

# Convolutional neural networks on the HEALPix sphere: a pixel-based algorithm and its application to CMB data analysis

N. Krachmalnicoff<sup>1,2</sup> and M. Tomasi<sup>3,4</sup>

<sup>1</sup> SISSA, Via Bonomea 265, 34136 Trieste, Italy  
e-mail: [nkrach@sisssa.it](mailto:nkrach@sisssa.it)

<sup>2</sup> IFPU, Institute for Fundamental Physics of the Universe, Via Beirut 2, 34014 Trieste, Italy

<sup>3</sup> Dipartimento di Fisica “Aldo Pontremoli”, Università degli Studi di Milano, Via Celoria 16, 20133 Milano, Italy

<sup>4</sup> Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Milano, Via Celoria 16, 20133 Milano, Italy

Received 5 February 2019 / Accepted 28 June 2019

## ABSTRACT

We describe a novel method for the application of convolutional neural networks (CNNs) to fields defined on the sphere, using the Hierarchical Equal Area Latitude Pixelization scheme (HEALPix). Specifically, we have developed a pixel-based approach to implement convolutional and pooling layers on the spherical surface, similarly to what is commonly done for CNNs applied to Euclidean space. The main advantage of our algorithm is to be fully integrable with existing, highly optimized libraries for NNs (e.g., PyTorch, TensorFlow, etc.). We present two applications of our method: (i) recognition of handwritten digits projected on the sphere; (ii) estimation of cosmological parameter from simulated maps of the cosmic microwave background (CMB). The latter represents the main target of this exploratory work, whose goal is to show the applicability of our CNN to CMB parameter estimation. We have built a simple NN architecture, consisting of four convolutional and pooling layers, and we have used it for all the applications explored herein. Concerning the recognition of handwritten digits, our CNN reaches an accuracy of  $\sim 95\%$ , comparable with other existing spherical CNNs, and this is true regardless of the position and orientation of the image on the sphere. For CMB-related applications, we tested the CNN on the estimation of a mock cosmological parameter, defining the angular scale at which the power spectrum of a Gaussian field projected on the sphere peaks. We estimated the value of this parameter directly from simulated maps, in several cases: temperature and polarization maps, presence of white noise, and partially covered maps. For temperature maps, the NN performances are comparable with those from standard spectrum-based Bayesian methods. For polarization, CNNs perform about a factor four worse than standard algorithms. Nonetheless, our results demonstrate, for the first time, that CNNs are able to extract information from polarization fields, both in full-sky and masked maps, and to distinguish between  $E$  and  $B$ -modes in pixel space. Lastly, we have applied our CNN to the estimation of the Thomson scattering optical depth at reionization ( $\tau$ ) from simulated CMB maps. Even without any specific optimization of the NN architecture, we reach an accuracy comparable with standard Bayesian methods. This work represents a first step towards the exploitation of NNs in CMB parameter estimation and demonstrates the feasibility of our approach.

**Key words.** methods: data analysis – methods: numerical – cosmic background radiation

## 1. Introduction

The field of astrophysics and cosmology is experiencing a rapidly increasing interest in the use of machine learning (ML) algorithms to analyze and efficiently extract information from data coming from experiments. This is happening because of two main factors: (i) the more and more extensive size of astronomical datasets and the ever-increasing requirements in instrumental sensitivities encourage the development of new algorithms to do the job; and (ii) ongoing research in ML continuously produces new concepts, approaches, algorithms, to classify and regress data. ML methods can demonstrate outstanding performance, and the astrophysical community is promptly taking the challenge to adapt them to the problems typically tackled in the field.

One of the most fruitful fields of ML is the domain of artificial neural networks (or neural networks for short, NNs). Roughly speaking, NNs are non-linear mathematical operators that can approximate functions that map inputs onto outputs through a process called training. NNs have already been applied in many fields of science, such as biology (Cartwright 2008; Lodhi 2012; Angermueller et al. 2016), economics (Lam 2004;

Hyun-jung & Kyung-shik 2007; Yan & Ouyang 2017; Tamura et al. 2018), medicine (Virmani et al. 2014; Kasabov 2014; Cha et al. 2016; Liu et al. 2017; Cascianelli et al. 2017), and of course astrophysics (Collister & Lahav 2004; Dieleman et al. 2015; Baccigalupi et al. 2000; Graff et al. 2014; Auld et al. 2007).

In this paper, we present a novel algorithm that extends commonly-used NN architectures to signals projected on the sphere. Spherical domains are of great interest for the astrophysical and cosmological community, which often observe vast portions of the celestial sphere whose curvature is not always negligible. The latter statement applies to several fields, for example: (i) the study of the cosmic microwave background (CMB), the relic thermal radiation coming from the early Universe, which carries crucial information concerning very high energy physical processes occurred just after the Big Bang, and (ii) the analysis of the Universe Large Scale Structure, that is, the mapping of the location and properties of galaxies over large portions of the sky, which yields fundamental clues about the dark cosmological components.

Given our background in the analysis of data coming from experiments observing the CMB radiation, we propose applications of our algorithm to the problem of estimating cosmological

parameters from CMB observations. While we provide only a proof-of-concept of such an application, this work demonstrates for the first time the feasibility of using deep CNNs for CMB parameter estimation on the very large angular scales, in both temperature and polarization. An implementation of the proposed algorithm, developed in Python, is available online<sup>1</sup>.

The paper is structured as follows: in Sect. 2, we introduce the basic concepts underlying the way standard and convolutional NNs (CNNs) work; in Sect. 3, we explain the problem of the application of CNNs on signals projected on the sphere and present our algorithm. Section 4 briefly describes the working environment used to develop and run our algorithm, as well as the NN architecture we have used in our analysis. Sections 5 and 6 deal with possible applications. In Sect. 5, we show the performance of our algorithm on a classification problem (recognition of hand-written digits from the MNIST dataset) on a spherical domain. Section 6 describes the application of our algorithm to CMB-related examples, namely the estimation of cosmological parameters applying CNNs to simulated CMB maps. Lastly, in Sect. 7 we discuss our main results and draw our conclusions.

## 2. Basic concepts of neural networks

In this section, we provide a brief introduction to the basic concepts of NNs. We stick to the two most common types of NNs, which are fully-connected networks and CNNs.

A NN is a computational tool able to approximate a non-linear function  $f$  that maps some inputs  $\mathbf{x} \in \mathbb{R}^n$  into outputs  $\mathbf{y} \in \mathbb{R}^m$ :  $\mathbf{y} = f(\mathbf{x})$ . The advantage of NNs is the lack of requirements on any a priori information for the function  $f$ , as the network can learn it from a (sufficiently large) set of labeled data, that is, a training set of inputs for which the corresponding outputs are known. Given this requirement, NNs fall into the category of supervised machine learning algorithms. NNs can approximate even extremely complicated mappings by recursively applying a non-linear function (called activation function) to linear combinations of the input elements.

### 2.1. Fully connected NN

Several network architectures have been developed in the last years; a particularly simple type builds on so-called fully connected layers. In this kind of NN, the architecture tries to approximate some function  $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$  via a sequence of multiple groups of artificial neurons; each group is called layer. Neurons of consecutive layers are linked to each other by connections. Each connection between neuron  $i$  in layer  $K$  and neuron  $j$  in layer  $K + 1$  is quantified by a scalar weight  $w_{ij}^{K(K+1)}$ , with  $w_{ij}^{K(K+1)} = 0$  if the two neurons are not connected directly. The first layer of a NN, called the input layer, contains  $n$  neurons and is fed with the input. The last layer, the output layer, produces the result of the computation and contains  $m$  neurons. All the other layers, sandwiched between the first and last, are called hidden layers. NNs are called shallow if they have few hidden layers, or deep if they contain several layers (up to a thousand and more He et al. 2016).

The network processes the input and produces the output via the following procedure:

1. The  $n$  elements of the input vector  $\mathbf{x} \in \mathbb{R}^n$  are associated to each neuron in the input layer ( $K = 0$ );

2. For every neuron  $j$  in the  $(K+1)$ th hidden layer, the following value is computed<sup>2</sup>, considering all the neurons of the  $K$ th layer pointing to  $j$ :

$$\text{net}_j^{(K+1)} = g_j \left( \sum_i \text{net}_i^K \cdot w_{ij}^{K(K+1)} \right), \quad (1)$$

where  $g_j$  is the non-linear activation function for neuron  $j$ ,  $\text{net}_i^K$  is the value associated with neuron  $i$  of the  $K$ th layer, and the sum is done over all the neurons  $i$  that are connected to  $j$ .

3. The output  $\tilde{\mathbf{y}} \in \mathbb{R}^m$  is the vector associated with the  $m$  neurons of the output layer, obtained after the computation of the previous step in all the hidden layers.

In order for the network to be able to approximate the function  $f$ , the connection weights must be optimized. This is done through the network training: starting from initialized values of the weights  $w_{ij}^{K(K+1)}$ , the network is fed with the elements  $\mathbf{x}^i$  of the training set and the output values  $\tilde{\mathbf{y}}^i$  are obtained. A cost function  $\mathcal{J}$  is computed:

$$\mathcal{J} = \frac{1}{N} \sum_{i=0}^N \mathcal{L}(\mathbf{y}^i, \tilde{\mathbf{y}}^i), \quad (2)$$

where,  $\mathcal{L}$  is a loss function that measures the distance between the true known output value  $\mathbf{y}^i$  and the one computed by the network  $\tilde{\mathbf{y}}^i$ , for the  $i$ th element of the training set; the sum is done over  $N$  elements of the training set. The minimum of the cost function is found in the following way: (i) the network runs over  $N$  elements of the training set; (ii) the cost function is computed; (iii) the values of the weights are updated moving towards the minimum of the loss function, usually with a gradient descent algorithm that requires the computation of the derivative of  $\mathcal{J}$  with respect to the weights; (iv) the full computation is iterated until the value of  $\mathcal{J}$  converges to its minimum. The application of steps (i)–(iii) on the entire training set is called epoch. Depending on the approach,  $N$  can be equal to one, to the total number of elements in the training set, or to a sub-fraction of it (randomly changing the training set elements included in this fraction at every iteration). In the latter case, the training set is divided in so-called mini-batches.


