DeepSphere: Efficient spherical Convolutional Neural Network with HEALPix sampling for cosmological applications

Nathanaël Perraudin^a, Michaël Defferrard^b, Tomasz Kacprzak^c, Raphael Sgier^c

^aSwiss Data Science Center (SDSC), Zurich, Switzerland ^bInstitute of Electrical Engineering, EPFL, Lausanne, Switzerland ^cInstitute for Particle Physics and Astrophysics, ETH Zurich, Switzerland

Abstract

Convolutional Neural Networks (CNNs) are a cornerstone of the Deep Learning toolbox and have led to many breakthroughs in Artificial Intelligence. So far, these neural networks (NNs) have mostly been developed for regular Euclidean domains such as those supporting images, audio, or video. Because of their success, CNN-based methods are becoming increasingly popular in Cosmology. Cosmological data often comes as spherical maps, which make the use of the traditional CNNs more complicated. The commonly used pixelization scheme for spherical maps is the Hierarchical Equal Area isoLatitude Pixelisation (HEALPix). We present a spherical CNN for analysis of full and partial HEALPix maps, which we call DeepSphere. The spherical CNN is constructed by representing the sphere as a graph. Graphs are versatile data structures that can represent pairwise relationships between objects or act as a discrete representation of a continuous manifold. Using the graph-based representation, we define many of the standard CNN operations, such as convolution and pooling. With filters restricted to being radial, our convolutions are equivariant to rotation on the sphere, and DeepSphere can be made invariant or equivariant to rotation. This way, DeepSphere is a special case of a graph CNN, tailored to the HEALPix sampling of the sphere. This approach is computationally more efficient than using spherical harmonics to perform convolutions. We demonstrate the method on a classification problem of weak lensing mass maps from two cosmological models and compare its performance with that of three baseline classifiers, two based on the power spectrum and pixel density histogram, and a classical 2D CNN. Our experimental results show that the performance of DeepSphere is always superior or equal to the baselines. For high noise levels and for data covering only a smaller fraction of the sphere, DeepSphere achieves typically 10% better classification accuracy than the baselines. Finally, we show how learned filters can be visualized to introspect the NN. Code and examples are available at https://github.com/SwissDataScienceCenter/DeepSphere.

Keywords: Spherical Convolutional Neural Network, DeepSphere, Graph CNN, Cosmological data analysis, Mass mapping

1. Introduction

Cosmological and astrophysical data often come in the form of spherical sky maps. Observables that cover large parts of the sky, such as the Cosmic Microwave Background (CMB) [1–3], neutral hydrogen [4, 5], galaxy clustering [6], gravitational lensing [7, 8], and others, have been used to constrain cosmological and astrophysical models. Cosmological information contained in these maps is usually extracted using summary statistics, such as the power spectra or higher order correlation functions. Convolutional Neural Networks (CNNs) have been proposed as an alternative analysis tool in cosmology thanks to their ability to automati-

Preprint submitted to Astronomy and Computing

cally design relevant statistics to maximise the precision¹ of the parameter estimation [9–16], while maintaining robustness to noise. This is possible as neural networks (NNs) have the capacity to build rich models and capture complicated non-linear patterns often present in the data. CNNs are particularly well suited for the analysis of cosmological data as their trainable weights are shared across the domain, i.e., the network does not have to relearn to detect objects or features at every spatial location.

¹Here, the word "precision" is to be understood as the final size of the posterior distribution on the measured parameters (including systematic errors).



Figure 1: Overall NN architecture, showing here three convolutional layers acting as feature extractors followed by three fully connected layers with softmax acting as the classifier. A convolutional layer is based on five operations: convolution, non-linearity, batch normalization, down-sampling, and pooling. While most operations are agnostic to the data domain, the convolution and the down-sampling have to be adapted. In this paper, we propose first to model the sphere with a graph and to perform the convolution on the graph. Graphs are versatile data structures which can model any sampling, even irregular or partial. Second, we propose to exploit a hierarchical pixelization of the sphere for the down-sampling operation. It allows the NN to analyze the data at multiple scales while preserving the spatial localization of features. This figure shows a network that operates on the whole sphere. The process is the same when working with partial observations, except that the graph is only built for the region of interest.

So far these algorithms have mostly been demonstrated on Euclidean domains, such as images. The main challenge in designing a CNN on the sphere is to define a convolution operation that is suitable for this domain, while taking care of the necessary irregular sampling. Moreover, the designed convolution and resulting NN should possess the following three key characteristics. First, the convolution should be equivariant to rotation, meaning that a rotation of the input implies the same rotation of the output. Sky maps are rotation equivariant: rotating a map on the sphere doesn't change its interpretation. Depending on the task, we want the CNN to be either equivariant or invariant to rotation.² Second, to be able to train the network in reasonable time, the convolution has to be computationally efficient. Third, a CNN should work well on parts of the sphere, as many cosmological observations cover only a part of the sky. For ground-based observations this can be due to limited visibility of the sky from a particular telescope location, and for space-based instruments due to masking of the galactic plane area (see Figure 2 for example maps).

Three ways of generalizing CNNs to spherical data have been pursued. One approach is to apply a standard 2D CNN to a grid discretisation of the sphere [20–22]. An alternative is to divide the sphere into small chunks and project those on flat 2D surfaces [9, 11, 12, 23].

While these approaches use the well-developed 2D convolution and hierarchical pooling, they are not equivariant to rotation. Another way is to leverage the spherical Fourier transform and to perform the convolution associated to the SO(3) rotation group in the spectral domain, thanks to the convolution theorem [24, 25]. While the resulting convolution is equivariant to rotation, this approach is computationally expensive, even if a fast spherical Fourier transform is used. Moreover, all those methods cannot be much accelerated when maps only span a part of the sky.

Our spherical CNN leverages convolutions on graphs and hierarchical pooling to achieve the following properties: (i) rotation equivariance, (ii) computational efficiency, and (iii) partial sky observations. The main idea is to model the discretised sphere as a graph of connected pixels: the length of the shortest path between two pixels is an approximation of the geodesic distance between them. We use the graph CNN formulation introduced in [26], and a pooling strategy that exploits a hierarchical pixelisation of the sphere to analyse the data at multiple scales. As the Equal Area isoLatitude Pixelisation (HEALPix) [27] is a popular sampling used in cosmology and astrophysics, we tailored the method to that particular sampling. DeepSphere is, however, easily used with other samplings as only two elements depend on it: (i) the choice of neighboring pixels when building the graph, and (ii) the choice of parent pixels when building the hierarchy. The flexibility of modeling the data domain with a graph allows one to easily model

 $^{^2 \}mbox{When only the statistics of the maps are relevant, they are rotation invariant.}$



Figure 2: Example maps on the sphere: (left) the CMB temperature (K) map from Planck [17], with galactic plane masked, (middle) map of galaxy number counts (number of galaxies per arcmin²) in SDSS DR14 [18], and (right) simulated weak lensing convergence map (dimensionless) simulated with DES DR1 mask [19]. These maps were pixelised using $N_{side} = 512$. The CMB and weak lensing mass maps were smoothed with Gaussian kernels with FWHM=1 deg, and the galaxy count map with FWHM=0.5 deg.

data that spans only a part of the sphere, or data that is not uniformly sampled. Using a *k*-nearest neighbours graph, the convolution operation costs $O(N_{pix})$ operations, where N_{pix} is the number of pixels. This is the lowest possible complexity for a convolution without approximations. DeepSphere is readily apt to solve four tasks: (i) global classification (i.e., predict a class from a map), (ii) global regression (i.e., predict a set of parameters from a map), (iii) dense classification (i.e., predict a class for each pixel of a map), and (iv) dense regression, (i.e., predict a set of maps from a map). Input data are spherical maps with a single value per pixel, such as the CMB temperature, or multiple values per pixel, such as surveys at multiple radio frequencies.

We give a practical demonstration of DeepSphere on cosmological model discrimination using maps of projected mass distribution on the sky [28]. These kind of maps can be created using the gravitational lensing technique [see 29, for review]. Our maps are similar to the ones used by [9]. In a simplified scenario, we classify partial sky convergence maps into two cosmological models. These models were designed to have very similar angular power spectrum. We compare the performance of DeepSphere to three baselines: a 2D CNN, and an SVM classifier that takes pixel histograms or power spectral densities (PSDs) of these maps as input. The comparison is made as a function of the additive noise level and the area of the sphere used in the analysis. Results show that our model is always better at discriminating the maps, especially in the presence of noise. DeepSphere is implemented with TensorFlow

[30] and is intended to be easy to use out-of-the-box for cosmological applications. Many plots and co PyGSP [31] for computations and plots. The Python Graph Signal Processing package (PyGSP) [31] is used to build graphs, compute the Laplacian and Fourier basis, and perform graph convolutions. Code and examples are available online.³

2. Method

A CNN is composed of the following main building blocks [32]: (i) a convolution, (ii) a non-linearity, and, optionally, (iii) a down-sampling operation, (iv) a pooling operation, and (v) a normalization.⁴ Our architecture is depicted in Figure 1 and discussed in greater details in Section 2.8. As operations (ii) and (v) are pointwise, they do not depend on the data domain. The pooling operation is simply a permutation invariant aggregation function which does not need to be adapted either. The convolution and down-sampling operations, however, need to be generalized from Euclidean domains to the sphere.

