $20 \pm 10 \mu\text{K}^2$, although the contamination levels are almost two orders of magnitude.

At small scales, the residuals show a slight bump in all cases, which is also visible in Fig. 4. It is due mainly to extragalactic foregrounds. However, it should be noted that for multipoles $l > 2500$, where the *Planck* beam caused the foregrounds to decrease, the residuals remain constant and at lower levels than at large scales. The continuous brown line in Fig. 5 shows that the instrumental noise is strong in this region, however. This is the first indication that these types of ML models might be able to handle instrumental effects. This behaviour of CENN should be studied in detail and tested in future works because it is crucial to constrain both instrumental and systematics effects for primordial B mode detection.

4.3. Contamination at small scales

In this subsection, we focus on studying the performance of the network against contamination at small scales by analysing a specific patch that was selected for its much higher visual contamination at small scales with respect to the average case. The patch is represented in the top panel in Fig. 6 (units in μK): from left to right, the total patch with all the contaminants and the CMB added together, the input CMB for the same patch, the network output, and the residual (the difference between the input CMB and the output) patches.

The small-scale structure can be seen in the pixels [237, 41], [100, 70], [169, 148], and [80, 185], mostly in the residual patch. Despite the high contamination, the residuals are relatively low. The mean value is $-0.72 \pm 15.52 \mu\text{K}$. The highest and lowest values that can be also seen by eye in the residual patch correspond to point-like objects that contaminate the network output at small scales. In particular, the pixels at [237, 41] are the brightest in the input total patch with a value of about $1000 \mu\text{K}$. This point-like structure is still present in the output patch with a value lower than $300 \mu\text{K}$ and is clearly visible in the residual patch with absolute values of about $100 \mu\text{K}$. These numbers suggests that the network still performs well even for this problematic patch because the contamination from small scales is notably reduced.

In the bottom panel in Fig. 6 we also compare the power spectra of the input, output, and residual patches of this sky area. The rebinning has a step of 50 until $l < 1000$, and a step of 200 above $l = 1000$. In the top panel, we represent the input and output CMB power spectra as blue and red lines, respectively, with their uncertainties, estimated as the standard deviation, plotted as blue and red areas, respectively. In the middle panel, we show the difference between input and output power spectra. The grey area is the uncertainty. In the bottom panel, we plot the power spectrum of the residual patch, and the green area is the uncertainty.