The NN parameters (e.g., the number of layers and neurons, the loss function, the dimension of mini-batches, the learning rate for the gradient descent algorithm, etc.) are optimized by evaluating the performance of the network over the validation set for which the outputs  $\mathbf{y}$  are also known. The final assessment of the ability of the NN to approximate the mapping  $f$  is done over the test set. Depending on the characteristics of the output and the function  $f$ , NNs can be used, for example, to solve classification problems or to perform parameter regression.

### 2.2. Convolutional NNs

A peculiar type of NN is the so-called convolutional neural network (Krizhevsky et al. 2012), which is the main topic of this paper. In CNNs, convolutional layers take the place of some or all of the fully connected hidden layers. A convolutional layer convolves its input with one or more kernels (filters), and the coefficients of these kernels are the weights that are optimized during training.

<sup>2</sup> The formula in Eq. (1) is called the weighted-average formula, and it is the most common. Depending on the domain and the purpose of the NN, other formulae can be used.

<sup>1</sup> <https://github.com/ai4cmb/NNhealpix>



**Fig. 1.** Application of CNN convolutional filter to 2D image. The input image (shown as a colored matrix) contains  $3 \times 5$  pixels, while the convolutional filter (in gray) is a  $3 \times 3$  matrix. The filter moves across the input image and each time a dot product between the filter and the sub  $3 \times 3$  portion of the image is performed. The result of the product is stored in the output. In this example the filter is not applied to all the possible  $3 \times 3$  partitions of the input image but jumps one pixel at every move (the way the filter moves across the image is controlled by the stride parameter, which in this case is equal to 2 as the filter is translated by two pixels at every move). The filter elements  $w_{ij}$  represent the weights of the network which are optimized during training, and do not change while the filter moves. In this example, the value of each pixel in the output layer is the result of the dot product between the filter and the image pixels in the two squares centered on (2, 2) and (2, 4).

We show a sketch of convolution in Fig. 1: the convolutional kernel is represented by a  $3 \times 3$  matrix, thus containing nine weights  $w_{ij}$ . The kernel scans the input coming from the previous layer, which, in this case, is a  $3 \times 5$  pixels image. The kernel is applied to each sub-region of  $3 \times 3$  pixels of the input image, and the result of the dot product between the two is saved in each pixel of the output layer:

$$p_i = \sum_j w_{ij} x_{ij}, \quad (3)$$

where indices  $x_{ij}$  are the elements of the  $3 \times 3$  portion of the input image. The number of pixels in the output depends on the dimension of the input, the dimension of the kernel, and the way the kernel moves across the input image (see caption of Fig. 1). Once the output is computed, the non-linear activation function is applied to its elements, and this becomes the input of the following layer. In CNNs, training is done as for the fully connected ones, by minimizing a cost function through the application of the network on a training set.

In Fig. 1, we applied the convolutional layer to a 2D image. Nevertheless, CNN can be applied to inputs with other dimensionality, such as 3D images (where, typically, the third dimension deals with the color channels) or to 1D vectors. In the latter case, the kernel is a 1D vector as well, scanning the input vector and applying the dot product on a portion of it.

The fact that CNNs retain the local information of the inputs while performing the convolution on adjacent pixels makes them an extremely powerful tool for pattern recognition. In classification and regression problems, CNNs are often used to extract synthetic information from input data. Typically, this requires some operation that reduces the dimensionality of the input, which is usually done using pooling layers. A pooling layer partitions the input coming from the previous layer into sets of contiguous pixels and applies some operation on each set. Therefore, a pooling layer performs some down-sampling operation; typical examples are: (1) the computation of the average over the pixels

within the partition (average pooling), or (2) the extraction of the pixel with the minimum/maximum value in the partition (min or max pooling).

### 3. CNNs on the HEALPix sphere

The CNN concepts and examples introduced in Sect. 2 are typically applied to Euclidean domains (e.g., 1D vectors and planar images). Nevertheless, it is of great interest to extend the method to signals projected on the sphere. This would have general applications<sup>3</sup> and would also be extremely interesting for those astrophysical and cosmological fields that need to study large parts of the celestial sphere.

In this work, we present a novel pixel-based algorithm to build spherical CNNs, having in mind applications related to cosmological studies. Our implementation relies on the HEALPix (Hierarchical Equal Area iso-Latitude Pixelization) tessellation scheme, which we briefly introduce in the next section.

#### 3.1. The HEALPix spherical projection

The HEALPix tessellation of the sphere was presented by Górski et al. (2005). HEALPix partitions off the spherical surface into 12 base-resolution pixels, and it recursively divides each of them into 4 sub-pixels till the desired resolution is reached. The number of time a sub-division is applied to the base pixels is equal to  $\log_2(N_{\text{side}})$ , where the parameter  $N_{\text{side}}$  is linked to the resolution of the map. The value of  $N_{\text{side}}$  must be a power of 2, and the number of pixels on the map is  $N_{\text{pix}} = 12 \times N_{\text{side}}^2$ .

A signal projected on the sphere with the HEALPix scheme is usually encoded in the computer's memory as a vector, with each element being associated to a single pixel on the map, following a specific order. Two ordering schemes exist; (a) ring order sorts pixels according to iso-latitude rings, and (b) nested order makes use of the hierarchical structure of the tessellation to enumerate pixels. Once the ordering scheme is picked, and for a fixed value of the  $N_{\text{side}}$  parameter, the association between the vector elements and the map pixels is uniquely defined<sup>4</sup>.

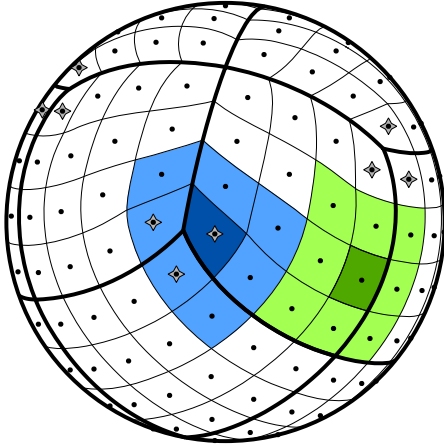
HEALPix has been widely used, with applications not only in multiple branches of astronomy and astrophysics (e.g. analysis of CMB maps Planck Collaboration I 2016, high-energy sources Fermi Collaboration 2015, large-scale structure Percival et al. 2007; Anderson et al. 2012, astroparticle physics IceCube Collaboration 2014, exoplanetary searches Robinson et al. 2011; Majeau et al. 2012; Sullivan et al. 2015, astrohydrodynamics Bryan et al. 2014, gravitational-wave astronomy Singer & Price 2016, etc.) but also in other fields, such as the geophysical one (Westerteiger et al. 2012; Shahram & Donoho 2007). Algorithms able to apply CNNs on spherical maps projected with the HEALPix scheme, as the one presented in this work, have therefore the potential of being of large interest, with possible applications in several different contexts.

#### 3.2. Description of the HEALPix CNN algorithm

The algorithm that we propose implements CNNs over HEALPix maps preserving the local information of the signal by applying kernels to pixels close to each other. We describe the details of this process in the following sections.

<sup>3</sup> One example is the use of computer vision applied to images taken using  $4\pi$  cameras mounted on flying drones.

<sup>4</sup> For a complete description of the HEALPix features and related libraries, we refer to <https://healpix.sourceforge.io>



**Fig. 2.** Tessellation of a sphere according to the HEALPix scheme. In this figure the pixel resolution corresponds to  $N_{\text{side}} = 4$  ( $N_{\text{pix}} = 192$ ). Almost every pixel in a HEALPix map has 8 neighbors; however, a few pixels only have 7 of them. In this figure, the blue pixel has 7 neighbors, while the green one has 8. Black dots represent pixel centers; seven-neighbor pixels are highlighted with grey diamonds. These pixels are at the intersection of three base-pixels (delimited using thick lines); this kind of intersection occurs only near the two polar caps, and in any map with  $N_{\text{side}} \geq 2$  there are exactly 8 of them. Thus, the number of seven-neighbor pixels is always 24 (because of the nature of the orthographic projection used in this sketch, only 9 seven-neighbor pixels are visible).

### 3.2.1. Convolutional layers on a sphere

The majority of pixels in the HEALPix scheme has eight adjacent neighbors (structured with a “diamond” shape). Standard HEALPix libraries implement algorithms to find and enumerate them. In particular, for a given pixel labeled by the index  $p$ , these functions return the indices of its eight closest neighbors following some pre-defined ordering (N-W, N, N-E, E, S-E, S, S-W, and W). However, a few pixels (24 for  $N_{\text{side}} \geq 2$ ) have only 7 neighbors: this happens for those pixels along some of the borders of the 12 base-resolution ones, as shown in Fig. 2.

Our convolutional algorithm makes use of the possibility to quickly find neighbor pixels in the following way (see Fig. 3):

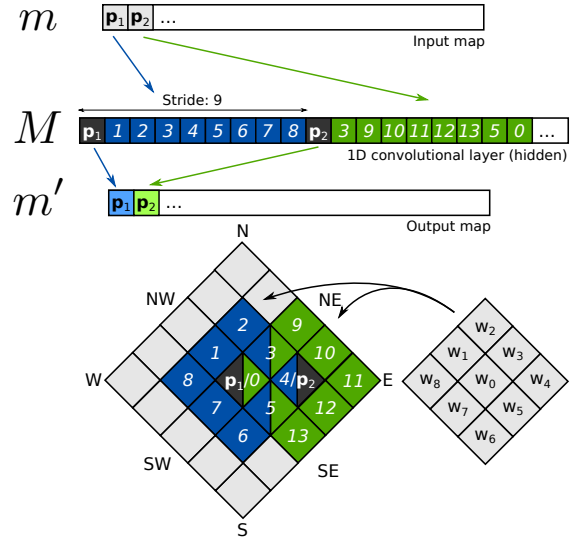
1. The input of the convolutional layer is a map (or a set of maps) at a given  $N_{\text{side}}$ , which is stored in a vector  $m$  of  $N_{\text{pix}}$  elements.