On regular Euclidean domains, such as 1dimensional time series or 2-dimensional images, a convolution can be efficiently implemented by sliding a localized convolution kernel (for example a patch

³https://github.com/SwissDataScienceCenter/ DeepSphere/

⁴Batch normalization has been shown to help training [33]. We verified this experimentally in our setting as well.

of 5×5 pixels) in the signal domain. Because of the irregular sampling, there is no straightforward way to define a convolution on the sphere directly in the pixel domain; convolutions are most often performed using spherical harmonics. In our method we also use the spectral domain to define the convolution. The implementation, however, does not need a direct access to the spectrum, which is computationally more efficient (see Section 2.5).

Down-sampling is achieved on regular Euclidean domains by keeping one pixel every n pixels in every dimension. That is again not a suitable strategy on the sphere because of the irregular sampling.

The gist of our method is to define the convolution operation on a sphere using a graph, and the downsampling operation using a hierarchical pixelisation of the sphere.

2.1. HEALPix sampling

Before doing any numerical analysis on the sphere, one first has to choose a tessellation, i.e., an exhaustive partition of the sphere into finite area elements, where the data under study is quantized. The simplicity of the spherical form belies the intricacy of global analysis on the sphere: there is no known point set that achieves the analogue of uniform sampling in Euclidean space. While our method is applicable to any pixelisation of the sphere, two details depend on the chosen sampling: (i) the choice of neighbours in the construction of the graph, and (ii) the choice of parent vertices when coarsening the graph. As HEALPix [27] is our target application, we tailor the method to that particular sampling in the subsequent exposition. Figure 2 shows three examples of HEALPix maps: the Cosmic Microwave Background [17], galaxies found in Sloan Digital Sky Survey Data Release 14 [18], and an example simulated mass map on the footprint of Dark Energy Survey Data Release 1 [19].

HEALPix is a particular case of a more general class of schemes based on a hierarchical subdivision of a base polyhedron. Another example is the geodesic grids which are based on geodesic polyhedrons, i.e., polyhedrons made of triangular faces. A counter-example is the equirectangular projection, which is not constructed from a base polyhedron, although it can be subdivided. In the particular HEALPix case, the base is a rhombic dodecahedron, i.e., a polyhedron made from 12 congruent rhombic faces. See Figure 5 for an illustration of the base rhombic dodecahedron and its subdivisions.

The HEALPix pixelisation produces a hierarchical subdivision of a spherical surface where each pixel covers the same surface area as every other pixel. A hierarchy is desired for the data locality in the computer memory. Equal area is advantageous because white noise generated by the signal receiver gets integrated exactly into white noise in the pixel space. Isolatitude is essential for the implementation of a fast spherical transform. HEALPix is the sole pixelisation scheme which satisfies those three properties.

The lowest possible resolution is given by the base partitioning of the surface into $N_{pix} = 12$ equal-sized pixels (right-most sphere in Figure 5). The resolution changes as $N_{pix} = 12N_{side}^2$ such that $N_{pix} = 48$ for $N_{side} = 2$ and $N_{pix} = 192$ for $N_{side} = 3$. High-resolutions maps easily reach millions of pixels.

2.2. Graph construction

Our graph is constructed as an approximation of the sphere S^2 , a 2D manifold embedded in \mathbb{R}^3 . Indeed, [34] showed that the graph Laplacian converges to the Laplace-Beltrami when the number of pixels goes to infinity providing uniform sampling of the manifold and a fully connected graph built with exponentially decaying weights. While our construction does not exactly respect their setting (the sampling is deterministic and the graph is not fully connected), we empirically observe a strong correspondence between the eigenmodes of both Laplacians (see Appendix A).

From the HEALPix pixelization, we build a weighted undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$, where \mathcal{V} is the set of $N_{pix} = |\mathcal{V}|$ vertices, \mathcal{E} is the set of edges, and \mathcal{W} is the weighted adjacency matrix. In our graph, each pixel *i* is represented by a vertex (also called vertex) $v_i \in \mathcal{V}$. Each vertex v_i is then connected to the 8 (or 7)⁵ vertices v_j which represent the neighboring pixels *j* of pixel *i*, forming edges $(v_i, v_j) \in \mathcal{E}$. Given those edges, we define the weighted adjacency matrix $\mathcal{W} \in \mathbb{R}^{N_{pix} \times N_{pix}}$ as

$$W_{ij} = \begin{cases} \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|_2^2}{\rho^2}\right) & \text{if pixels } i \text{ and } j \text{ are neighbors,} \\ 0 & \text{otherwise,} \end{cases}$$

where x_i is a vector encoding the 3-dimensional coordinates of pixel *i*, and

$$\rho = \frac{1}{|\mathcal{E}|} \sum_{(v_i, v_j) \in \mathcal{E}} ||\mathbf{x}_i - \mathbf{x}_j||_2$$

is the average Euclidean distance over all connected pixels. This weighting scheme is important as distances between pixels are not equal. Other weighting schemes are

⁵The $12 \times 4 = 48$ pixels at the corner of each rhombus of the base dodecahedron only have 7 neighboring pixels. See Figures 6 and 5.

possible. For example, [35] uses the inverse of the distance instead. We found out that the one proposed above works well for our purpose, and did not investigate other approaches, leaving it to future work. Figure 6 shows a graph constructed from the HEALPix sampling of a sphere.

2.3. Graph Fourier basis

Following [36], the normalized graph Laplacian, defined as $L = I - D^{-1/2}WD^{-1/2}$, is a second order differential operator that can be used to define a Fourier basis on the graph. Here **D** is the diagonal matrix where $D_{ii} = d_i$ and $d_i = \sum_j W_{ij}$ is the weighted degree of vertex v_i . By construction, the Laplacian is symmetric positive semi-definite and hence can be decomposed as $L = U\Lambda U^{\mathsf{T}}$, where $U = [u_1, \ldots, u_{N_{pix}}]$ is an orthonormal matrix of eigenvectors and Λ is a diagonal matrix of eigenvalues. The graph Fourier basis is defined as the Laplacian eigenvectors, motivated by the fact that a Fourier basis should diagonalize the Laplacian operator. The graph Fourier transform of a signal $f \in \mathbb{R}^{N_{pix}}$ is simply its projection on the eigenvectors given by $f = \mathcal{F}_{\mathcal{G}}\{f\} = U^{\mathsf{T}}f$. It follows that the inverse graph Fourier transform reads $\mathcal{F}_{\mathcal{G}}^{-1}{\hat{f}} = U\hat{f} = UU^{\mathsf{T}}f = f$. Note that the Fourier modes are ordered in the increasing order of the Laplacian eigenvalues Λ , which can be interpreted as squared frequencies. Indeed,

$$\boldsymbol{\Lambda}_{ii} = \boldsymbol{u}_i^{\mathsf{T}} \boldsymbol{L} \boldsymbol{u}_i = \sum_{(v_j, v_k) \in \mathcal{E}} \frac{\boldsymbol{W}_{jk}}{\sqrt{\boldsymbol{d}_j \boldsymbol{d}_k}} (\boldsymbol{U}_{ji} - \boldsymbol{U}_{ki})^2$$

is a measure of the variation of the eigenvector u_i on the graph defined by the Laplacian L.

Figure 3 shows the Fourier modes of a HEALPix graph, created using the graph construction described above. The graph Fourier modes resemble the spherical harmonics. That is a strong hint that the graph is able to capture the spherical properties of the HEALPix sampling. This topic is further discussed in Appendix A.

2.4. Convolution on graphs

As there is no notion of translation on a graph, we cannot convolve two graph signals in a strict sense. We can, however, convolve a signal with a kernel defined in the spectral domain. More precisely, we can filter a graph signal by a kernel. Given the convolution kernel $h : \mathbb{R}_+ \to \mathbb{R}$, a signal $f \in \mathbb{R}^{N_{ptx}}$ on the graph is filtered as

$$h(L)f = Uh(\Lambda)U^{\mathsf{T}}f,\tag{1}$$

where $h(\Lambda)$ is a diagonal matrix where $h(\Lambda)_{ii} = h(\Lambda_{ii})$.



Figure 3: The first 16 eigenvectors of the graph Laplacian, an equivalent of Fourier modes, of a graph constructed from the HEALPix sampling of the sphere ($N_{side} = 16$). Eigenvectors 1–3 could be associated with spherical harmonics of degree $\ell = 1$ and order |m| = (0, 1), eigenvectors 4–8 with degree $\ell = 2$ and order |m| = (0, 1, 2), and eigenvectors 9–15 with degree $\ell = 3$ and order |m| = (0, 1, 2, 3). Nevertheless, graph eigenvectors are only approximating spherical harmonics.

Contrary to classical signal processing on Euclidean domains, the kernel h has no single representation in the vertex domain and cannot be translated on the graph. It can however be *localized* on any vertex v_i by the convolution with a Kronecker delta⁶ $\delta^i \in \mathbb{R}^{N_{pix}}$. The localization operator \mathcal{T}_i reads $\mathcal{T}_i h = h(L)\delta^i = (h(L))_i$, the *i*th column of h(L). This localization of the kernel h can be useful to visualize kernels, as shown in an example of heat diffusion presented in Appendix B. If the graph is not regular, i.e., all vertices do not have the same number of neighbors, and all distances are not equal, the effect of the kernel will slightly differ from one vertex to another. While there is no perfect sampling of the sphere, these differences are negligible as the structure of the whole graph is very regular. However, when considering only parts of the sphere, one can observe important border effects (see Appendix C).