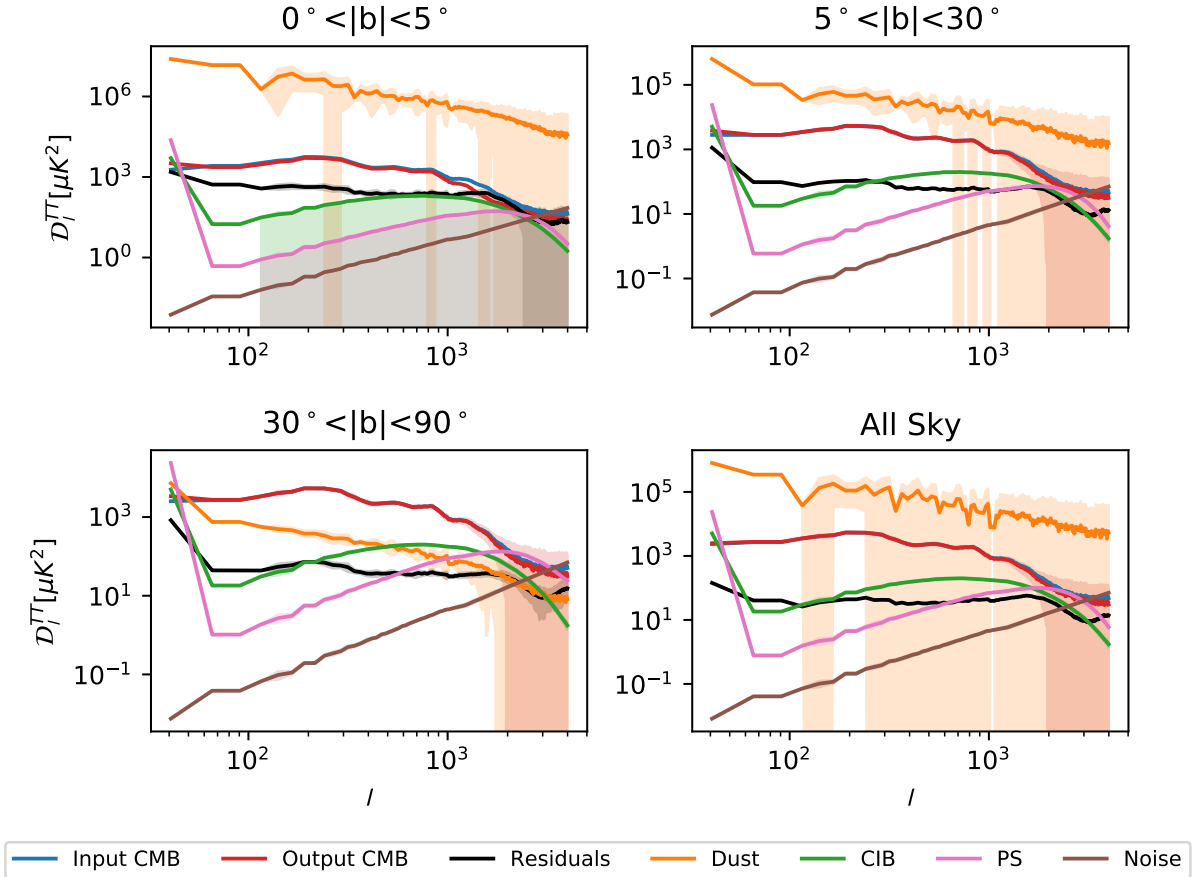


Fig. 5. Mean power spectrum of the residuals (black) and its uncertainty (grey area) for $0^\circ < |b| < 5^\circ$ (top left panel), $5^\circ < |b| < 30^\circ$ (top right panel), $30^\circ < |b| < 90^\circ$ (bottom left panel) and for the whole sky (bottom right panel) against the input and output CMB (blue and red lines, respectively) and their uncertainties (blue and red areas, respectively). In all cases, the uncertainties are the standard deviation of each bin. The contribution of each foreground in the simulations through their power spectra is also plotted. The thermal dust is represented by the orange line, the PS and CIB by the pink and green lines, respectively, and the instrumental noise is shown by the brown line.

CENN recovers the CMB power spectrum with a mean difference between 0 and $30 \mu\text{K}^2$ (black line) for the entire range of multipoles. In particular, between $l = 1200$ and $l = 1800$, the difference between input and output is higher, with a value of about $50 \mu\text{K}^2$. The estimated power spectrum of the residuals (green line) at the same multipoles has $40 \mu\text{K}^2$, decreasing at higher multipoles.

When our network is applied without any mask, the residual PS slightly affects the network output, but mainly at multipoles around 2000. The residuals at higher multipoles are negligible, and no residual contamination is related to the instrumental noise, at least at the *Planck* noise level. Moreover, we tested the performance of the network in patches with five and ten times higher instrumental noise than the *Planck* noise, finding that the performance at high multipoles is similar in the three cases, and the error is almost negligible. This behaviour is very different from that of the traditional component separation methods, for which above a certain multipole, $l > 2500$, the power spectrum is completely dominated by the PS and instrumental noise residuals, even after a masking procedure.

5. Conclusions

We developed a new component separation method based on artificial neural networks for future CMB experiments. More precisely, we trained a fully convolutional neural network called

CENN with realistic simulations of the central channels of the *Planck* mission to extract the CMB signal in total intensity from the other foregrounds, which are thermal dust from our Galaxy, the CIB signal, the thermal Sunyaev–Zel’dovich effect, the contribution from PS (radio and IR late-type galaxies), and instrumental noise. The frequencies are the 143, 217, and 353 GHz channels.