2. Before applying the convolution, the vector  $m$  is replaced by a new vector  $M$  with  $9 \times N_{\text{pix}}$  elements. We associate each element  $p$  in  $m$  with nine elements  $P_i^p$  of  $M$  ( $i = 0 \dots 8$ ), setting  $P_0^p$  equal to the value of the pixel  $p$  itself and the others equal to the values of the eight neighbors of  $p$ , following the clockwise order specified above. Pixels with 7 neighbors are still associated with 9 elements of  $M$ , but the element corresponding to the missing pixel is set to zero.

3. Once the vector  $M$  is generated, a 1D convolution is done in the standard way: a nine element long 1D kernel, with elements  $w_i$ , is applied to  $M$ . This convolution is done with a stride parameter equal to 9, meaning that the kernel is not applied to all the elements of the  $M$  vector, but it makes a jump of nine of them at each move. This is equivalent to the convolution of each pixel with its 8 closest neighbors. The elements  $w_i$  of the kernel represent the weights that the network needs to optimize during the training phase.

4. The output of this operation is a new vector  $m'$  with  $N_{\text{pix}}$  elements, corresponding to the convolved map.

To visualize how convolution works we have prepared a simple example described in Appendix A.



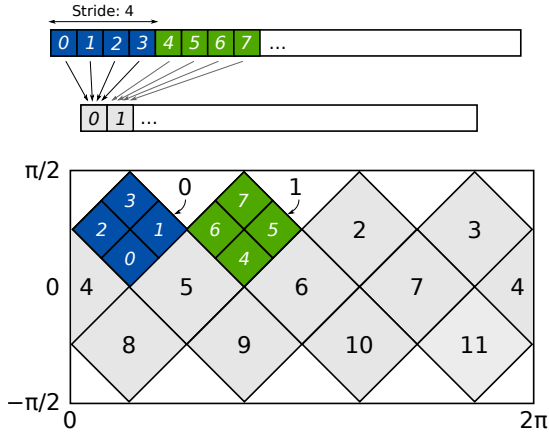
**Fig. 3.** Sketch of the convolution algorithm developed in this work. The elements  $p_1$  and  $p_2$  of the vector  $m$  correspond to two pixels on map (in this case 0 and 4).  $p_1$  and  $p_2$  are associated to nine elements each, in vector  $M$ . The first of this nine elements corresponds to the pixel itself, while the subsequent ones corresponds to its eight neighbors. The neighbors are visited, and unrolled in the  $M$  vector, following a clockwise order starting from the NW direction. The filter (in grey) is also unrolled in a nine element long vector, using the same convention. Convolution is then implemented as a standard 1D convolutional layer with a stride of 9, and the output is the  $m'$  convolved vector (map), which has the same number of elements (pixels) as the input  $m$ . In the figure, numbers over the pixels do not represent the values of the pixels but their unique index: in practical applications, these indices will be those associated to an ordering scheme and a map resolution  $N_{\text{side}}$ . Colors are the same as in Fig. 1, to underline the similarity.

The algorithm presented above has a few interesting properties. Firstly, the search of pixel neighbors needs to be done only once for each  $N_{\text{side}}$  value: the re-ordering scheme that brings a map  $m$  into  $M$  is always the same once  $N_{\text{side}}$  is fixed, and re-ordering operation of a vector is computationally extremely fast. Secondly, the convolution is pixel based, therefore there is an immediate analogy between standard Euclidean convolution and this implementation of a spherical one. Lastly, since the operation is the 1D convolution between a kernel and a vector, it can be performed using the highly-optimized NN libraries commonly available today (e.g., TensorFlow, PyTorch, Theano, Flux, etc.)

On the other hand, our algorithm suffers from one main downside: unlike other spherical CNNs recently introduced in literature (e.g. Cohen et al. 2018; Perraudin et al. 2019), we lose the notation of rotation equivariance. This makes our network less efficient, meaning that it needs a larger set of data to be properly trained and that its final accuracy could be sub-optimal. We discuss this point in Sect. 3.2.2.

### 3.2.2. Pooling layers on a sphere

As we have seen for standard CNNs (Sect. 2.2), pooling layers are useful for lowering the resolution of the network inputs. The hierarchical nature of the HEALPix scheme naturally supports pooling operations, as each pixel in a low-resolution map is mapped to four (or more) sub-pixels in a higher resolution map. In our algorithm, pooling is easily implemented in one dimension using the same approach described in Sect. 3.2.1 for convolution. For each pooling layer, the input vector  $m$  containing



**Fig. 4.** Sketch for the implementation of pooling layers on HEALPix maps. Since a map with  $N_{\text{side}} > 1$  can be degraded to a map with  $N_{\text{side}}/2$ , thanks to HEALPix hierarchical construction, pooling can be implemented naturally using standard 1D pooling layers, if pixels of the input vector (map)  $m$  are properly ordered and a stride equal to 4 is used.

a map with  $N_{\text{pix}}$  pixels is converted to HEALPix nested ordering, and a standard 1D pooling (either min, max, or average) is applied with a stride of 4. This process turns a map with resolution  $N_{\text{side}}$  to a map with  $N_{\text{side}}/2$ . Fig. 4 sketches the process.

### 3.2.3. Related works on spherical CNNs

In the last few years, several works have studied the possibility to apply CNNs to extract informations from signals projected on the sphere. The simplest idea for achieving the goal relies on using standard 2D Euclidean CNNs on planar projection of the sphere, such as the equirectangular one (e.g., Hu et al. 2017; Boomsma & Frellsen 2017). These approaches are easy to implement, but they suffer from the fact that any planar projection of the sphere causes distortions. This problem has been addressed by the CNN implementation described in Coors et al. (2018) where the sampling location of the convolutional filters is adapted in order to compensate distortions. A second approach described in literature requires to project the spherical signal to tangent planes around the sphere, and then to apply standard 2D CNNs to each of these planes (e.g., Xiao et al. 2012). In the limit of making the projection on every tangent plane, this strategy is accurate, but requires large computational resources. Su et al. (2017) proposed a combination of the two approaches by processing spherical signals in their equirectangular projection and mimicking the filter response as if they were applied to the tangent planes.

All the methods described above propose CNNs that are not rotationally equivariant. Rotational equivariance means that if  $R \in \text{SO}(3)$  is a rotation and  $f$  is the convolution operation applied to some input  $x$ , its output  $y = f(x)$  is such that  $f(R(x)) = R(y)$ ,  $\forall R$ . A non-equivariant CNN architecture causes the learning process to be sub-optimal, as the efficiency in filter weight sharing is reduced.

Cohen et al. (2018) proposes the first spherical CNN with the property of being rotational equivariant. Its approach uses a generalized fast Fourier Transform (FFT) algorithm to implement convolutional layers in harmonic domain. Their network is very accurate in 3D model recognition; however, it is computationally expensive, as the computation of spherical harmonics is a  $O(N_{\text{pix}}^2)$  operation.

While we were in implementing the spherical CNN described in this paper, Perraudin et al. (2019) presented an

algorithm to apply the HEALPix tessellation of the sphere to CNNs. Their implementation, called DeepSphere, uses graphs to perform convolutions. It has the advantage of being computationally efficient, with a computational complexity of  $O(N_{\text{pix}})$  and to be adaptable to the analysis of signals that only partially cover the spherical surface. Moreover, by restricting to filters with radial symmetry, their network is close to be rotational equivariant.

Our CNN is specifically built to run on HEALPix, without requiring any other projection of sphere. Similarly to DeepSphere, it is computational efficient, scaling as  $O(N_{\text{pix}})$ , and it is applicable to partially-covered spherical surfaces. Moreover, since the convolution operation ends up in being a traditional 1D convolution, our algorithm has the advantage of being fully integrated in existing libraries for NNs, like TensorFlow, and can therefore rely on years of algorithmic optimization of the convolution in the context of deep learning. On the other hand, the main drawback is the lack of rotation equivariance, which is due to three factors. Firstly, although covering all the same area, pixels in the HEALPix tessellation have different shapes, depending on their position on the sphere (see Fig. 2). This causes some degree of distortion, especially close to the poles. Secondly, as described in Sect. 3.2.1 and Fig. 2, in every HEALPix sphere there are 24 pixels with only seven neighbors instead of eight. This causes distortions that mostly affects convolution on low resolution maps (i.e., small values of  $N_{\text{side}}$ ). Lastly, in this first implementation of our algorithm we do not force filters to have radial symmetry. This makes the efficiency of our network sub-optimal. However, the good performance of our algorithm in the applications that we present in Sects. 5 and 6 shows the validity of our approach in the present case.

It would be interesting to compare our results with the ones of other networks that encode equivariance. However, the spherical CNN introduced by Cohen et al. (2018) is not developed to run on HEALPix; adapting the code would require an amount of work that is outside the scope of this paper. On the other hand, DeepSphere computes convolutions on the HEALPix sphere, similarly to our approach. Nevertheless, the main goal of this paper is to test the feasibility of our algorithm in estimating cosmological parameters that affect the large scales of CMB polarized maps (tensor fields), as described in Sect. 6.2. DeepSphere has not been tested on this kind of analysis yet, as Perraudin et al. (2019) only applied their code to a classification problem performed on scalar fields. Therefore, we defer this comparison to some future works.

## 4. Working environment and NN architecture

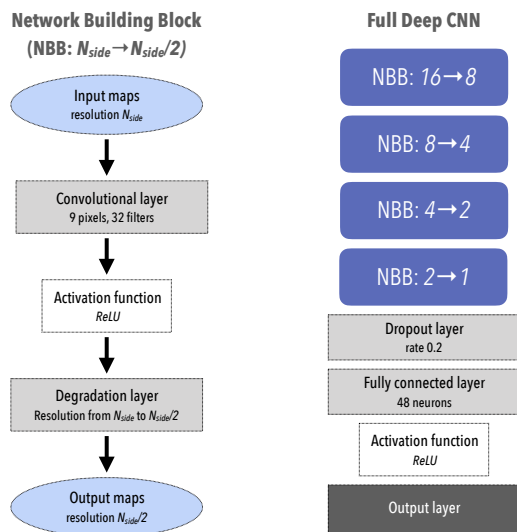
In this section we briefly describe the working environment used to develop and train the NNs that implement the spherical convolutional layers introduced previously. We also present the network architecture used in the applications described in Sects. 5 and 6.