Finally, the graph convolution can be interpreted in the vertex domain as a scalar product with localizations $\mathcal{T}_i h$ of the kernel *h*. Indeed, the result of the convolution of the signal *f* with the kernel *h* is

$$(h(L)f)_i = \langle \mathcal{T}_i h(L), f \rangle = \langle h(L)\delta^i, f \rangle.$$
(2)

To make the parallel with the classical 1D convolution, let $f, g \in \mathbb{Z} \to \mathbb{R}$ be two 1D discrete signals. Their convolution can be written in the same form as (2):

$$(f * g)[i] = \sum_{j=-\infty}^{\infty} f[j]g[i-j] = \langle T_ig, f \rangle,$$

⁶A Kronecker delta is the signal $\delta^i \in \mathbb{R}^{N_{pix}}$ that is zero everywhere except on vertex v_i where it takes the value one.



Figure 4: Comparison of filtering speed for Gaussian smoothing of maps of various sizes. The fast spherical harmonic transform (SHT) is implemented by the healpy Python package (via the healpy.sphtfunc.smoothing function). The graph filtering is defined by (4) and implemented with the numpy and scipy Python packages. Both are executed on a single core. The theoretical cost of filtering on the graph is $O(KN_{pix})$ and $O(\ell_{max}^3) = O(N_{pix}^{3/2})$ for the spherical harmonics, where ℓ_{max} is the largest angular frequency. The timings for the partial graphs correspond to a convolution on two fractions (1/12 and 1/192) of the sphere, and illustrates the $O(N_{pix})$ scaling of graph convolutions.

where $T_ig[j] = g[i - j]$ is, up to a flip (i.e., g[i - j] instead of g[i + j]), a translation operator. Similarly as (2), the convolution of the signal f by a kernel g is the scalar product of f with translated versions T_ig of the kernel g. Additionally, it turns out that the localization operator \mathcal{T}_ih is a generalization of the translation operator on graphs. In the particular case where the Laplacian matrix L is circulant, \mathcal{T}_ih is a translated version of \mathcal{T}_jh for all i, j and both convolutions are equivalent. We refer the reader to [37, Sec 2.2] for a detailed discussion of \mathcal{T}_i .

To shed some light on the meaning of the convolution on a graph, we show in Appendix B that the diffusion of heat on a graph can be expressed as the convolution of an initial condition f with a heat kernel h.

2.5. Efficient convolutions

While (1) is a well justified definition of the convolution, it is computationally cumbersome. As no efficient and general fast Fourier transform (FFT) exists for graphs [38], the execution of the Fourier transform by multiplication of the signal f by the dense matrix Ucosts $O(N_{pix}^2)$ operations. In a NN, this operation has to be performed for each forward and backward pass. As current training procedures require processing of many samples, that would be very slow. Moreover, the eigendecomposition of the Laplacian L is needed to obtain the Fourier basis U in the first place. That has a unique cost of $O(N_{pix}^3)$ operations.

Fortunately, both of these computational issues are overcome by defining the convolution kernel *h* as a polynomial $h_{\theta}(\lambda) = \sum_{k=0}^{K} \theta_k \lambda^k$ of degree *K* parametrised by *K* + 1 coefficients θ . The filtering operation (1) becomes

$$h_{\theta}(\boldsymbol{L})\boldsymbol{f} = \boldsymbol{U}\left(\sum_{k=0}^{K} \theta_{k} \boldsymbol{\Lambda}^{k}\right) \boldsymbol{U}^{\mathsf{T}} \boldsymbol{f} = \sum_{k=0}^{K} \theta_{k} \boldsymbol{L}^{k} \boldsymbol{f}, \quad (3)$$

where L^k captures k-neighborhoods. The entry $(L^k)_{ii}$ is the sum of all weighted paths of length k between vertices v_i and v_j , where the weight of a path is the multiplication of all the edge weights on the path. Hence, it is non-zero if and only if vertices v_i and v_j are connected by at least one path of length k.⁷ Filtering with a polynomial convolution kernel can thus be interpreted in the pixel (vertex) domain as a weighted linear combination of neighboring pixel values. In the classical setting, convolutions are also weighted local sums. The weights are however given by the filter coefficients only, and there is one coefficient per pixel in the patch. In the graph setting, the weights are determined by the filter coefficients θ and the Laplacian L, and there is one coefficient per neighborhood, not vertex. That defines radial filters, values of which depend only on the distance to the center, and not on the direction. While it may seem odd to restrict filters to be 1D while the sphere is 2D, radial filters are direction-less and result in rotation equivariant convolutions. We note that restricting the graph convolutional kernel to a polynomial is similar to restricting the classical Euclidean convolution to a fixed-size patch. Similarly, each column of the matrix $\sum_{k=0}^{K} \theta_k L^k$ defines an irregular patch of radius K. Hence, filters designed as polynomials of the Laplacian have local support in the vertex domain.

Following [26], we define our filters as Chebyshev polynomials. The filtering operation (1) becomes

$$h_{\theta}\left(\tilde{\boldsymbol{L}}\right)\boldsymbol{f} = \boldsymbol{U}\left(\sum_{k=0}^{K}\theta_{k}T_{k}\left(\tilde{\boldsymbol{\Lambda}}\right)\right)\boldsymbol{U}^{\mathsf{T}}\boldsymbol{f} = \sum_{k=0}^{K}\theta_{k}T_{k}\left(\tilde{\boldsymbol{L}}\right)\boldsymbol{f}, \quad (4)$$

where

$$\tilde{\boldsymbol{L}} = \frac{2}{\lambda_{max}} \boldsymbol{L} - \boldsymbol{I} = -\frac{2}{\lambda_{max}} \boldsymbol{D}^{-1/2} \boldsymbol{W} \boldsymbol{D}^{-1/2}$$

is the rescaled Laplacian with eigenvalues $\tilde{\Lambda}$ in [-1, 1]. $T_k(\cdot)$ is the Chebyshev polynomial of degree k defined

 $^{^{7}}$ The length of a path between two vertices defines a distance on the graph.

by the recursive relation $T_k(\tilde{L}) = 2\tilde{L}T_{k-1}(\tilde{L}) - T_{k-2}(\tilde{L})$, $T_1(\tilde{L}) = \tilde{L}$, $T_0(\tilde{L}) = I$. While definitions (3) and (4) both allow the representation of the same filters, we found in our experiments that optimizing θ in (4) is slightly more stable than θ in (3). We believe that this is due to (i) their almost orthogonality in the spectral and spatial domains, and (ii) their uniformity.⁸ Finally, note that while the graph convolution (1) is motivated in the spectral domain, definitions (3) and (4) are implementations in the vertex domain.

Exploiting the recursive formulation of Chebyshev polynomials, evaluating (4) requires O(K) multiplications of the vector f with the sparse matrix \tilde{L} . The cost of one such multiplication is $O(|\mathcal{E}| + |\mathcal{V}|)$. By construction of our graph, $|\mathcal{E}| < 8N_{pix}$ and the overall computational cost of the convolution reduces to $O(N_{pix})$ operations and as such is much more efficient than filtering with spherical harmonics, even though HEALPix was designed as an iso-latitude sampling that has a fast spherical transform. This is especially true for smooth kernels which require a low polynomial degree K. Figure 4 compares the speed of low-pass filtering for Gaussian smoothing using the spherical harmonics and the graph-based method presented here. On a single core, a naive implementation of our method is ten to twenty times faster for $N_{side} = 2048$, with K = 20 and K = 5, respectively, than using the spherical harmonic transform at $\ell_{max} = 3N_{side}$ implemented by the highly optimized libsharp [39] library used by HEALPix. The further important speed-up of graph convolutions on fractions of the sphere is a direct reflection of the $O(N_{pix})$ complexity. On graphs, you only pay for the pixels that you use. Note that the two might however scale differently, given that libsharp can be distributed on CPUs through MPI while our method can be distributed on GPUs by TensorFlow.

2.6. Coarsening and Pooling

Coarsening can be naturally designed for hierarchical pixelisation schemes, where each subdivision divides a cell in an equal number of child sub-cells. To coarsen, the sub-cells are merged to summarise the data supported on them. Merging cells lead to a coarser graph. Coarsening defines C(i), the set of children of vertex v_i . For the HEALPix subdivision scheme, the number of children is constant, i.e., $|C(i)| = 4^p \forall i$, for some p.



Figure 5: Two levels of coarsening and pooling: groups of 4 cells are merged into one, then the data on them is summarized in one. The coarsest cell covers 1/12 of the sphere.

Pooling refers to the operation that summarizes the data supported on the merged sub-cells in one parent cell. Given a map $x \in \mathbb{R}^{N_{pix}}$, pooling defines $y \in \mathbb{R}^{N'_{pix}}$ such that

$$y_i = f\left(\left\{x_j : j \in C(i)\right\}\right), \ \forall i \in [N'_{pix}],\tag{5}$$

where f is a function which operates on sets (possibly of varying sizes) and N_{pix}/N'_{pix} is the down-sampling factor, which for HEALPix is

$$|C(i)| = N_{pix}/N_{pix}' = (N_{side}/N'_{side})^2 = 4^p$$

where $p = \log_2(N_{side}/N'_{side})$. That operation is often taken to be the maximum value, but it can be any permutation invariant operation, such as a sum or an average. Figure 5 illustrates the process.