To train CENN, we used 60 000 realistic simulations. Each simulation had three image patches with an area of 256×256 pixels containing all the emissions and an additional map with only the CMB contribution at 217 GHz, which was used to learn this signal by the neural network via the optimiser and the loss function. In addition to this dataset, we used 6000 simulations to test the network during the training. After the training, we validated the network with different simulations, which had patches of the CMB signal extracted from a different PLA simulation than the one used for training the network. We developed four different validation datasets: one formed by 6000 random simulations at all sky, and another three with 2000 simulations each set at $0^\circ < |b| < 5^\circ$, $5^\circ < |b| < 30^\circ$, and $30^\circ < |b| < 90^\circ$ degrees of latitude. After validating CENN, we analysed the results by rebining the data and computing the mean power spectra of the input and output CMB from the neural network, and the difference between them, in all cases considering the standard deviation as the uncertainty of each bin. We obtain a mean difference of $13 \pm 113 \mu\text{K}^2$ for multipoles up to above 4000.

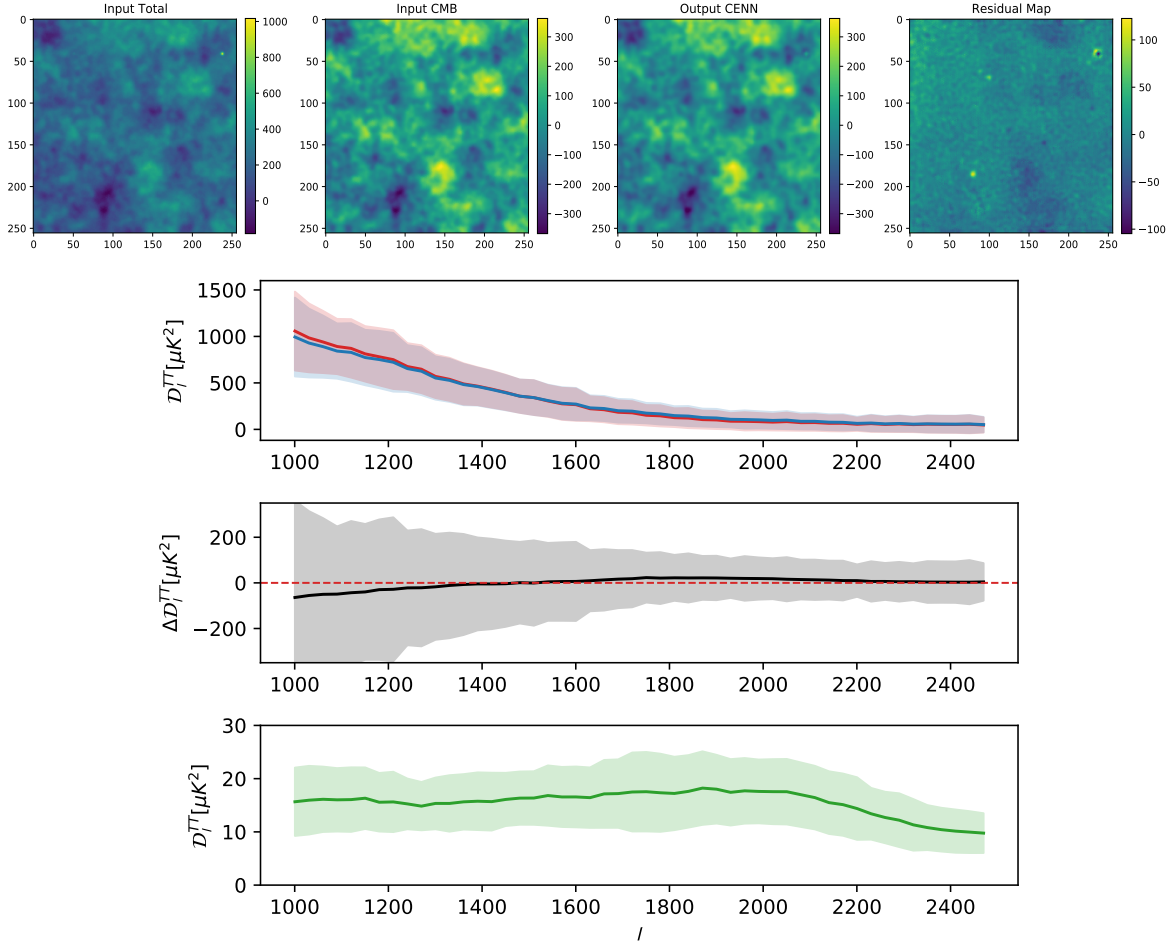


Fig. 6. Analysis of a patch highly contaminated at small scales. *Top panel:* sky patch selected with high contamination at small scales. The patches represent *from left to right* all the emission, the CMB label, the neural network output, and the residual map computed as the difference between the second and third columns. The frequency for all patches is 217 GHz, and their temperature (in μK_{CMB}) is shown in the right vertical bars. *Bottom panel:* CMB power spectrum comparison computed over the patch shown in the *top panel*. The true CMB is represented in the top panel as a blue line, and the output from CENN is shown as a red line. The corresponding uncertainties (i.e. the standard deviation) are shown as blue and red areas, respectively. The difference (ΔD_l^{TT}) between input and output is plotted in the *middle panel* as a black line, and its uncertainty is the grey area. The power spectrum of the residual map is represented as a green line in the *bottom panel*. The green area is its uncertainty (also the standard deviation).

We computed the mean power spectrum of the residuals (i.e. the difference between input and output CMB), obtaining $700 \pm 60 \mu K^2$ for $0^\circ < |b| < 5^\circ$, $80 \pm 30 \mu K^2$ for $5^\circ < |b| < 30^\circ$, and $30 \pm 20 \mu K^2$ for $30^\circ < |b| < 90^\circ$. We also analysed the mean and standard deviation patches for each latitude and for the whole sky. We obtain a RMS value of $7.732 \mu K$ for $0^\circ < |b| < 5^\circ$, $7.536 \mu K$ for $5^\circ < |b| < 30^\circ$, $6.85 \mu K$ for $30^\circ < |b| < 90^\circ$, and $7.372 \mu K$ for all sky. Furthermore, we computed their power spectra, obtaining $7 \pm 25 \mu K^2$ for $0^\circ < |b| < 5^\circ$, $2 \pm 10 \mu K^2$ for $5^\circ < |b| < 30^\circ$, and $2 \pm 3 \mu K^2$ for $30^\circ < |b| < 90^\circ$.