### 4.1. Working environment

We have implemented the algorithm to apply CNNs on the HEALPix sphere in the Python programming language. We have taken advantage of the existing healpy package<sup>5</sup>, a Python library specifically built to handle HEALPix maps.

As we emphasized in Sect. 3, our implementation of convolutional and pooling layers ends up in being a sequence of operations on 1D vectors, which can be applied using existing NN

<sup>5</sup> <https://healpy.readthedocs.io/en/latest/index.html>



**Fig. 5.** Left panel: basic Network Building Block (NBB). Right panel: full deep CNN architecture.

libraries. In our work, we have employed the Keras package<sup>6</sup>, with TensorFlow backend<sup>7</sup> to build, test, and train CNNs.

We have used the computing facilities provided by the US National Energy Resource Scientific Computing Center (NERSC) to perform training. As NERSC does not provide GPU computing nodes at the moment, we relied on CPU-based training architectures. In order to parallelize the computation and make the training more efficient, we used the Horovod package (Sergeev & Balso 2018) to distribute the workload over cluster nodes using MPI<sup>8</sup>.

## 4.2. Network architecture

Using the convolutional and pooling layers described in Sects. 3.2.1 and 3.2.2, we have built a full deep CNN that we have applied to the examples reported in the following sections. For all our applications we have used a single network architecture, changing only the output layer depending on the problem addressed.

We have implemented this CNN architecture by stacking several instances of a Network Building Block (NBB) with the following structure (see left panel of Fig. 5):

1. The first layer in the NBB accepts one or more HEALPix maps with the same resolution parameter  $N_{\text{side}}$  and performs a convolution with  $N$  filters, according to the algorithm described in Sect. 3.2.1 and Fig. 3, then applying an activation function;
2. The next layer contains an average-pooling operation (see Sect. 3.2.2 and Fig. 4);
3. As a result, the NBB lowers the resolution of the input maps from  $N_{\text{side}}$  to  $N_{\text{side}}/2$ .

In all our applications, the input of the network is a set of HEALPix maps at  $N_{\text{side}} = 16$  (each map has 3072 pixels, each of which covers an area of about  $13 \text{ deg}^2$  on the sphere). We have built the full deep CNN by stacking together four NBBs, going from maps at  $N_{\text{side}} = 16$  to  $N_{\text{side}} = 1$ . In each convolutional layer, we convolve the input maps with  $N = 32$  filters. We use

a rectified linear unit (ReLU) as the activation function in the NBB.

After the four building blocks, we include a dropout layer, with a dropout rate set to 0.2, in order to avoid overfitting<sup>9</sup>. Subsequently, the 32 output maps are flattened into one vector, which is fed to a fully connected layer with 48 neurons activated via a ReLU function. The shape of the output layer and its activation function depend on the kind of problem considered. The full CNN architecture is shown on the right panel of Fig. 5.

This architecture, which convolves maps at decreasing angular resolution while moving forward in the network, makes each layer sensitive to features on map laying at different angular scales. As it will be described in the following, this is an important property for the proposed applications.

## 5. Application to the MNIST dataset

The first application of the algorithms presented in this work is the classification of handwritten digits projected over the sphere. We have used images from the MNIST image database<sup>10</sup>, a publicly available database containing 70 000 grayscale  $28 \times 28$  pixels images with handwritten digits (having a repartition between training and test set with a proportion 6:1).

We have chosen this example, despite being of little relevance for astrophysics, because the MNIST dataset is widely used to test the performance of machine learning algorithms, and it is currently considered one of the standard tests to check how well automatic classifiers perform. Moreover, the performance of CNNs depends critically on the existence of a large and reliable training set, and MNIST satisfies this requirement.

### 5.1. Projection

To use MNIST images in our tests we need to project them on the HEALPix maps. We have therefore developed a simple algorithm to project a rectangular image over a portion of the sphere, as shown in Fig. 6. We consider the image as placed on a plane in 3D space, perpendicular to the Equatorial plane of the celestial sphere; the point at the center of the image is on the Equatorial plane. We use a ray tracing technique, firing rays from the center of the sphere against pixels in the plane and finding their intersection with the sphere. We fill each pixel on the HEALPix map hit by a ray with the value of the corresponding pixel in the plane figure. We fire a number of rays sufficient to cover all the HEALPix pixels in the portion of the sphere where we want to project the image, and we bin multiple hits that fall in the same spherical pixel with an average. In this way, the image on the sphere is smooth and has no holes. After the image has been projected with its center on the Equator, the code applies a rotation, in order to center the image on a user-defined point on the sphere.

### 5.2. Training and results

We have performed the training of the NN on the Cori cluster at NERSC, with the working environment and network architecture


<sup>9</sup> The dropout regularization randomly drops neurons and their connections in a NN layer during training. The dropout rate represents the percentage of the neurons of the layers randomly switched off in each training epoch.

<sup>10</sup> The MNIST (Modified National Institute of Standards and Technology) database can be downloaded from <http://yann.lecun.com/exdb/mnist/>

<sup>6</sup> <https://keras.io>

<sup>7</sup> <https://www.tensorflow.org>

<sup>8</sup> For a tutorial on how to work with this environment at NERSC see <https://docs.nersc.gov/analytics/machinelearning/tensorflow/>



**Fig. 6.** Method used to project MNIST images on the sphere. The image is aligned so that the center of the pixel in the middle (marked with a  $\times$ ) is on the Equatorial plane. A number of rays are fired, originating from the center of the sphere, targeting each of the pixels in the image, and the intersection with each pixel on the sphere is computed. In order to cover all the pixels on the sphere, the angular separation among rays is  $1/2$  of the angular resolution of the HEALPix map.


described in Sect. 4. In particular, since this application of our algorithm is a classification problem the last layer of the CNN contains 9 densely-connected neurons, corresponding to the nine different classes of our problem<sup>11</sup>. We have used a Softmax activation function for the last layer, as this is the typical choice for classifiers.

To generate the training set, we have taken 40 000 images out of the MNIST database, and we have projected them on the HEALPix sphere, randomizing on the position, rotation, and size of each image. Each projected image has random dimensions with  $120^\circ < \theta < 180^\circ$  and  $120^\circ < \varphi < 360^\circ$ , therefore covering a large portion of the sphere. We have used this basic data-augmentation to produce a training set of 100 000 different images. The validation set is composed by 10 000 randomly projected images, taken from the remaining 10 000 ones of the MNIST training set.

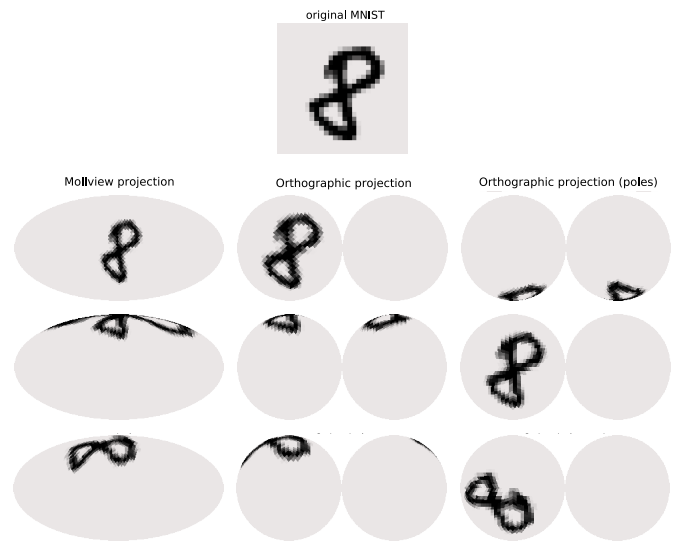
We have adopted the mean squared error as the loss function for training, and categorical accuracy as the metric to assess the NN performance. We have used the Adam optimizer on mini-batches of 32 samples, with learning rate equal to  $10^{-3}$  in the first step and then divided by 10 every time the validation loss has not improved for 5 consecutive epochs. We stopped the training after twenty epochs without improvement. Each epoch required  $\sim 350$  s; training stopped after 50 epochs, for a total of less than 5 h of wall-clock time. Figure 7 reports the training history of the network. The accuracy achieved on the validation set is  $\sim 95\%$ . The network did not show overfitting, but it generalized well the performance from the training set to the validation one.

We have generated 10 000 randomly projected images from the MNIST database to be used as test set, using the same data-augmentation described above. The error we achieved on the test set is about 4.8%. This error is comparable to the one obtained for the classification of the MNIST dataset projected on the sphere by Cohen et al. (2018).

As discussed in Sect. 3.2.2 our CNN is not rotationally equivariant and distortions, due to the different pixel shapes, could affect the network performance, especially close to the poles. For this reason we have checked that, once trained, the network is able to classify handwritten digits with the same accuracy, independently from their position on the sphere. To do so, we have tested the network on three different groups of images, each containing 100 different sets of 300 images for a total of



**Fig. 7.** Learning history for the CNN applied to the MNIST dataset projected on the sphere.



**Fig. 8.** Projection of a MNIST image (upper panel) on the HEALPix sphere. The same image is projected along the Equator (first row), around the north pole (second row) and at a random position on the sphere (third row). For each case, we show the map using the Mollweide projection (first column), an Orthographic projection centered on the Equator and at the north and south poles. The resolution of the projected maps corresponds to  $N_{\text{side}} = 16$ .

30 000 samples for group. The three groups differed only in the region of the sphere where we projected the image:

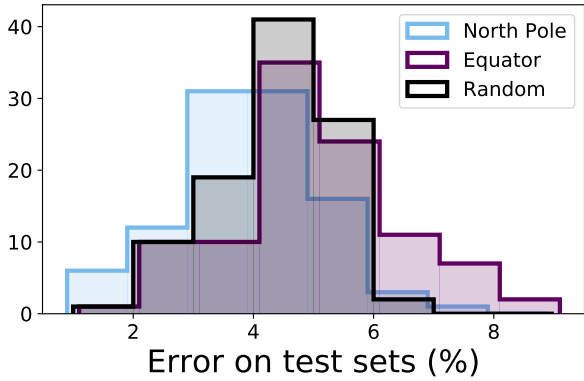
1. images in the first group were projected at random points along the Equator;
2. images in the second group were projected near the North Pole;
3. images in the third group were projected in random positions on the sphere.