2.7. Layers

Neural networks are constructed as stacks of layers which sequentially transform the data from its raw representation to some predictions. The general Deep-Sphere architecture, pictured in Figure 1, is composed of many layers. The convolutional part, the head of the NN, is composed of graph convolutional layers (GC), pooling layers (P), and batch normalization layers (BN). The tail is composed of multiple fully connected layers (FC) followed by an optional softmax layer (SM) if the network is used for discrete classification. A non-linear function $\sigma(\cdot)$ is applied after every linear GC and FC layer, except for the last FC layer. That operation is point-wise, i.e., $y_{ii} = \sigma(x_{ii})$ and $y_i = \sigma(x_i)$ for matrices X, Y and vectors x, y. The rectified linear unit (ReLU) $\sigma(\cdot) = \max(\cdot, 0)$ is a common choice, and is the one we adopted in this contribution.

Given a matrix $X = [x_1, \ldots, x_{F_{in}}] \in \mathbb{R}^{N_{pix} \times F_{in}}$, a *GC* layer computes $Y = GC(X) = [y_1, \ldots, y_{F_{out}}] \in \mathbb{R}^{N_{pix} \times F_{out}}$, where N_{pix} is the number of pixels (and vertices), F_{in} is the number of input features, and F_{out} is the number of output features. Using the efficient graph

⁸The amplitude of the Chebyshev polynomials $T_k(x)$ is mostly constant over the domain [-1, 1], independently of the order k. On the contrary, the amplitude of the monomials x^k is very different for $|x| \approx 0$ and $|x| \approx 1$.

convolution from (4), each output feature map is computed as

$$\mathbf{y}_i = \sum_{j=1}^{F_{in}} h_{\theta_{ij}}(\tilde{L}) \mathbf{x}_j + b_i \in \mathbb{R}^{N_{pix}}, \ \forall i \in [F_{out}].$$

As such, a *GC* layer is composed of $F_{in} \times F_{out}$ filters, each parameterized by *K* numbers (see Section 2.5). A bias term $\boldsymbol{b} \in \mathbb{R}^{F_{out}}$ is jointly optimized.

Given a matrix $X \in \mathbb{R}^{N_{pix} \times F}$, a pooling layer computes $Y = P(X) \in \mathbb{R}^{N'_{pix} \times F}$ by reducing its spatial resolution $(N'_{pix} < N_{pix})$ according to (5). The batch normalization layer [33] computes Y = BN(X) such as

$$\mathbf{y}_i = \gamma_i \frac{\mathbf{x}_i - \mathbf{E}(\mathbf{x}_i)}{\sqrt{\operatorname{Var}(\mathbf{x}_i) + \epsilon}} + \beta_i, \ \forall i \in [F],$$

where γ_j and β_j are parameters to be learned and ϵ is a constant added for numerical stability. The empirical expectation $E(\mathbf{x}_i) \in \mathbb{R}$ and variance $Var(\mathbf{x}_i) \in \mathbb{R}$ are taken across training examples and pixels.

The layer $FC : \mathbb{R}^{F_{in}} \to \mathbb{R}^{F_{out}}$ is defined as

$$\mathbf{y} = FC(\mathbf{x}) = \mathbf{W}\mathbf{x} + \mathbf{b},\tag{6}$$

where $\boldsymbol{W} \in \mathbb{R}^{F_{out} \times F_{in}}$ and $\boldsymbol{b} \in \mathbb{R}^{F_{out}}$ are the parameters to be learned. Note that the output $\boldsymbol{Y} \in \mathbb{R}^{N_{pix} \times F_{out}}$ of the last *GC* is vectorized as $\boldsymbol{x} = \text{vec}(\boldsymbol{X}) \in \mathbb{R}^{F_{in}}$ before being fed to the first *FC*, where $F_{in} = N_{pix} \times F_{out}$.

The softmax layer is the last layer in a NN engineered for classification. Given the output $x \in \mathbb{R}^{N_{classes}}$ of the last *FC*, called the logits in the deep learning literature, the softmax layer outputs y = S M(x) such that

$$y_i = \frac{\exp(x_i)}{\sum_{j=1}^{N_{classes}} \exp(x_j)} \in [0, 1], \ \forall i \in [N_{classes}].$$

where $N_{classes}$ is the number of classes to discriminate. Thanks to the softmax, the output $\mathbf{y} \in \mathbb{R}^{N_{classes}}$ of a NN is a discretised conditional distribution for the class given the data and trained parameters. That is, y_i is the confidence of the network that the input sample belongs to class *i*. This last layer is actually normalizing \mathbf{x} into \mathbf{y} such that $||\mathbf{y}||_1 = \sum_i y_i = 1$.

2.8. Network architectures

Given a map $X \in \mathbb{R}^{N_{pix} \times F_{in}}$, a neural network computes $NN_{\theta}(X)$, where NN is a composition of the above layers and θ is the set of all trainable parameters. The number of input features F_{in} depends on the data. For the CMB radiation temperature, $F_{in} = 1$. For observations in radio frequencies, F_{in} would be equal to the

| | Classification | Regression |
|--------|--|---|
| Global | $NN_{	heta}(X) \in \mathbb{R}^{N_{classes}}$ | $NN_{\theta}(X) \in \mathbb{R}^{F_{out}}$ |
| Dense | $NN_{\theta}(X) \in \mathbb{R}^{N_{pix} \times N_{classes}}$ | $NN_{\theta}(\mathbf{X}) \in \mathbb{R}^{N_{pix} \times F_{out}}$ |

Table 1: The output's size of the neural network $NN_{\theta}(X)$ depends on the task to be solved. N_{pix} is the number of pixels in a HEALPix map, $N_{classes}$ is the number of classes in a classification task, and F_{out} is the number of variables to be predicted in a regression task. The output size of a global task does not depend on the input size, whereas a dense task asks for one prediction per pixel. A classification task asks for discrete (or probabilistic) predictions, whereas a regression task asks for continuous variables.

number of surveyed frequencies. F_{in} might also be the number of slices in the radial direction.

DeepSphere can perform dense or global predictions, for regression or classification. A typical global classification task is to classify maps into cosmological model classes [9]. A typical global regression task is to infer the parameters of a cosmological model from maps [11, 23]. Some dense regression tasks are denoising, interpolation of missing values, or inpainting parts of a map [40]. Segmentation and feature detection [41] are examples of dense classification tasks. The size of the output of the NN depends on the task. See Table 1.

Fully convolutional networks (FCNs) have been introduced by [42] and are mostly used for the semantic segmentation of images, a dense classification task. An example FCN for dense regression is

$$\mathbf{Y} = NN_{\theta}(\mathbf{X}) = (GC \circ \sigma \circ BN \circ GC)(\mathbf{X}) \in \mathbb{R}^{N_{pix} \times F_{out}},$$

where \circ denotes composition, i.e., $(f \circ g)(\cdot) = f(g(\cdot))$. If *P* layers are used, they have to be inverted via upsampling in an encoder-decoder architecture.

The set of input pixels which influence the value of an output pixel forms a *receptive field*. That field is isotropic and local, i.e., it forms a disk centered around the output pixel. The field's radius is influenced by the polynomial order K of GC layers (setting how far the convolution operation looks around), and the downsampling factor of P layers, if any. This radius should be large enough to capture statistics of interest. For example, a partial sky observation can provide only limited information of cosmological relevance. On the other hand, looking at the whole sky is often superfluous and waste computations, as most interactions are local. A data and task dependent trade-off is to be found. Note here that while global predictions use, by definition, the whole sky, dense predictions are not necessarily local.

By treating each output pixel independently and in parallel, this architecture is a principled way to perform rotation equivariant operations. Rotation equivariance means that the rotation operation commutes with the NN, i.e., a rotation of the input implies the same rotation of the output.⁹

Global tasks can be solved by averaging dense predictions [43, 44]. Global averaging is computed by the layer $AV : \mathbb{R}^{N_{pix} \times F} \to \mathbb{R}^{F}$. Doing so assumes the data is locally independent, and form different observations of an unknown process. Averaging over independent observations is a common way to reduce variance. An example FCN for global regression is

$$NN = AV \circ GC \circ \sigma \circ GC \circ P \circ \sigma \circ BN \circ GC.$$

Here, the FCN predicts parameters for many overlapping parts of the sky in parallel, then average those predictions to get one set of parameters for the whole sky. Global predictions are made invariant to rotation on the SO(3) group by averaging dense equivariant predictions. Rotation invariance means that rotating the input does not impact the prediction.