Finally, we studied the performance of CENN in recovering the CMB signal in a patch with high contamination at small scales. We obtain a value of $50 \pm 250 \mu K^2$ for the difference between input and output CMB power spectra, and residuals of about $40 \pm 10 \mu K^2$ at multipoles between $l = 1200$ and $l = 1800$.

These results show that our model is reliable not only at the multipoles analysed in *Planck* ($l \leq 2500$), but also for higher multipoles. Therefore, it seems to be a promising model for extracting the CMB signal for future experiments with higher angular resolution than *Planck*.

Moreover, it has other advantages: the first advantage is that our model is evaluated in the whole sky without the use of any confidence mask to avoid strong Galactic contamination regions. Therefore, a more realistic CMB map might be obtained with respect to the inpainted version at the Galactic plane used in *Planck*. The second advantage is that the PS or galaxy cluster contamination is very small and affects the output power spectrum only in a short multipole range around $l \sim 2000$. These results can easily be upgraded using a dedicated masking scheme or providing information about the PS and galaxy clusters during the training. These potential upgrades are beyond the scope of the current paper and will be considered for the application of the network to real data.

Both conclusions about the reliability of CENN with the foregrounds at large and small scales might be a first indication that the network might be able to separate the CMB from the foregrounds by learning their non-Gaussian structures (Coulton & Spergel 2019), which could be an advantage of this approach with respect to the traditional approaches. Efforts to study the network performance with these non-Gaussian structures will

be crucial in order to detect primordial B-modes on polarisation data. However, this is beyond the scope of this work and will be considered in future developments. Moreover, our results suggest that CENN might be able to perform a component separation well even in single-frequency maps, which is worth to be tested in the future as well. Lastly, when it is trained, our model can extract the CMB in microwave maps almost immediately, which is fundamental for future satellites with high amounts of data.