We show examples from each group in the first, second and third row of Fig. 8 respectively. We report results in Fig 9 for each group, we provide the histogram of the final error for the one hundred different test sets. The three distributions are close to each other, showing that the network performs equally and independently on the position of the object on the sphere.

## 6. Parameter estimation for CMB

As a second application of our algorithm, we have applied CNNs on the sphere to a regression problem. Our objective was to assess the performance of the NN when applied to the estimation of parameters defining the properties of a Gaussian scalar

<sup>11</sup> Once projected on the sphere with random rotation, numbers 6 and 9 cannot be distinguished. Therefore, we drop out all images containing digit 9 from the MNIST dataset and only use nine classes, corresponding to integer numbers from 0 to 8.



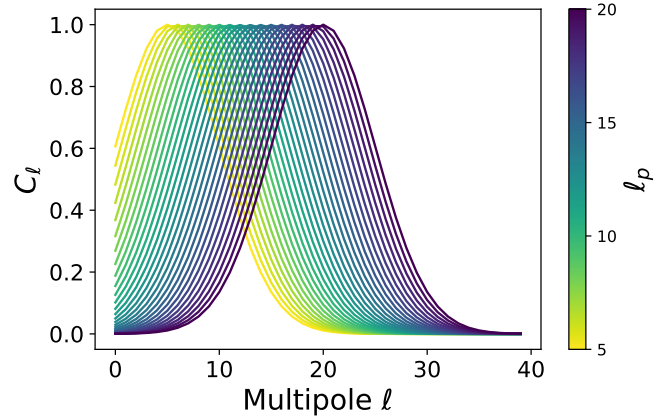
**Fig. 9.** Distribution of NN error on three groups of test sets, where MNIST images have been projected: (i) at the north pole (blue), (ii) at the equator (purple) and (iii) on random positions on the sphere (black). Each group is composed by 100 test sets of 300 images each.

or tensor field on the sphere. This represents the main goal of this explorative work: to understand whether our spherical CNN implementation is suitable to be applied to this kind of problem.

The CMB represents the relic radiation from the Big Bang, emitted when the universe was about 380 000 years old, after the decoupling of matter and photons. It has a blackbody spectrum at the temperature of  $\sim 2.7$  K and it shows a high level of isotropy in the sky. Anisotropies in its temperature, which represent the trace of primordial fluctuations, are of the order of  $\Delta T/T \simeq 10^{-5}$ . The CMB signal is linearly polarized due to Thomson scattering of photons with free electrons in the primordial universe. A net linear polarization can be generated only when this scattering occurred in the presence of a quadrupole anisotropy in the incoming radiation. Since only a small fraction of photons lastly scattered under this condition the polarization fraction of the CMB radiation is rather small, at the level of about 10%.

In the last decades, several experiments have measured the CMB anisotropies<sup>12</sup>, in both total intensity and polarization, making the study of the CMB radiation one of the fundamental branches of modern observational cosmology. In particular, information on the origin, evolution and composition of our Universe can be extracted with a statistical analysis of the CMB temperature and polarization fields used to estimate the values of the cosmological parameters). As primordial CMB fluctuations can be considered as a Gaussian field all the information is encoded in their angular power spectrum. Classical methods for parameter fitting rely on Bayesian statistics: the values of cosmological parameters, defining the properties of our Universe, are typically extracted through the maximization of a likelihood function computed from the angular power spectra of CMB sky maps.

Recently, several works have explored the possibility to estimate cosmological parameters directly from CMB maps using NNs without the computation of angular power spectra. He et al. (2018) have successfully used deep residual CNNs to estimate the value of two cosmological parameters from CMB temperature sky maps, in the flat sky approximation. the DeepSphere algorithm presented by Perraudin et al. (2019) have been tested to classify weak lensing convergence maps (which are scalar fields similarly to CMB temperature maps) generated from two sets of different cosmological parameters. Caldeira et al. (2019) also have applied CNNs to the field of CMB science, extracting



**Fig. 10.** Power spectra computed according to Eq. (4). The  $\ell_p$  parameter, defining the position of the peak, varies in the multipole range 5–20.

the projected gravitational potential from simulated CMB maps whose signal is distorted by the presence of gravitational lensing effect.

As pointed out by the authors of these papers, the possibility to complement traditional Bayesian methods for cosmological parameter estimations with innovative network-based alternatives could represent an important tool to cross-check results, especially in the case where the hypothesis of signal Gaussianity does not hold. This happens if the CMB primordial radiation is contaminated by non-Gaussian spurious signals, such as residual instrumental systematic effects or Galactic emissions.

In the following sections, we demonstrate that the algorithm presented in this work has the potentiality to be used for this purpose. We have first applied the NN to a “mock” cosmological parameter, in order to assess its ability to perform regression and estimate it directly from maps (Sect. 6.1). We have then tested the NN performance on the estimation of the value of  $\tau$ , one of the six fundamental cosmological parameters of the  $\Lambda$ CDM model, which quantifies the Thomson scattering optical depth at reionization, on simulated maps (Sect. 6.2).

### 6.1. Position of a peak in the power spectrum

We have assessed the ability of our CNN to estimate the value of a mock cosmological parameter in pixel-space. This parameter,  $\ell_p$ , defines the angular scale of the peak of a random Gaussian field in the power spectrum. We have constructed simulated maps starting from power spectra having the following form:

$$C_\ell = \exp\left(-\frac{(\ell - \ell_p)^2}{2\sigma_p^2}\right) + 10^{-5}. \quad (4)$$

The peak in the power spectrum is centered around multipole  $\ell_p$ , with standard deviation<sup>13</sup>  $\sigma_p$ . We have fixed the value of  $\sigma_p = 5$ , while we have varied the value of the parameter  $\ell_p$  in the range of multipoles between 5 and 20. Examples of these power spectra, for different values of  $\ell_p$ , are shown in Fig. 10.

To assess the feasibility of CNNs to retrieve cosmological parameters under different conditions, we have estimated the value of  $\ell_p$  in the following cases: on scalar and tensor fields, with or without noise, in the case of full or partial sky coverage, as described in the following sections.

<sup>12</sup> See <https://lambda.gsfc.nasa.gov/product/expt/> for a list of all previous and operating experiments.

<sup>13</sup> We have added a constant small plateau at  $10^{-5}$  in order to avoid convergence issues.



### 6.1.1. Application to a scalar field

We have constructed the training, validation and test sets by simulating maps as random Gaussian realizations of the power spectra defined in Eq. (4) using the synfast module in the healpy package. We have generated a scalar field projected on the sphere for each spectrum, corresponding to a temperature map in the CMB terminology. Maps have been simulated from spectra with  $\ell_p$  randomly chosen from a uniform distribution in the interval 5–20 at the resolution corresponding to  $N_{\text{side}} = 16$ . We have not included noise in this first test.

We have used the working environment described in Sect. 4. We have used 100 000 maps for the training set, and 10 000 maps for the validation set; we have considered different seeds and  $\ell_p$  values for each each map. The network architecture is the same as the one applied to the MNIST example (see Sect. 4.2), with the only exception of the last layer: as our goal was to estimate the parameter  $\ell_p$ , the output layer contained one single neuron. We have trained the network using a mean squared error as loss function, and we have monitored the accuracy of the network using the mean percentage error on the validation set. Training has been done with Adam optimizer and learning rate decay, as for the MNIST case, on mini-batches of 32 sample. The convergence took place after 70 epochs<sup>14</sup>, for a total of about 6 wall-clock hours of training.

We have applied the trained network on a test set, composed by 1000 maps simulated as before, and we have compared the results reached with the network with those obtained with standard Bayesian method. In the latter case, we have estimated the  $\ell_p$  parameter from the power spectrum of each map of the test set (computed with the anafast module within healpy), using a Markov chain Monte Carlo (MCMC) algorithm to maximize the following  $\chi^2$  likelihood:

$$-2 \ln \mathcal{L} = \sum_{\ell} (\widehat{C}_{\ell} - \widetilde{C}_{\ell})^2 / \sigma(\widetilde{C}_{\ell})^2, \quad (5)$$

where  $\widehat{C}_{\ell}$  is derived from the reference model (in this case,  $\widehat{C}_{\ell} = C_{\ell}$  from Eq. (4)),  $\widetilde{C}_{\ell}$  are computed from the simulated map and  $\sigma(\widetilde{C}_{\ell})$  represent their uncertainties. Since we have been running on simulations, we assume that  $\sigma(\widetilde{C}_{\ell}) = \sigma(\widehat{C}_{\ell})$ , with the signal variance equal to

$$\sigma(\widehat{C}_{\ell}) = \widehat{C}_{\ell} \sqrt{\frac{2}{(2\ell + 1)}}. \quad (6)$$

We fit the power spectra up to  $\ell_{\text{max}} = (3N_{\text{side}} - 1) = 47$ .

The mean error obtained with CNN on the test set is  $\sim 1.3\%$ , while the error of the MCMC fit is  $\sim 0.7\%$ . Although the standard Bayesian method performs about twice as good as the NN, there are two important considerations to be made, which make the network performance remarkable:

i. We have not optimized our network architecture, as we have used the same architecture described in Sect. 4.2 adapted to the regression problem. Indeed, as already emphasized, the goal of our work is not to find the best architecture for a given problem, but to prove the feasibility of our approach. An optimized architecture could therefore lead to better results in the estimation of the  $\ell_p$  parameter.

ii. Contrary to the standard maximum likelihood approach applied to power spectra to fit for the value of  $\ell_p$ , the estimation

<sup>14</sup> The code stopped training after 20 epochs without improvement of the validation loss.

**Table 1.** Summary of the results obtained in the estimation of the  $\ell_p$  parameter from the neural network and the MCMC fit, for the noiseless case and the three considered noise levels.