By replacing the average by a *FC* layer, one gets the standard convolutional neural network (CNN), which is a generalization of the FCN. Indeed, the *AV* layer is a *FC* layer where all the entries of *W* in (6) are equal to $1/N_{pix}$ and b = 0. By dropping rotation invariance, CNNs learn where to put attention on the domain. That is useful when pixels are not of the same importance, with respect to the task. For example, on images, the subject of interest is most often around the center of the picture. Hence, those pixels are more predictive than the ones in the periphery when considering image classification. On the sphere, location is important when analyzing weather data on the Earth for example, as oceans and mountains play different roles. An example of such an architecture is

$$NN = FC \circ \sigma \circ F \circ P \circ \sigma \circ BN \circ GC.$$

There are in general two ways to deal with symmetries one wants to be invariant to: (i) build them into the architecture, or (ii) augment the dataset such that a (more general) model learns them. With respect to rotation invariance, the FCN architecture is of the first kind, while the CNN is of the second kind. For tasks that are rotation invariant, FCNs hence need less training data as the rotation symmetry is backed in the architecture and need not be learned. This can be seen as intrinsic data augmentation, as a CNN would need to see many rotated versions of the same data to learn the invariance. Moreover, FCNs can accommodate maps with a varying number of pixels. Such global summarization as the AV layer is commonly used along graph convolutions to classify graphs of varying sizes (as it is invariant to vertex permutation) [45, 46]. As such, CNNs should only be used if rotation invariance is undesired.

All the above architectures can be used for classification (instead of regression) by appending a *S M* layer. An example FCN for global classification is therefore

$$NN = SM \circ AV \circ GC \circ \sigma \circ GC \circ P \circ \sigma \circ BN \circ GC.$$

Similarly, we emphasize that the use of FC and AV layers is the sole difference between a NN engineered for global or dense prediction.

2.9. Training

The cost (or loss) function $C(Y, \overline{Y}) = C(NN_{\theta}(X), \overline{Y})$ measures how good the prediction Y is for sample X, given the ground truth \overline{Y} . For a classification task, the cost is usually taken to be the cross-entropy

$$C(\boldsymbol{Y}, \bar{\boldsymbol{Y}}) = -\sum_{i=1}^{N_{pix}} \sum_{j=1}^{N_{classes}} \bar{y}_{ij} \log(y_{ij}),$$

where $\bar{Y} \in \mathbb{R}^{N_{pix} \times N_{classes}}$ is the ground truth label indicator, i.e., $\bar{y}_{ij} = 1$ if pixel *i* of sample *X* belongs to class *j* and is zero otherwise. For global prediction, we have $N_{pix} = 1$. For a regression task, a common choice is the mean squared error (MSE)

$$C(\mathbf{Y}, \bar{\mathbf{Y}}) = \frac{\|\mathbf{Y} - \bar{\mathbf{Y}}\|_{2}^{2}}{N_{pix}F_{out}} = \frac{1}{N_{pix}F_{out}} \sum_{i=1}^{N_{pix}} \sum_{j=1}^{F_{out}} (y_{ij} - \bar{y}_{ij})^{2},$$

where \bar{Y} is the desired output. Again, take $N_{pix} = 1$ for global regression. We emphasize that the cost function and the *S M* layer are the sole differences between a NN engineered for classification or regression.

The goal of training is to find the parameters θ of the NN that minimize the risk $R(\theta) = E\left[C\left(NN_{\theta}(X), \bar{Y}\right)\right]$, where E is the expectation over the joint distribution (X, \bar{Y}) . In general, that expectation cannot be computed as the data distribution is unknown. We can however minimize an approximation, the empirical risk over the training set $\{(X_i, \bar{Y}_i)\}_{i=1}^{N_{sumples}}$:

$$\hat{\theta} = \arg\min_{\theta} \sum_{i=1}^{N_{samples}} C\left(NN_{\theta}(\boldsymbol{X}_{i}), \bar{\boldsymbol{Y}}_{i}\right).$$

The optimization is performed by computing an error gradient w.r.t. all the parameters by back-propagation

⁹Small errors will however appear as *GC* layers are not exactly equivariant due to the small discrepancy between the graph Fourier modes and the spherical harmonics. See Appendix A.



Figure 6: Some pixelizations of the sphere. Left: the equirectangular grid, using equiangular spacing in a standard spherical-polar coordinate system. Middle: an equiangular cubed-sphere grid, as described in [47]. Right: graph built from a HEALPix pixelization of half the sphere ($N_{side} = 4$). By construction, each vertex has eight neighbors, except the highlighted ones which have only seven.⁵ Left and middle figures are taken from [20].

and updating them with a form of stochastic gradient descent (SGD):

$$\theta \leftarrow \theta - \frac{\eta}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \frac{\partial C\left(NN_{\theta}(X_i), \bar{Y}_i\right)}{\partial \theta}$$

where η is the learning rate, and \mathcal{B} is the set of indices in a mini-batch. Batches are used instead of single samples to gain speed by exploiting the parallelism afforded by modern computing platforms.

3. Related work

3.1. 2D convolutional neural networks

A first approach, explored by [20] for molecular modeling and [21, 22] for omnidirectional imaging, is to use a 2D CNN on a discretisation of the sphere that is a grid, such as the equirectangular projection (Figure 6, left panel), or a set of grids, such as the cubedsphere defined by [47] (Figure 6, middle panel). As this formulation uses the standard 2D convolution, all the optimizations developed for images can be applied, which makes it computationally efficient. This approach is applicable to the many pixelizations that are based on a regular subdivision of a base polyhedron, such as HEALPix. Each base polyhedron then forms a grid. Care has to be taken to handle boundary conditions: for example by padding a grid with the content of the opposite side (equirectangular) or of the neighboring grids (cubed-sphere, HEALPix). That incurs some computational losses. For samplings that are not equal area, such as the equirectangular projection, the convolution operation should be adjusted to take the induced distortion into account [21, 22].

Another approach to leverage 2D CNNs is to project spherical data onto many tangent planes, which are flat 2D surfaces. This approach has been extensively used for cosmological maps [9, 11, 12, 23] and omnidirectional imaging [48, 49]. This idea has been generalized to arbitrary 2D manifolds for shape alignment and retrieval [50–52].

The main issue with the above two approaches is that they depend on a (local) coordinate system to define anisotropic filters, i.e., filters which are direction dependent. Direction is well defined and matters for some applications, such as the analysis of weather and climate data on the Earth (north, south, east, west), and omnidirectional imaging (up, down). Indeed, rotation invariance has been shown to reduce discriminative power for omnidirectional imaging [22]. Some problems are, however, intrinsically invariant (or equivariant) to rotation. Examples include the analysis of cosmological maps, and the modeling of atoms and molecules. In such cases, directions are arbitrarily defined when setting the origin of the pixelization. Therefore, a convolution operation has to be isotropic to be equivariant to rotation.

3.2. Spherical neural networks

Rotation equivariance was addressed by leveraging the convolution associated to the 3D rotation group SO(3), with applications to atomization energy regression and 3D model classification, alignment and retrieval [24, 25]. The resulting convolution is performed by (i) a spherical harmonic transform (SHT), i.e., a projection on the spherical harmonics, (ii) a multiplication in the spectral domain, and (iii) an inverse SHT. Note the similarity with the naive graph convolution defined in (1). Likewise, the computational cost of a convolution is dominated by the two SHTs, and a naive implementation of the SHT costs $O(N_{pix}^2)$ operations. Accelerated schemes however exist for some sampling sets [see 39, 53, 54, for examples]. The convolutions remain nevertheless expensive, limiting the practical use of this approach. For example with HEALPix, which was designed to have a fast SHT by being iso-latitude, the computational cost of the SHT is $O(N_{pix}^{3/2}) = O(N_{side}^3) =$ $O(\ell_{max}^3)$, where ℓ_{max} is the largest angular frequency [27, 39].¹⁰ While the Clebsh-Gordan transform can be

¹⁰All pixels are placed on $N_{ring} = 4N_{side} - 1 = O(\sqrt{N_{pix}})$ rings of constant latitude. Each ring has $O(\sqrt{N_{pix}})$ pixels. Thanks to this iso-latitude property, the SHT is computed using recurrence relations for Legendre polynomials on co-latitude and fast Fourier transforms (FFTs) on longitude. The computational cost is thus $\sqrt{N_{pix}}$ FFTs for a total cost of $O(N_{pix} \log \sqrt{N_{pix}})$, plus $\sqrt{N_{pix}}$ matrix multiplications of size $\sqrt{N_{pix}}$ for a total cost of $O(N_{pix}^{3/2})$ operations.

leveraged to avoid SHTs between layers (it does a nonlinear transformation in the spectral domain), the transform itself costs $O(N_{pix}^{3/2})$, for no reduction of overall complexity [55]. This work is another indication that there might be a $\Omega(N_{pix}^{3/2})$ computational lower bound for the proper treatment of rotation equivariance with the spherical harmonics. In comparison, DeepSphere scales as $O(N_{pix})$ (see Section 2.5), at the expense of not being exactly equivariant (see Appendix A). Being a mathematically well-defined rotation equivariant network is the main advantage of methods based on spherical harmonics, though the fact that convolution kernels are learned probably compensates for inexact equivariance. See [56] for a rigorous treatment of convolution and equivariance in NNs.

While defining the convolution in the spectral domain avoids all sampling issues, filters so defined are not naturally localized in the original domain. Localization is desired for the transformation to be stable to local deformation [57], and most interactions in the data are local anyway. A straightforward way to impose locality in the original domain is to impose smoothness in the spectral domain, by Heisenberg's uncertainty principle.¹¹ A more elegant approach is to define filters that are provably localized. The filters defined in (3) and (4), by being polynomials of the Laplacian, are of this kind.