As we explained in Sect. 2, we have trained CENN with realistic simulations from the Planck Sky Model. However, these simulations do not take the non-Gaussian behaviours of the foregrounds of the microwave sky properly into account, especially at the frequencies analysed in this work, where thermal dust is the most significant foreground and it is known to have a strong non-Gaussian behaviour. Therefore, our results could be biased by this approximation of the microwave sky when validating it with real data, even if, as commented above, our network seems to rely on such non-Gaussian structures of the different emissions to recover the CMB signal. Because this issue is one of the major limitations of this kind of ML models, the network performance in total intensity and polarisation must be further tested in the future by training it with more realistic simulations, for example the outputs by ForSE (the generative adversarial neural network developed by Krachmalnicoff & Puglisi 2021), which can include small-scale non-Gaussian features extended to 12 arcmin.

Acknowledgements. We warmly thank the anonymous referee for the very useful and constructive comments on the original manuscript. J.M.C., L.B., J.G.N., M.M.C. and D.C. acknowledge financial support from the PGC 2018 project PGC2018-101948-B-I00 (MICINN, FEDER). J.M.C. also acknowledges financial support from the SV-PA-21-AYUD/2021/51301 project. M.M.C. also acknowledges to be granted by PAPI-20-PF-23 and PAPI-21-PF-04 (Universidad de Oviedo). C.B. acknowledges support from the ASI COSMOS and LiteBIRD Networks (cosmosnet.it), as well as by the INFN INDARK and LiteBIRD grants. J.D.S., M.L.S. and J.D.C.J. acknowledge financial support from the I+D 2017 project AYA2017-89121-P and support from the European Union's Horizon 2020 research and innovation programme under the H2020-INFRAIA-2018-2020 grant agreement no. 210489629. This research has made use of the python packages `matplotlib` (Hunter 2007), `Keras` (Chollet et al. 2015), and `Numpy` (Oliphant 2006), also the `HEALPix` (Górski et al. 2005) and `healpy` (Zonca et al. 2019) packages.