	Noiseless $\sigma_n = 0$	$S/N = 1$ $\sigma_n = 5$	$S/N = 1/2$ $\sigma_n = 10$	$S/N = 1/3$ $\sigma_n = 15$
NN	1.3%	2.9%	5.2%	8.4%
MCMC	0.7%	2.5%	4.8%	7.8%

**Notes.** The values reported represent the mean percentage error (percentage difference between the input and the estimated  $\ell_p$  value) computed over the entire test set of 1000 samples.

with the network does not make any assumption on the Gaussianity of the signal.

We have also tested the network performance in the presence of white noise. We have considered three levels of noise on maps, with standard deviation  $\sigma_n = 5, 10, 15$ , corresponding to a signal-to-noise ratio<sup>15</sup> of about 1, 1/2 and 1/3. For each of the three noise levels, we have trained a new NN, with the same architecture defined above and a training set containing 100 000 signal+noise maps. We have generated each map as a Gaussian realization of the power spectrum of Eq. (4) plus Gaussian noise. The training strategy has been the same used for the noiseless case.

As in the previous case, we have compared the results from CNN with the ones coming from standard MCMC fitting on a test set of 1000 maps, for each of the three noise levels. When fitting for the value of  $\ell_p$  from the power spectra, we have taken into account the fact that the presence of noise on maps induces a bias on the auto-spectra  $\widetilde{C}_{\ell}$ , which must be considered while maximizing the likelihood in Eq. (5). Therefore, in this case, the reference model becomes:

$$\widetilde{C}_{\ell} = C_{\ell} + N_{\ell}, \quad (7)$$

with  $C_{\ell}$  from Eq. (4) and  $N_{\ell}$  being the noise power spectrum. In the case of homogeneous white noise with standard deviation  $\sigma_n$  on a map with  $N_{\text{pix}} = 12 \times N_{\text{side}}^2$ , the noise power spectrum is:

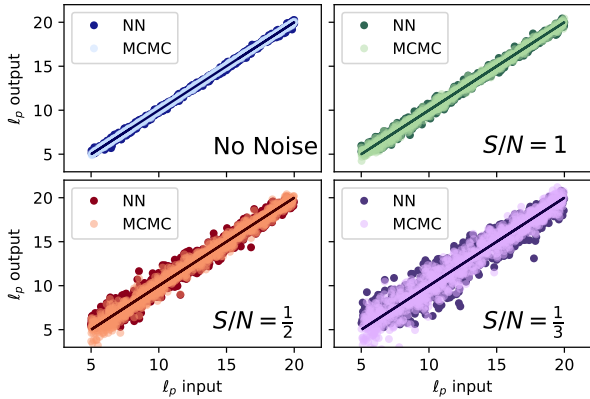
$$N_{\ell} = \frac{4\pi\sigma_n^2}{N_{\text{pix}}}. \quad (8)$$

We present a summary of the results obtained with the NN and with the MCMC approach in Table 1. For each noise level used in the simulation, we report the mean percentage error on the estimated values of the  $\ell_p$  parameter computed over the entire test set. In Fig. 11, we show the distribution of the estimated values of  $\ell_p$  around the true ones for each noise level, considering both approaches. Even in the presence of noise, the accuracy is comparable with the one reached with the MCMC fit.

### 6.1.2. Application to a tensor field

We have tested the ability of the network to estimate  $\ell_p$  also in the case of a tensor field projected on the sphere. This is analog to the case of CMB polarization measurements, and it is of great importance, as current CMB experiments are mainly focusing on polarization observations.

<sup>15</sup> The signal amplitude is defined as the standard deviation of the noiseless maps computed on the whole training set of 100 000 maps. This means that the signal-to-noise ratio is slightly different on each map of the training set, as the signal is generated from spectra with different value of  $\ell_p$ , while the noise standard deviation is the same for all maps.



**Fig. 11.** Comparison of the results obtained in the estimation of the  $\ell_p$  parameter with the neural network and with the MCMC fit, for the noiseless case and for the three considered noise levels.

The CMB signal is linearly polarized, with a polarization fraction of about 10%, and with amplitude and orientation defined by the Stokes parameter  $Q$  and  $U$ . Although CMB experiments directly measure the amplitude of Stokes parameters across the sky, producing  $Q$  and  $U$  maps, these quantities are coordinate-dependent. Therefore, it is more convenient to describe the CMB field using a different basis, dividing the polarization pattern according to its parity properties and using the so-called  $E$ -modes (even) and  $B$ -modes (odd). The use of the  $(E, B)$  basis is convenient also from a physical point of view, as these modes are sourced by different types of perturbations in the primordial universe, with  $B$ -modes produced by tensor (metric) perturbations.


Using tensor spherical harmonics, the polarization tensor field can be described in terms of  $E$  and  $B$  power spectra, which, as for the case of a scalar (temperature) field, give a complete description of the signal statistics in case of a Gaussian field.

For our exercise, we have used the model in Eq. (4) to build  $E$  and  $B$ -mode power spectra: from those, we derive Stokes  $Q$  and  $U$  maps with  $N_{\text{side}} = 16$ . For each pair of maps we have considered polarization power spectra with values of the parameters  $\ell_p^E$  and  $\ell_p^B$  different from each other and randomly draw from a uniform distribution in the multipole interval 5–20.

In this application, the purpose of the NN was to estimate the values of both  $\ell_p^E$  and  $\ell_p^B$  having as input a pair of  $Q$  and  $U$  maps. In order to do so, we slightly change the input and output layers of the network: the input layer accepts two maps ( $Q$  and  $U$ ) instead of one, and the output layer emits two scalar values ( $\ell_p^E$  and  $\ell_p^B$ ). The other network layers share the same architecture used for the recognition of the MNIST dataset in Sect. 5 and for the regression problem on a scalar field in Sect. 6.1.1.

The training set was composed by 100 000 pairs of  $Q$  and  $U$  maps, generated from 100 000 pairs of  $E$  and  $B$  spectra, and the validation set contained 10 000 pairs of maps. We have trained the network using the same procedure described in Sect. 6.1.1: we have minimized a mean squared error loss function, reaching convergence on the validation set after about 7 h of training.

As before, we have assessed the performance of the NN on a test set of 1000 samples. The mean percentage error on the estimation of  $\ell_p^E$  and  $\ell_p^B$  is about 2.7% on the test set. This number needs to be compared with the accuracy that can be reached with standard MCMC fitting. In this case we have estimated  $\ell_p^{E/B}$  by computing the  $E$  and  $B$ -mode power spectra of the test maps, and by minimizing, separately for the two spectra, a chi squared likelihood (see Eqs. (5) and (6)) as we did for the scalar case.



**Fig. 12.** Four circular masks applied to the estimation of the  $\ell_p^{E/B}$  parameters (see Sect. 6.1.2), with a retained sky fraction of about 50, 20, 10 and 5%. On the image masks shown with dark colors retain also the portion of the sky included in those with lighter colors.

The mean percentage error that we get in on both  $\ell_p^E$  and  $\ell_p^B$  from this standard fitting procedure is about 0.7% (analogous to the scalar case).

Although it is true that for polarization field the NN performs worse than the standard MCMC approach, it is important to notice that, with this simple exercise, we demonstrate for the first time that the network can still discern between  $E$  and  $B$ -modes and fit for two parameters that independently affect the statistics of the two polarization states. Also, as already emphasized, we have not optimized the network architecture, but have stuck to the one used also for the other examples presented in this work. However, as in this case the network needs to map a more complicated function to distinguish the two polarization modes, it is worth expecting that larger network could perform better, possibly approaching the MCMC accuracy.

Another important test that can be done with tensor fields projected on the sphere, is the evaluation of parameters from partially covered maps. Ideally, one would expect the uncertainty on cosmological parameter to scale as  $f_{\text{sky}}^{-0.5}$ , at the first order. However, the computation of power spectra on patches of the sky is problematic, as masking induces correlation between Fourier modes and mixing between polarization states. These effects can lead to biases in the estimation of cosmological parameters from spectra, as well as to an increase variance, especially when large scales (low multipoles) are considered. Sophisticated power spectrum estimators, correcting for multipole and polarization state mixing, are currently used to mitigate this problem (Tristram et al. 2005; Grain et al. 2009). It is therefore interesting to understand how the NN accuracy scales as a function of the considered sky fraction.

In our test, we have considered four sky masks, obtained as circular patches in the sky, with retained sky fraction<sup>16</sup> varying between  $\sim 50\%$  and  $\sim 5\%$  (Fig. 12). For each of the four masks, we have trained a new CNN, where the input simulated maps have been obtained as described before, but the  $Q/U$  signal outside the mask have been set to zero. The output of the networks is, as before, the pair of estimates for  $\ell_p^E$  and  $\ell_p^B$ .

Results of are summarized in Table 2. The mean percentage error on the recovered  $\ell_p^{E/B}$  parameter scales roughly as  $f_{\text{sky}}^{-0.36}$ , meaning, therefore, that the NN accuracy is closer to the optimal one for smaller masks (this could be due to the smaller amount of data that it needs to process) and that mixing between different modes and polarization states does not affect its performance.

<sup>16</sup> Masks are obtained by retaining the portions of the sky included in circular regions, centered on the central pixel of a map at  $N_{\text{side}} = 16$ , with radii equal to 90, 53, 37 and 26° and have sky fraction of about 50, 20, 10 and 5% respectively.

**Table 2.** Accuracy (mean percentage error) reached on the estimation of the  $\ell_p^{E/B}$  parameters with CNN for the full sky case and the four sky masks shown in Fig. 12.

$f_{\text{sky}} = 1$	$f_{\text{sky}} = 0.5$	$f_{\text{sky}} = 0.2$	$f_{\text{sky}} = 0.1$	$f_{\text{sky}} = 0.05$
2.7%	3.9%	5.3%	6.4%	8.4%

## 6.2. Estimation of $\tau$


In Sect. 6.1.2, we have showed that our network can estimate the value of simple parameters defined in Fourier space using data in pixel space as input. This proves that the spherical CNN algorithm presented in this paper works well for this kind of regression problems, making it suitable for estimating real cosmological parameters directly from CMB maps (either in total intensity or in polarization). Consequently, we have tested our network architecture on a more realistic case: the estimation of the value of the Thomson scattering optical depth at reionization.