All the presented 2D and spherical CNNs cannot be easily accelerated when the data lies on a part of the sphere only. That is an important use case in cosmology as measurements are often partial, i.e., whole sky maps are rare. One could still fill the unseen part of the sphere with zeros, and avoid computing empty integrals. It is, however, not straightforward to identify empty space, and computations would still be wasted (for example on a ring that mostly contains zeros but a few measurements). With graphs, however, computations are only performed for used pixels. While it results in some distortions due to border effects (see Figure C.15 and Appendix C), these can be mitigated by padding with zeros a small area around the measurements.

3.3. Graph neural networks

The use of a graph to model the discretised sphere was also considered for omnidirectional imaging [35]. This work is the closest to our method, with three differences. First, they parametrize their convolution kernel with (3) instead of (4) (see Section 2.5 for a discussion and comparison). Second, they did not take

advantage of a hierarchical pixelization of the sphere and resorted to dynamic pooling [58]. While that operation has proved its worth to pool sequences of varying length, such as sentences in language models, it is undesired in our context as it is not local and destroys spatial coherence. Third, they introduced a statistical layer an operation that computes a set of statistics from the last feature maps — to provide invariance to rotation. We propose to use the idea of fully convolutional networks (FCNs) instead (see Section 2.8). While statistics have to be hand-chosen to capture relevant information for the task, the filters in a FCN are trained end-to-end to capture it.

Many formulations of graph neural networks, reviewed by [59] and [60], have been proposed. For this contribution, we chose the formulation of [26] as its root on strong graph signal processing theory makes the concept of convolutions and filters explicit [36]. As the convolution is motivated by a multiplication in the graph Fourier spectrum, it is close in spirit, and empirically, to the formulation based on the spherical harmonics, which is the ideal rotation equivariant formulation.

Thanks to their versatility, graph neural networks have been used in a variety of tasks, such as identifying diseases from brain connectivity networks [61] or population graphs [62], designing drugs using molecular graphs [63], segmenting 3D point clouds [64], optimizing shapes to be aerodynamic [65], and many more. By combining graph convolutional layers and recurrent layers [66], they can, for example, model structured time series such as traffic on road networks [67], or recursively complete matrices for recommendation [68]. Another trend, parallel to the modeling of structured data, is the use of graph neural networks for relational reasoning [69].

4. Experiments

The performance of DeepSphere is demonstrated on a discrimination problem: the classification of convergence maps into two cosmological model classes. The experiment presented here is similar to the one by [9]. These maps are similar to those created with gravitational lensing techniques [28]. The two sets of maps were created using the standard cosmological model with two sets of cosmological parameters (see Section 4.1 for details). The classification methods were trained to predict a label from a HEALPix map. Our Python implementation to reproduce those experiments

¹¹Heisenberg's uncertainty principle states that a filter cannot be arbitrarily concentrated in one domain without being de-concentrated in the other.



Figure 7: Power spectral densities of the noiseless maps. To prevent the cosmological models from being distinguished by their power spectra alone, the maps have been smoothed with a Gaussian kernel of radius 3 arcmins to remove high frequencies ($\ell > 1000$).

is openly available online.¹² The data is available upon request.¹³

While we only demonstrate a classification task here, other tasks, such as regression or segmentation, are also possible (see Section 2.8). The regression task will most likely be the most practical cosmological application, as described in [11] and [23]. Implementation of full cosmological inference typically requires, however, many more simulations, building the likelihood function, and several other auxiliary tasks. The classification problem is much more straightforward to implement and execute, and can be used to fairly compare the accuracy of the algorithm against benchmarks [9]. For these reasons we decided to use the classification problem in this work, and we expect the relative performance of the methods to generalise to the regression scenario.

4.1. Data

Convergence maps represent the dimensionless distribution of over- and under-densities of mass in the universe, projected on the sky plane. The 3D structures are projected using a geometric kernel, the value of which depends on the radial distance. In gravitational lensing, this kernel is dependent on the radial distances between the observer, the mass plane, and the plane of source galaxies [see 70, for a review of gravitational lensing].

We make whole sky N-body simulations for two parameter sets of the Λ CDM cosmological model: model 1 ($\Omega_m = 0.31, \sigma_8 = 0.82$) and model 2 ($\Omega_m = 0.26, \sigma_8 = 0.91$), where Ω_m is the matter density in the universe and σ_8 is the normalisation of the matter

power spectrum. Other parameters are set to: Hubble constant $H_0 = 70 km/s/Mpc$, spectral index $n_s = 0.96$, and Baryon density today $\Omega_b = 0.05$. The parameters Ω_m and σ_8 were chosen for the maps to have the same spherical harmonic power spectrum. That means that it is difficult to distinguish between these cosmological models. We found that the differences in power spectrum is 5% for ℓ > 1000. To remove this information, we additionally smooth the spherical maps with a Gaussian kernel of radius 3 arcmin. The resulting power spectral density (PSD), computed using the anafast function of the HEALPix package, are displayed in Figure 7. We also subtract the mean of each map and downsample them to a resolution of $N_{side} = 1024$, which corresponds to maps of $12 \times 1024^2 \approx 12 \times 10^6$ pixels. As shown by the occupied spectrum (Figure 7), a larger resolution would waste memory and computation, while a lower resolution would destroy information.

The simulations are created using the fast lightcone method UFalcon described in [71]. A brief overview about the map making procedure used in UFalcon as well as the simulation parameters are given in Appendix F. We however use a single simulation box, as opposed to two used in that work, as we use source galaxies at a lower redshift of z = 0.8, instead of z = 1.5. L-PICOLA [72] is used for fast and approximate N-body simulations. We generate 30 simulations for each of the two classes. Out of the 60 simulations, 20 are kept as the test set, and 20% of the remaining training data is used as a validation set, to monitor the training process and select the hyper-parameters. Figure 8 shows the whole sky simulations and a zoom region for both models. Initial conditions for these simulations are the same, so the differences in structures can only be attributed to different cosmological parameters used to evolve the particle distribution.

4.2. Problem formulation

As the distribution of matter in the universe is homogeneous and isotropic, no pixel is more informative than any other.¹⁴ As such, we can control the difficulty of the classification problem by limiting the number of pixels available to the algorithms, i.e., extracting partial maps from whole sky maps. Using the properties of HEALPix, we split maps into $12 \times o^2$ independent samples of $(N_{pix}/o)^2$ pixels (for o = 1, 2, 4, ...). The resulting partial maps are large enough to suffer from the effects of the spherical geometry, and cover 8.3%

¹²https://github.com/SwissDataScienceCenter/ DeepSphere

¹³https://doi.org/10.5281/zenodo.1303272

¹⁴Contrast that with images, where the subject of interest is most often around the center of the picture.



Figure 8: Example maps from two classes to be discriminated. Left: model 1 with $\Omega_m = 0.31$ and $\sigma_8 = 0.82$. Right: model 2 with $\Omega_m = 0.26$ and $\sigma_8 = 0.91$. The initial conditions for both simulations are the same, so differences only arise due to different cosmological parameters.

($\approx 1 \times 10^6$ pixels), 2.1% ($\approx 260 \times 10^3$ pixels), and 0.5% ($\approx 65 \times 10^3$ pixels) of the sphere for o = 1, 2, 4, respectively. Corresponding areas can be seen in Figure 5: the surface of pixels in the left, middle, and right spheres correspond to samples at order o = 4, 2, 1, respectively. We report results for o = 1, 2, 4 only, as full sphere classification is easy at such resolution (perfect classification accuracy is already obtained at o = 1, see Figure 9). The published code nonetheless includes an example demoing classification on full spheres of lower resolution ($N_{side} = 64$).

To make the discrimination harder and the problem more realistic, centered Gaussian noise is added to the simulated maps. The standard deviation of the noise varies from zero (i.e., no noise) to 2× the standard deviation of pixel's values in the noiseless maps. While the noise model of real maps often has a slightly different distribution, Gaussian noise should be a sufficient model to demonstrate the performance of our method. To avoid over-fitting, random noise is generated during training, so that no two samples are exactly the same (i.e., two samples might be from the same simulation, but the added noise will be different). That is a data augmentation scheme, as it creates more training examples than we have simulations.

4.3. Baselines

DeepSphere is first compared against two simple yet powerful benchmarks. The two baselines are based on features that are (i) the power spectral densities (PSDs) of maps, and (ii) the histogram of pixels in the maps [73]. After normalization, those features are used by a linear support vector machine (SVM) trained for classification. A linear kernel is used as other kernels did not provide better results, while having significantly worse scaling properties. For fairness, the training set was augmented in a similar way as for DeepSphere: we created samples by adding new random realizations of the noise. We stopped adding new samples to the training data once the validation error converged. This process is detailed in Appendix E. As for DeepSphere, the SVM regularization hyper-parameter was tuned based on the performance over the validation set. The classification accuracy of DeepSphere and both baselines was checked as a function of the used area (the order *o*), and the relative level of additive noise.

DeepSphere is further compared to a classical CNN for 2D Euclidean grids, later referred to as 2D ConvNet. To be fed into the 2D ConvNet, samples have to be transformed into flat images. As described in Appendix D, we use the property that HEALPix is defined as the subdivision of 12 square surfaces to efficiently project the spherical maps. This property has been independently exploited (while this paper was under review) to define another spherical CNN [74]. In our experiments, we used the same architecture as DeepSphere (described in Section 4.4), with convolution kernels of size 5×5 and an $\ell 2$ weight decay with a weight of 3 as regularization. The number of filters was reduced in order to match the number of trainable parameters between DeepSphere and the 2D ConvNet. (A larger architecture resulted in unstable training for a marginal performance gain.)