References

- Abazajian, K., Addison, G., Adshead, P., et al. 2019, CMB-S4 Decadal Survey APC White Paper
- Ade, P., Aguirre, J., Ahmed, Z., et al. 2019, *J. Cosmol. Astropart. Phys.*, 2019, 056
- Baccigalupi, C. 2003, *New A Rev.*, 47, 1127
- Bennett, C. L., Bay, M., Halpern, M., et al. 2003, *ApJ*, 583, 1
- Bonavera, L., Suarez Gomez, S. L., González-Nuevo, J., et al. 2021, *A&A*, 648, A50
- Boudet, N., Mutschke, H., Nayral, C., et al. 2005, *ApJ*, 633, 272
- Cai, Z.-Y., Lapi, A., Xia, J.-Q., et al. 2013, *ApJ*, 768, 21
- Carlstrom, J. E., Holder, G. P., & Reese, E. D. 2002, *ARA&A*, 40, 643
- Casas, J. M., González-Nuevo, J., Bonavera, L., et al. 2022, *A&A*, A110
- Chollet, F. 2015, *Keras*, <https://github.com/fchollet/keras>
- Coulton, W. R., & Spergel, D. N. 2019, *J. Cosmol. Astropart. Phys.*, 2019, 056
- de Zotti, G., Massardi, M., Negrello, M., & Wall, J. 2010, *A&A Rev.*, 18, 1
- Delabrouille, J., Cardoso, J. F., & Patanchon, G. 2003, *MNRAS*, 346, 1089
- Delabrouille, J., Cardoso, J. F., Le Jeune, M., et al. 2009, *A&A*, 493, 835
- Delabrouille, J., Betoule, M., Melin, J. B., et al. 2013, *A&A*, 553, A96
- Dole, H., Lagache, G., Puget, J.-L., et al. 2006, *A&A*, 451, 417
- Duchi, J., Hazan, E., & Singer, Y. 2011, *J. Mach. Learn. Res.*, 12, 2121
- Eriksen, H. K., Dickinson, C., Lawrence, C. R., et al. 2006, *New A Rev.*, 50, 861
- Eriksen, H. K., Jewell, J. B., Dickinson, C., et al. 2008, *ApJ*, 676, 10
- Farsian, F., Krachmalnicoff, N., & Baccigalupi, C. 2020, *J. Cosmology Astropart. Phys.*, 2020, 017
- González-Nuevo, J., Toffolatti, L., & Argüeso, F. 2005, *ApJ*, 621, 1
- Goodfellow, I. J. 2010, Technical Report: Multidimensional, Downsampled Convolution for Autoencoders, Université de Montréal
- Goodfellow, I. J., Bengio, Y., & Courville, A. 2016, *Deep Learning* (Cambridge, MA, USA: MIT Press)
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
- Hauser, M. G., & Dwek, E. 2001, *ARA&A*, 39, 249
- Hensley, B. S., & Draine, B. T. 2021, *ApJ*, 906, 73
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, 9, 90
- Jeffrey, N., Boulanger, F., Wandelt, B. D., et al. 2022, *MNRAS*, 510, L1
- Krachmalnicoff, N., & Puglisi, G. 2021, *ApJ*, 911, 42
- Krachmalnicoff, N., & Tomasi, M. 2019, *A&A*, 628, A129
- Krachmalnicoff, N., Baccigalupi, C., Aumont, J., Bersanelli, M., & Mennella, A. 2016, *A&A*, 588, A65
- Krachmalnicoff, N., Carretti, E., Baccigalupi, C., et al. 2018, *A&A*, 618, A166
- Lapi, A., González-Nuevo, J., Fan, L., et al. 2011, *ApJ*, 742, 24
- Leach, S. M., Cardoso, J. F., Baccigalupi, C., et al. 2008, *A&A*, 491, 597
- LeCun, Y., Boser, B., Denker, J. S., et al. 1989, *Neural Comput.*, 1, 541
- Lecun, Y., Bottou, L., Bengio, Y., & Haffner, P. 1998, *Proc. IEEE*, 86, 2278
- Long, J., Shelhamer, E., & Darrell, T. 2015, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 3431
- Martínez-González, E., Diego, J. M., Vielva, P., & Silk, J. 2003, *MNRAS*, 345, 1101
- Meny, C., Gromov, V., Boudet, N., et al. 2007, *A&A*, 468, 171
- Negrello, M., Clemens, M., Gonzalez-Nuevo, J., et al. 2013, *MNRAS*, 429, 1309
- Oliphant, T. 2006, *NumPy: A Guide to NumPy* (USA: Trelgol Publishing)
- Paradis, D., Bernard, J.-P., Meny, C., & Gromov, V. 2011, *A&A*, 534, A118
- Peebles, P. J. E., & Yu, J. T. 1970, *ApJ*, 162, 815
- Petroff, M. A., Addison, G. E., Bennett, C. L., & Weiland, J. L. 2020, *ApJ*, 903, 104
- Planck Collaboration IX. 2016, *A&A*, 594, A9
- Planck Collaboration X. 2016, *A&A*, 594, A10
- Planck Collaboration XXVI. 2016, *A&A*, 594, A26
- Planck Collaboration I. 2020, *A&A*, 641, A1
- Planck Collaboration IV. 2020, *A&A*, 641, A4
- Ronneberger, O., Fischer, P., & Brox, T. 2015, Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986, *Nature*, 323, 533
- Schneider, P. 2006, *Extragalactic Astronomy and Cosmology* (Berlin, Heidelberg: Springer Berlin Heidelberg)
- Smoot, G. F., Bennett, C. L., Kogut, A., et al. 1992, *ApJ*, 396, L1
- Sugai, H., Ade, P. A. R., Akiba, Y., et al. 2020, *J. L. Temp. Phys.*, 199, 1107
- Toffolatti, L., Argüeso Gomez, F., de Zotti, G., et al. 1998, *MNRAS*, 297, 117
- Tucci, M., Toffolatti, L., de Zotti, G., & Martínez-González, E. 2011, *A&A*, 533, A57
- Wang, G.-J., Shi, H.-L., Yan, Y.-P., et al. 2022, *ApJS*, 260, 13
- Wei, X.-S., Song, Y.-Z., Mac Aodha, O., et al. 2021, *IEEE Transactions on Pattern Analysis and Machine Intelligence*
- Zonca, A., Singer, L. P., Lenz, D., et al. 2019, *J. Open Source Softw.*, 4, 1298