CMB photons, released at the last scattering surface in the early universe, interact with free electrons of the intergalactic medium, ionized by the first emitting objects, at a redshift  $z$  between about 11 and 6, during the so-called reionization epoch. This causes modifications in the CMB signal both in total intensity and in polarization. In particular, for what concern the polarized signal, the effect is visible especially at the larger angular scales (for multipoles  $\ell \lesssim 20$ ) as a bump in the  $E$ -mode power spectrum, caused by the new Thomson scattering events experienced by the CMB photons. The optical depth of Thomson scattering at reionization, usually indicated with the letter  $\tau$ , parametrizes the amplitude and shape of this low- $\ell$  bump.

In order to constrain the value of  $\tau$ , high sensitivity polarization observations on large portion of the sky are needed. Moreover, at the large angular scales affected by the reionization bump, spurious signals, coming either from instrumental systematic effects or from residual foreground emission, can strongly contaminate the measurements, making the estimation of this parameter particularly tricky. Currently, the tightest constraint available on the value of  $\tau$  comes from the full sky observation of the *Planck* satellite, with  $\tau = 0.054 \pm 0.007$  (68% confidence level; [Planck Collaboration VI 2018](#)). Among the six cosmological parameters of the standard  $\Lambda$ CDM model,  $\tau$  is the parameter whose value is currently constrained with the largest uncertainty.

The peculiarity of  $\tau$  makes it an interesting test case for parameter estimation using the algorithm presented in this work. Firstly, the fact that  $\tau$  affects the large angular scales requires algorithms defined on the sphere, as the flat-sky approximation cannot be applied. Secondly, since mainly multipoles at  $\ell \lesssim 20$  are affected by the parameter, low resolution maps (at  $N_{\text{side}} = 16$ ) are sufficient to fit for  $\tau$ : this makes the problem computationally feasible for CNNs. Lastly, the complexity in the estimation of  $\tau$  calls for new analysis techniques that can complement standard parameter estimation routines, thus increasing the confidence on the results.

To estimate  $\tau$  we have used a similar approach as the one described for the  $\ell_p$  parameter, generating training, validation, and test sets from simulations. We have computed a set of five thousands  $E$ -mode power spectra, using the CAMB code ([Lewis & Bridle 2002](#)). Each spectrum has a different value of  $\tau$ , uniformly distributed in the range 0.03–0.08, while the other cosmological parameters are fixed to the best  $\Lambda$ CDM model from *Planck* ([Planck Collaboration VI 2018](#)) (we show a subset of these spectra in Fig. 13). From these spectra, we have generated 100 000 pairs of full sky  $Q$  and  $U$  maps with  $N_{\text{side}} = 16$  for the



**Fig. 13.** Subset of the  $E$ -modes power spectra used to generate to polarization maps for training, validation and test set on which we run the neural network. The value of  $\tau$  changes in the range 0.03–0.08 while all the other cosmological parameters are fixed to the best  $\Lambda$ CDM model from *Planck* results.

training set, and 10 000 and 1000 for the validation and test sets respectively<sup>17</sup>. We have not included any primordial or lensing  $B$ -mode signal.

We have used the same network architecture, working environment and training procedure described in Sect. 4, feeding the network with the simulated  $Q$  and  $U$  maps and producing an estimate for  $\tau$  as the output. By running the CNN on full sky maps, we have obtained a network accuracy of about 4% (mean percentage error of  $\tau$  values over the test set).

We have compared this number with the accuracy reachable through standard MCMC method. As for the previous example, we have maximized the simple chi-squared likelihood of Eq. (5) to fit for the best value of  $\tau$ , where the model  $\widehat{C}_\ell$  represents the theoretical  $E$ -mode power spectrum, and  $\widehat{C}_\ell$  are computed from the test set polarization maps. The accuracy reached with this method is about 2.8%, a factor  $\sim 1.5$  better than the NN one.

Although the NN average error is larger, we believe that this result demonstrates that our CNN algorithm could be a feasible approach to constraint the value of  $\tau$ . Obviously, a complete analysis to understand its real potentiality is needed, including the addition of realistic noise, sky masking, optimization of the NN architecture, error estimation, and so on. However, this is outside the objective of the present work and is hence deferred to a dedicated paper.

## 7. Discussion and conclusions

In this work, we have presented a novel algorithm for the application of CNNs to signals projected on the sphere. We considered the HEALPix tessellation, which allowed us to easily implement convolutional and pooling layers on a pixel domain. The HEALPix pixelization scheme is used in several scientific fields; therefore, the implementation of efficient and compatible CNN algorithms can be of great interest in several contexts.

Our algorithm presents some main advantages. Firstly, convolution is done in the pixel domain, as for standard 1D Euclidean cases: this allows the construction, training, and testing of the CNNs using the existing highly-optimized neural network libraries, such as PyTorch and TensorFlow. Moreover, the algorithm is computationally efficient, with a computational

<sup>17</sup> Each pair of maps is simulated by randomly choosing one of the 5000 spectra with different  $\tau$  and with a different seed.

complexity that scales as  $O(N_{\text{pix}})$ . Lastly, the process can be easily optimized to run only on portion of the sphere.

On the other hand and contrary to other spherical CNNs recently introduced in literature (Perraudin et al. 2019; Cohen et al. 2018), our implementation is not rotationally equivariant. Although this makes the network sub-optimal, the good results obtained on our applications show that training makes the CNN able to overcome this deficiency.

We have built a simple network architecture, which includes four convolutional and pooling layers. Therefore, convolution is applied to maps at decreasingly lower resolutions, making the network sensitive to feature at different angular scales. The same architecture has been used for all the applications presented in this work.

We have first tested the feasibility of our approach on the recognition of handwritten digits (MNIST dataset) projected on the HEALPix sphere. Results show that the trained CNN can recognize handwritten digits that cover large portion of the sphere, with an accuracy of about 95%, a performance comparable with other kinds of spherical NNs. As our network is not spherical invariant, it must be trained with images projected at different positions and orientations on the sphere. When properly done, the performance of the CNN is independent on the position and orientation of images.

We have then moved to applications related to the Cosmology field. In particular, we have applied CNNs to the estimation of cosmological parameters from CMB simulated observations, impacting the very large angular scales ( $\ell \lesssim 20$ ), directly from temperature or polarization sky maps. This work represents the first one in which NNs are tested on such a task.

We have started with the estimation of a mock parameter,  $\ell_p$ , defining the angular scales at which the power spectrum of a Gaussian field projected on the sphere peaks. We have used the same network architecture applied the MNIST recognition, changing only input and output layers depending on whether we fit for the value of  $\ell_p$  in temperature or polarization.

For the temperature case (scalar field projected on the sphere) we have applied the network to both noiseless or noisy maps, showing that it is able to correctly retrieve the value of  $\ell_p$ . The reached accuracy is comparable with the one obtained from the estimation with standard bayesian methods applied to the angular power spectra computed from maps. In particular, for the noiseless case, CNNs perform about a factor two worse than standard methods, but this factor move close to one for the noisy maps.

We have applied the same CNN architecture also to maps in polarization (tensor field projected on the sphere). In this case we fed the network with pairs of Stokes  $Q$  and  $U$  maps, retrieving the values of two parameters  $\ell_p^E$  and  $\ell_p^B$ , representing the angular scales at which the polarization  $E$  and  $B$ -mode power spectra peak. The accuracy reached is about four times worse than for the standard bayesian estimation (percentage error on the retrieved  $\ell_p^{E/B}$  values is  $\sim 2.7\%$  against  $\sim 0.7\%$ ). Although not being optimal, this result is indeed remarkable, as it proves for the first time that CNNs are able to distinguish and separate  $E/B$  polarization in pixel space. We have repeated the test in case of partial sky coverage, applying sky masks, with retained sky fraction between 50 and 5%, to the input  $Q$  and  $U$  maps. The network performs well also under these conditions, with achieved accuracy closer to the optimal one for the smaller masks. This shows that mixing between multipoles and polarization states, due to the sky cut, does not impact the network performance.

Lastly, we have tested the network on the estimation of a real cosmological parameter: the optical depth of Thomson scat-

tering at reionization ( $\tau$ ). Although in this work we only present preliminary results on the estimation of  $\tau$  as a proof-of-concept of our approach, this represents one possible concrete and important application of our network, to be fully explored in the future. The  $\tau$  parameter mostly impacts the large angular scales of  $E$ -mode polarized emission. Its estimation is complicated by the fact that spurious signals, coming from instrumental systematic effects or Galactic residual emissions, can strongly contaminate the cosmological one at low multipoles. As a matter of fact, among the six cosmological parameters of the standard  $\Lambda$ CDM model,  $\tau$  is the parameter whose value is currently constrained with the largest uncertainty (Planck Collaboration VI 2018). For these reasons, pairing standard estimation methods with new ones based on NNs, would be of great importance to achieve a better constraint and cross check results.

We have applied the same CNN architecture used in the previous examples to simulated maps, with values of  $\tau$  in the range 0.03–0.08. We have tested the performance of the network in the simplest case of full sky and noiseless observations, and without applying any specific optimization to the CNN architecture. Results shows that the network is able to reach an accuracy on the estimation of  $\tau$  that is  $\sim 1.5$  worse than standard Bayesian fitting.

Our findings represent a first step towards the possibility to estimate  $\tau$  with NNs, and show that our implementation of CNNs on the sphere is a valid tool to achieve the goal. Obviously, a more sophisticated analysis is needed, including complications to make simulations more representative of real data, such as the introduction of realistic noise, instrumental systematic effects and foreground signals. In order to have a reliable estimate of  $\tau$ , a full characterization if the results obtained with NNs is also needed, especially regarding the estimation of uncertainties. Moreover, specific optimization of the CNN architecture is necessary. We defer all these studies to a subsequent work.


*Acknowledgements.* This research was supported by the ASI-COSMOS Network (<http://cosmosnet.it>). We acknowledge support from the RADIO-FOREGROUNDS project, funded by the European Commission's H2020 Research Infrastructures under the Grant Agreement 687312. The authors thank Prof. Carlo Baccigalupi for useful discussion and early paper review. We acknowledge the use of NERSC for the simulations done in this work.