Ideally, we would compare DeepSphere to alternative spherical CNNs like [24] and [25]. There are however two issues with that comparison. First, these CNNs were developed for the equirectangular sampling. An option is to modify the available implementations to use HEALPix, which is a major undertaking due to their complexity. Another option is to transform the cosmological maps to the equirectangular sampling, which is of disputable usefulness as the field is unlikely to use that sampling. Second, those methods, based on SHTs, only work on the full sphere (at least in their current state). As shown in Figure 4, using the full sphere at a resolution of $N_{side} = 1024$ would result in a slow-down of multiple orders of magnitude compared to using a partial graph. The authors of [75] could not use [24] for CMB analysis because of that.¹⁵ For those reasons and because we think that a proper comparison is better carried out on diverse datasets and tasks, we leave it as future work.

4.4. Network architecture and hyper-parameters

As discussed in Section 2.8, the choice of a network architecture depends on the data and task at hand. For our problem, where we assume that each pixel carries the same quantity of information about the cosmological model, a rotation invariant FCN architecture should be best. The selected network is defined as

$$FCN = \underbrace{SM \circ AV \circ GC_2}_{\text{classifier}} \circ \underbrace{L_5 \circ L_4 \circ L_3 \circ L_2 \circ L_1}_{\text{feature extractor}}$$
with
$$L_1 = P_4 \circ \sigma \circ BN \circ GC_{16},$$

$$L_2 = P_4 \circ \sigma \circ BN \circ GC_{32},$$

$$L_3, L_4, L_5 = P_4 \circ \sigma \circ BN \circ GC_{64},$$

where GC_F indicates a graph convolutional layer with F output feature maps, P_4 is a pooling layer that divides the number of pixels by 4, σ is a ReLU, and the Chebyshev polynomials in (4) are of degree K = 5. The layers L_1 to L_5 build the statistical evidence to classify each pixel of a down-sampled sphere. The last GC layer acts as a linear classifier and outputs two predictions. The AV layer averages these predictions and the SM layer normalizes them.

For comparison, we tried the more conventional CNN architecture, where the classification is performed by a fully connected layer from the last feature maps instead of the average of local classifications. This architecture is not invariant to rotation. The selected network is defined as

$$CNN = \underbrace{SM \circ FC_2}_{\text{classifier}} \circ \underbrace{L_5 \circ L_4 \circ L_3 \circ L_2 \circ L_1}_{\text{feature extractor}},$$

where FC_2 denotes a fully connected layer with two outputs. Architectures were selected over their performance on the validation set.

The Adam optimizer [76] with $\beta_1 = 0.9$, $\beta_2 = 0.999$ was used with an initial learning rate of 2×10^{-4} that is

exponentially decreased by a multiplication with 0.999 at each step. All models were trained for 80 epochs, which corresponds to 1920 steps. The batch size is set to $16 \times o^2$ to keep the amount of supervision used to estimate the gradient identical, irrespective of the chosen order *o* (the learning rate would otherwise need to be adapted). By this choice, DeepSphere is trained with the same amount of information, i.e., the number of pixels, across all variants of the problem. Training took approximately 1.5 hour using a single Nvidia GeForce GTX 1080 Ti and 5 hours with a Tesla K20.

4.5. Results

Figure 9 compares the classification accuracy of DeepSphere (FCN and CNN variants), and the three baselines across five noise levels and three sample sizes. We indeed observe that the problem is made more difficult as the sample size decreases, and the noise increases. As fewer pixels are available to make a decision about a sample, the algorithms have access to less information and thus cannot classify as well. As the noise increases, the useful information gets diluted in irrelevant data, i.e., the signal-to-noise ratio (SNR) diminishes.

As the cosmological parameters were chosen for the maps to have similar PSDs (see Figure 7), it is reassuring to observe that an SVM with those as features has difficulties discriminating the models. Other statistics are therefore needed to solve the problem. Using histograms as features gives a significant improvement. Performance is very good for larger maps and deteriorates for the smaller cases with increased noise level. reaching 80% at the highest noise level for o = 4. The histogram features contain information about the distribution of pixels, which clearly varies for the two classes. They do not, however, include any spatial information. DeepSphere (both the FCN and CNN variants) shows superior performance over all configurations. The gap widens as the problem becomes more difficult. This is likely due to DeepSphere learning features that are adapted to the problem instead of relying on hand-designed features. The accuracy of Deep-Sphere is > 97% for orders o = 1 and o = 2, for all noise levels, and starts to deteriorate for o = 4, reaching 90% for the highest noise level.

It may seem unfair to compare DeepSphere, who has access to raw pixels, with SVMs who only see limited features (histograms and PSDs). We however trained an SVM on the raw pixels and were unable to obtain over 60% accuracy in any of the three noiseless cases.

The 2D ConvNet fairs in-between the SVMs and DeepSphere. Its relatively low performance probably

¹⁵Note that reported results in [24] and [25] are for maps of at most $(2 \times 128)^2 = 65,536$ pixels. That is the size of our smallest partial map, while the largest is of approximately one million pixels.



Figure 9: Classification accuracy of the fully convolutional variant of DeepSphere (DeepSphere FCN), the standard convolutional variant of DeepSphere (DeepSphere CNN), the fully convolutional 2D ConvNet (2D ConvNet FCN), the standard 2D ConvNet (2D ConvNet CNN), the support vector machine (SVM) with the power spectral density (PSD) as features, and the SVM with the histogram as features. The difficulty of the task depends on the level of noise and the size of a sample (that is, the number of pixels that constitute the sample to classify). Order o = 1 corresponds to samples which area is $\frac{1}{12} = 8.1\%$ of the sphere, order o = 2 to $\frac{1}{12\times2^2} = 2.1\%$, and order o = 4 to $\frac{1}{12\times4^2} = 0.5\%$. The standard deviation of the added Gaussian noise varies from zero to 2× the standard deviation of pixel's values in the noiseless maps. It is clear from those results that the noise makes the problem harder, and that having more pixels available to classify a sample makes the problem easier (the classifier having more evidence to make a decision). The FCN variant of DeepSphere beats the CNN variant by being invariant to rotation. Both variants largely beat the 2D ConvNet and the two SVM baselines.

come from the following drawbacks. First, it does not exploit the rotational invariance of the problem. We observed that the learned convolution kernels were not radial. Second, the 2D projection distorts the geometry, and the NN has to learn to compensate for it, which comparatively requires more training data.

As expected, the FCN variants outperform the CNN variants (both for DeepSphere and the 2D ConvNet). This may seem counterintuitive as the CNN is a generalization of the FCN and hence should be able to learn the same function and provide at least equivalently good results. Nevertheless, as discussed in Section 2.8, the larger number of parameters incurred by the increased flexibility requires more training data to learn the rotation symmetry. The superior performance of the FCN is an empirical validation that these maps are stationary with a small-support radial autocorrelation function, stemming from the fact that the mass distribution is homogeneous and isotropic. It also implies that this classification problem is invariant to rotation. Hence, a pixel can be statistically classified using only its surrounding, and those local predictions can be averaged to vote for a global consensus. The CNN variant, however, may be better for data that does not have this property, as the architecture should always be adapted to the data and task (see Section 2.8).

While testing the FCN and CNN variants of Deep-Sphere, we made the following empirical observations. First, training was more stable (in the sense that the loss decreased more smoothly) when using Chebyshev polynomials, as in (4), rather than monomials, as in (3). Nevertheless, we could not observe a significant difference in accuracy after convergence. Second, using ℓ_2

regularization does not help either with performance or training of the models, mostly because the models are not over-fitting. Third, we recommend initializing the Chebyshev coefficients with a centered Gaussian distribution with standard deviation $\sqrt{2/(F_{in} \times (K + 0.5))}$, where *K* is the degree of the Chebyshev polynomials and F_{in} the number of input channels. This standard deviation has the property of keeping the energy of the signal more or less constant across layers.¹⁶ Finally, we observe that scaling the Laplacian's eigenvalues between [-a, a], where $0 < a \le 1$, significantly helps in stabilizing the optimization. We use a = 0.75 in our experiments.

4.6. Filter visualization

A common visualization to introspect and try to understand how a CNN function is to look at the learned filters. Since our construction leads to almost (up to sampling irregularities) spherical filters, we plot in Figure 10 both the radial profile and a gnomonic projection on a plane of a random selection of learned filters from the last layer of the network. While those particular filters were obtained from the experiment with order o = 2 and a relative noise level of 2, all trained networks presented visually similar patterns. Details of how the convolution kernels are plotted are described in Appendix B. While it is usually difficult to interpret the shape of the filters, especially given the type of data, we can notice that they often have a "peak"-like structure. An example of filter interpretation was demonstrated in [78].

¹⁶We derived this rule from the principles presented in [77], the Chebyshev polynomial equations, and some empirical experiments.



Figure 10: Random selection of 24 learned filters from the fifth spherical convolutional layer L_5 . Top: section. Bottom: gnomonic projection. The structure in the filters resembles peaks, which is not unexpected, given that the convergence maps largely consist of concentrated clumps.