Appendix A: Learning process of CENN

Kernels and filters cited in section 3.1 are adjusted in each epoch by updating non-linear parameters θ_{t+1} as

$$\theta_{t+1} = \theta_t + \Delta\theta_t, \quad (\text{A.1})$$

where θ_t is the parameter of the previous epoch, and

$$\Delta\theta_t = -\alpha g_t, \quad (\text{A.2})$$

where g_t is the gradient value at the epoch t . Furthermore, α depends on the selected optimisation algorithm. Because our training was performed using the adaptive gradient algorithm (AdaGrad, Duchi et al. 2011), the parameter α in Eq. (A.2) varies as

$$\alpha = \frac{\eta}{\sqrt{G_t + \epsilon}}, \quad (\text{A.3})$$

where G_t is a diagonal matrix at epoch t , for which the diagonal elements correspond to the sum of the squares of the past gradients, ϵ is a smoothing term that avoids division by zero, and η is the learning rate, with a value of 0.05 in our model.

When all minibatches of the training dataset flow through all the layers, the gradient in equation (A.2) is then computed in the last deconvolutional block by taking the derivative of the mean-squared error loss function (E) of equation (7) over the output patch O_{v_0} , where the subindex v_0 represents the frequency channel of the output patch (217 GHz in this work). Therefore, the gradient at the last deconvolutional block is easily estimated by using

$$g_t = \frac{\partial E}{\partial O_{v_0}} = \frac{\partial}{\partial O_{v_0}} \left(\frac{1}{2} |O_{v_0} - l_{v_0}|^2 \right) = |O_{v_0} - l_{v_0}|, \quad (\text{A.4})$$

where l_{v_0} is the true value (also known as label or target), that is, the signal we wish to extract from the maps, which in our case is the CMB. For the previous deconvolutional block, the model uses this value, that is, the gradient of the subsequent block, and the chain rule to update its weights and filters. Therefore, for the previous deconvolutional block the gradients are

$$g_{t,W} = \frac{\partial O_{v_0,LD}}{\partial W_{LD-1}} = \frac{\partial O_{v_0,LD}}{\partial O_{LD-1}} \frac{\partial O_{LD-1}}{\partial W_{LD-1}} = g_{t,LD} \frac{\partial O_{LD-1}}{\partial W_{LD-1}} = g_{t,LD} f_{LD-1} \quad (\text{A.5})$$

$$g_{t,f} = \frac{\partial O_{v_0,LD}}{\partial f_{LD-1}} = \frac{\partial O_{v_0,LD}}{\partial O_{LD-1}} \frac{\partial O_{LD-1}}{\partial f_{LD-1}} = g_{t,LD} \frac{\partial O_{LD-1}}{\partial f_{LD-1}} = g_{t,LD} W_{LD-1}, \quad (\text{A.6})$$

where $g_{t,W}$ and $g_{t,f}$ are the gradients used to update the weights and the filters, respectively. The final derivatives in equations (A.5) and (A.6) are computed using the fact that a matrix convolution operation is defined as

$$O_{i,j} = \sum_{m=0}^M \sum_{n=0}^N f(i-m, j-n) W(m, n). \quad (\text{A.7})$$

Moreover, the achronyms LD and LD-1 correspond to the last deconvolutional and its previous blocks, respectively. The remaining deconvolutional blocks have similar expressions. The convolutional blocks have the same derivation as Eqs. (A.5) and (A.6) to update their corresponding weights and filters, respectively. More precisely, for the last convolutional block,

$$g_{t,W} = g_{t,FD} f_{LC} \quad (\text{A.8})$$

$$g_{t,f} = g_{t,FD} W_{LC}, \quad (\text{A.9})$$

where the achronyms FD and LC correspond to the first deconvolutional and the last convolutional block, respectively. For the remaining convolutional blocks, both filters and weights are updated following the same way as in Eqs. (A.8) and (A.9).