## References


- Anderson, L., Aubourg, E., Bailey, S., et al. 2012, *MNRAS*, **427**, 3435  
 Angermueller, C., Pärnamaa, T., Parts, L., & Stegle, O. 2016, *Mol. Syst. Biol.*, **12**, 878  
 Auld, T., Bridges, M., Hobson, M. P., & Gull, S. F. 2007, *MNRAS*, **376**, L11  
 Baccigalupi, C., Bedini, L., Burigana, C., et al. 2000, *MNRAS*, **318**, 769  
 Boomsma, W., & Frellsen, J. 2017, in *Advances in Neural Information Processing Systems 30*, eds. I. Guyon, U. V. Luxburg, S. Bengio, et al. (Curran Associates, Inc.), 3433  
 Bryan, G. L., Norman, M. L., O'Shea, B. W., et al. 2014, *ApJS*, **211**, 19  
 Caldeira, J., Wu, W. L. K., Nord, B., et al. 2019, *Astron. Comput.*, **28**, 100307  
 Cartwright, H. M. 2008, *Methods in Molecular Biology*<sup>TM</sup> (Humana Press), 1  
 Cascianelli, S., Scialpi, M., Amici, S., et al. 2017, *Curr. Alzheimer Res.*, **14**, 198  
 Cha, K. H., Hadjiiski, L., Samala, R. K., et al. 2016, *Med. Phys.*, **43**, 1882  
 Cohen, T. S., Geiger, M., Koehler, J., & Welling, M. 2018, ArXiv e-prints [arXiv:1801.10130]  
 Collister, A. A., & Lahav, O. 2004, *PASP*, **116**, 345  
 Coors, B., Condurache, A. P., & Geiger, A. 2018, *European Conference on Computer Vision (ECCV)*  
 Dieleman, S., Willett, K. W., & Dambre, J. 2015, *MNRAS*, **450**, 1441  
 Fermi Collaboration 2015, *ApJS*, **218**, 23  
 Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759  
 Graff, P., Feroz, F., Hobson, M. P., & Lasenby, A. 2014, *MNRAS*, **441**, 1741

- Grain, J., Tristram, M., & Stompor, R. 2009, *Phys. Rev. D*, **79**, 123515
- He, K., Zhang, X., Ren, S., & Sun, J. 2016, in *Computer Vision – ECCV 2016*, eds. B. Leibe, J. Matas, N. Sebe, & M. Welling (Cham: Springer International Publishing), 630
- He, S., Ravanbakhsh, S., & Ho, S. 2018, *International Conference on Learning Representations*
- Hu, H. N., Lin, Y. C., Liu, M. Y., et al. 2017, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*
- Hyun-jung, K., & Kyung-shik, S. 2007, *Appl. Soft Comput.*, **7**, 569
- IceCube Collaboration 2014, *Phys. Rev. Lett.*, **113**, 101101
- Kasabov, N. K. 2014, *Neural Netw.*, **52**, 62
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. 2012, *Proceedings of the 25th International Conference on Neural Information Processing Systems – Volume 1, NIPS'12* (USA: Curran Associates Inc.), 1097
- Lam, M. 2004, *Decis. Support Syst.*, **37**, 567
- Lewis, A., & Bridle, S. 2002, *Phys. Rev. D*, **66**, 103511
- Liu, F., Zhou, Z., Jang, H., et al. 2017, *Magn. Reson. Med.*, **79**, 2379
- Lodhi, H. 2012, *Comput. Stat.*, **4**, 455
- Majeau, C., Agol, E., & Cowan, N. B. 2012, *ApJ*, **747**, L20
- Percival, W. J., Cole, S., Eisenstein, D. J., et al. 2007, *MNRAS*, **381**, 1053
- Perraudin, N., Defferrard, M., Kacprzak, T., & Sgier, R. 2019, *Astron. Comput.*, **27**, 130
- Planck Collaboration I. 2016, *A&A*, **594**, A1
- Planck Collaboration VI. 2018, *A&A*, submitted [arXiv:1807.06209]
- Robinson, T. D., Meadows, V. S., Crisp, D., et al. 2011, *Astrobiology*, **11**, 393
- Sergeev, A., & Balso, M. D. 2018, ArXiv e-prints [arXiv:1802.05799]
- Shahram, M., & Donoho, D. 2007, *Proceedings of SPIE – The International Society for Optical Engineering*, 6701
- Singer, L. P., & Price, L. R. 2016, *Phys. Rev. D*, **93**, 024013
- Su, Y. C., & Grauman, K. 2017, in *Advances in Neural Information Processing Systems 30*, eds. I. Guyon, U. V. Luxburg, S. Bengio, et al. (Curran Associates, Inc.), 529
- Sullivan, P. W., Winn, J. N., Berta-Thompson, Z. K., et al. 2015, *ApJ*, **809**, 77
- Tamura, K., Uenoyama, K., Itsuka, S., & Matsuo, Y. 2018, *Trans. Jpn. Soc. Artif. Intell.*, **33**, A
- Tristram, M., Macías-Pérez, J. F., Renault, C., & Santos, D. 2005, *MNRAS*, **358**, 833
- Virmani, J., Kumar, V., Kalra, N., & Khandelwal, N. 2014, *J. Digit. Imaging*, **27**, 520
- Westerteiger, R., Gerndt, A., & Hamann, B. 2012, in *OpenAccess Series in Informatics (OASIs), Vol. 27, Visualization of Large and Unstructured Data Sets: Applications in Geospatial Planning, Modeling and Engineering – Proceedings of IRTG 1131 Workshop 2011*, eds. C. Garth, A. Middel, & H. Hagen (Dagstuhl, Germany: Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik), 13
- Xiao, J., Ehinger, K., Oliva, A., & Torralba, A. 2012, *Proceedings, IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2695
- Yan, H., & Ouyang, H. 2017, *Wirel. Pers. Commun.*, **102**, 683

## Appendix A: Detecting stripes in a map



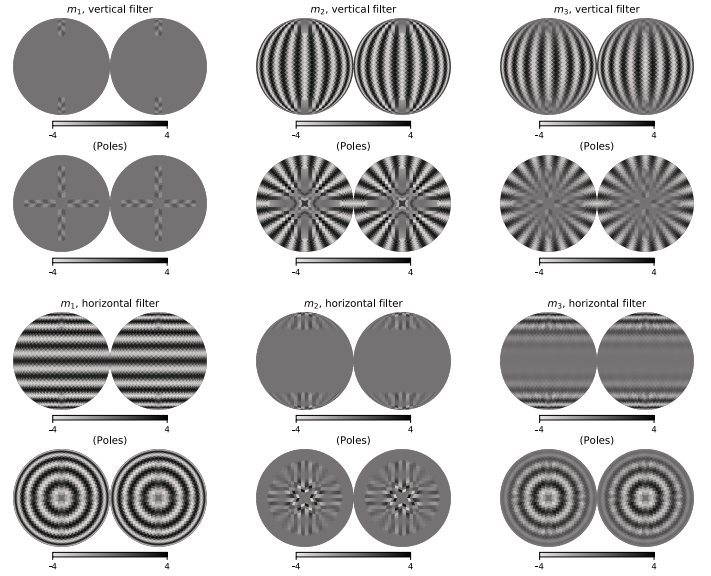
**Fig. A.1.** Filters used in the example of convolution described in Sect. 3.2.1. Black pixels are set to 1, white pixels to  $-1$ , gray pixels to zero. The filter on the left is designed to pick vertical features, i.e., stripes that are aligned with meridians. The filter on the right picks features running along parallels. These filters are here shown with a diamond shape, but practically they are applied as 1D vectors to maps, following the algorithm sketched in Fig. 3. The unrolling into vectors is done clockwise, starting from the NW direction.



**Fig. A.2.** Maps used in the explanation of convolution in Sect. 3.2.1. All the maps have  $N_{\text{side}} = 16$ . Each map is shown using two orthographic projections: in the *upper row*, maps are represented as seen from the Equatorial plane; in the *lower row*, maps are centered around the poles.

In this appendix, we show a pedagogical application of the convolution algorithm presented in Sect. 3.2.1. We have created two nine-elements filters, shown in Fig. A.1. In this figure, filters are shown with a diamond shape, but they are flattened into a nine-element vector to compute the convolution, following the clockwise order of elements. The sum of the values of the weights  $w_i$  is zero in both filters, and the only non-zero pixels are aligned along vertical or horizontal lines. In this way, the application of a filter to a block of nine pixels on a map will be significantly different from zero only if the pixels in this block show some vertical or horizontal features.

We have produced three maps  $m_1$ ,  $m_2$ ,  $m_3$ , where the value associated with each pixel  $p$  is given, respectively, by:



**Fig. A.3.** Result of the application of the two filters in Fig. A.1 to the three maps shown in Fig. A.2.

$$p_1(\theta, \varphi) = \sin(20\theta), \quad (\text{A.1})$$

$$p_2(\theta, \varphi) = \sin(20\varphi), \quad (\text{A.2})$$

$$p_3(\theta, \varphi) = p_1(\theta, \varphi) \sin^2(\theta) + p_2(\theta, \varphi) \cos^2(\theta), \quad (\text{A.3})$$

with  $(\theta, \varphi)$  being the colatitude and longitude of the pixel center ( $0 \leq \theta \leq \pi$ ,  $0 \leq \varphi < 2\pi$ ). The maps are shown in Fig. A.2; map  $m_3$  is a weighted combination of  $m_1$  and  $m_2$ , where the weights  $\sin^2(\theta)$  and  $\cos^2(\theta)$  make horizontal stripes (from  $m_1$ ) and vertical stripes (from  $m_2$ ) negligible at the Equator and at the poles, respectively.

The application of the horizontal and vertical filters in Fig. A.1 to the maps in Fig. A.2 is shown in Fig. A.3. The result of applying the horizontal filter to the vertically striped map is a map with almost no feature: the only residuals are shaped like a plus sign around the pole and correspond to those pixels that only have 7 neighbors instead of 8 (see Fig. 2), as the sum of the weights is no longer zero in this case. Similar results occur in the other maps. In the case of map  $m_3$ , which contains both horizontal and vertical stripes, the two filters correctly pick the right feature depending on the latitude, as expected.