5. Conclusion

We present DeepSphere, a convolutional neural network defined on the sphere with HEALPix sampling, designed for the analysis of cosmological data. The main contributions of this paper are (i) to show that spherical CNNs are a great NN architecture for cosmological applications, and (ii) that a graph-based spherical CNN has certain undeniable advantages. The core of our method is the use of a graph to represent the discretised sphere. This allows us to leverage the advantages of the graph convolution. It is both efficient, with a complexity of $O(N_{pix})$, and flexible, which allows DeepSphere to efficiently work on a partial sphere. The spherical properties of the domain are well captured: the graph Fourier modes are close to the spherical harmonics. Filters are restricted to be radial for the convolution operation to be equivariant to rotation. DeepSphere can then be made either equivariant or invariant to rotation. Equivariance is not perfect as small imprecisions due to the sampling cause the action of a graph filter to slightly depend on the location. However, we do not expect that to cause problems for practical applications.

We demonstrate that DeepSphere systematically and significantly outperforms three benchmark methods on an example problem of cosmological model discrimination with weak lensing mass maps, designed similarly to [9]. The maps were produced from two cosmological models with varying σ_8 and Ω_m , chosen to follow the typical weak lensing degeneracy in these parameters, so that the power spectra for these models are similar. We compared the performance of DeepSphere versus that of a classical 2D CNN and two SVM classifiers (one trained on the spherical harmonics power spectrum, and the other on the pixel density histogram). DeepSphere performs better than the three baselines for all considered cases. The advantage is small for large, noise-free maps, and grows up to 10% for smaller, noisier data.

Spherical CNNs are so far mostly used for omnidirectional imaging. Many scientific fields, such as weather or climate modelling, would however benefit from an efficient and versatile spherical CNN. With this work, we hope to demonstrate and help democratize the use of those tools for spherical data analysis. We publish the code as a small and easy-to-use python package. The code was designed so that DeepSphere can easily be used for typical machine learning tasks, such as classification and regression, both dense and global.

As future work, it would be interesting to further investigate the relation between the graph Fourier basis and the spherical harmonics. The goal would be to find edge weights such that the graph Laplacian converges (or is equivalent) to the Laplace-Beltrami (up to a certain bandwidth). Ideally, that should work for any sampling of the sphere. That would enable a truly rotation equivariant graph convolution and high-precision filtering using the graph rather than the SHT (for speed). Finally, a comparison of DeepSphere to other spherical CNN formulations, with different sampling schemes, would be worthwhile.

The fact that a graph representation of the discretised sphere enables an efficient convolution relaxes the isolatitude constraint on the design of sampling schemes which aim to enable fast convolutions. In the long term, graphs might enable researchers to consider sampling schemes with different trade-offs, or remove the need for schemes and interpolation altogether and simply consider the positions at which measures were taken.

Acknowledgements

This work was supported by a grant from the Swiss Data Science Center (SDCS) under project *DLOC: Deep Learning for Observational Cosmology* and grant number 200021_169130 from the Swiss National Science Foundation (SNSF). We thank Jean-François Cardoso for helpful discussions about the spherical harmonics. We thank Alexandre Refregier, Adam Amara, Thomas Hofmann, and Fernando Perez-Cruz for advice and helpful discussions. We thank Hamsa Padmanabhan for testing use cases of the code. Finally, we thank the two anonymous reviewers who provided extensive feedback that greatly improved the quality of this paper.

Appendix A. Graph Fourier modes and spherical harmonics

The first 16 eigenvectors $[u_1, \ldots, u_{16}]$ of the graph Laplacian L, forming the lower part of the graph Fourier basis U, are shown in Figure 3. Let us further observe the spectral properties of our constructed spherical graph laplacian L. Its eigenvalues, shown in Figure A.11, are clearly organized in frequency groups of $2\ell + 1$ orders for each degree ℓ . We remind the reader that the amplitude of the Laplacian eigenvalue is proportional to the sum of the variations of its associated eigenvector. Furthermore, all spherical harmonics with the same order ℓ have the same variation. Hence, the fact that the Laplacian eigenvalues are grouped in blocks of size $2\ell + 1$ is a hint that the graph eigenvectors approximate the spherical harmonics. In order to push the comparison one step further, we show the correspondence between the subspaces spanned by





Figure A.11: The eigenvalues Λ of the graph Laplacian $L = U\Lambda U^{\mathsf{T}}$, which corresponds to squared frequencies, are clearly organized in groups. Each group corresponds to a degree ℓ of the spherical harmonics. Each degree has $2\ell + 1$ orders. See also Figure 3.



Figure A.12: Correspondence between the subspaces spanned by the graph Fourier modes and the spherical harmonics. First, we compute the power spectral density (PSD) of each graph eigenvector with the SHT. Second, as there is $2\ell + 1$ spherical harmonics for each degree ℓ , we sum the contributions of the corresponding $2\ell + 1$ graph eigenvectors. The matrix shows how the subspaces align: the Fourier modes span the same subspaces as the spherical harmonics in the low frequencies, and the Fourier modes leak towards adjacent frequency bands at higher frequencies. While there is a systematic error, the Fourier modes align at higher frequencies as N_{side} increases.

the graph Fourier modes and the spherical harmonics in Figure A.12.

While these indications show that the constructed graph Fourier basis approximates well the spherical harmonics, one should not forget that the small irregularities in the sampling (non-constant number of neighboring pixels and varying distance between pixels, see Figure 6) have an important effect on the graph Fourier modes. First, we believe that they are responsible for energy leaking across frequency bands in Figure A.12. Second, as the resolution is increased with $N_{side} \rightarrow \infty$ and $N_{pix} \rightarrow \infty$, we are still unsure if the eigenvectors would converge towards the spherical harmonics. The theoretical study of those phenomenons is left as future work. Third, counter-intuitively, some eigenvectors will be localized [79], i.e., they will span a small part of the sphere. Those discrepancies result in a convolution operation that is not exactly equivariant to rotation. Nevertheless, our experiments suggest that these downsides do not have an important effect on the convolution nor hinder the performance of the NN.

Appendix B. Example: heat diffusion

Let us consider the heat diffusion problem

$$\tau L f(t) = -\partial_t f(t), \qquad (B.1)$$

where $f : \mathbb{R}_+ \to \mathbb{R}^{N_{pix}}$. Given the initial condition f(0), the solution of (B.1) can be expressed as

$$f(t) = e^{-L\tau t} f(0) = U e^{-\Lambda\tau t} U^{\mathsf{T}} f(0) = K_t(L) f(0),$$

which is, by definition, the convolution of the signal f(0) with the kernel $K_t(\lambda) = e^{-\tau t \lambda}$. Since the kernel K_t is applied to the graph eigenvalues Λ , which can be interpreted as squared frequencies, it can also be considered as a generalization of the Gaussian kernel on the sphere. Figure B.13 shows the effect of the convolution by diffusing a unit of heat for $\tau = 1$ at various times t. By comparing the graph convolution with the spherical symmetric Gaussian smoothing, we observe that both techniques lead to similar results (see Figure B.13). While the graph convolution remains different from the spherical convolution, this small experiment shows that, providing the correct parameters, the graph convolution can approximate the spherical convolution. Figure B.14 should give more insights about the graph kernel K_t .



Figure B.13: Comparison of convolution with the graph and the spherical harmonics ($N_{side} = 16$). Top: Diffusion of a unit of heat for different times *t* using the graph. Bottom: spherical symmetric Gaussian smoothing for different σ (arcmin). Relative difference between graph convolution and spherical smoothing: 10.4%, 4.8%, 3.8%.

Furthermore, the harmonic resemblance is another sign that the constructed graph is able to capture the spherical structure of the HEALPix sampling. In some applications, where the exactitude of the convolution is not a requirement, such as de-noising, graph convolution could be used instead of using the spherical harmonics. In a neural network setting, the fact that the graph convolution is not exactly equivariant to rotation is probably compensated by the fact that convolution kernels are learned.

We further leverage this example to present three visualizations of a convolution kernel. The first visualization (Figure B.14, top) is a plot of the kernel function, i.e., it shows K_t evaluated at the graph eigenvalues



Figure B.14: Visualization of the convolution kernel $K_t(x) = e^{-\tau t x}$. Top: spectral domain. Middle: gnomonic projection. Bottom: cross-section along the equator.

diag(Λ). Note that small eigenvalues λ correspond to small frequencies, i.e., to spectral modes of low variation. The second visualization (Figure B.14, middle) shows the kernel localized on the sphere (gnomonic projection). As the kernel K_t is defined in the graph spectral domain, it was convolved with a Kronecker δ to obtain a representation on the sphere (see Section 2.4). The third visualization (Figure B.14, bottom) plots the section of the filter. This visualization might be easier to read as the filters are isotropic (up to the irregularities of the sampling). The plot is obtained by convolving a Kronecker δ on the vertices along the equator. Note that, because of the small irregularities in the HEALPix sampling, the second and third methods are subject to small variations depending on the chosen position of the Kronecker δ .

Appendix C. Border effects

The graph setting used throughout this contribution corresponds to assumed reflective border conditions. While that is irrelevant when working on the complete sphere (as it has no border), it slightly affects the convolution operation when only a part of the sphere is considered. As depicted in Figure C.15, a filter localized near a border (via $h(L)\delta_i$) is no longer isotropic. These border effects can, however, be mitigated by padding

with zeros a small area around the part of interest (in which case they become similar to border effects in classical CNNs). We however do not expect these effects to cause any problem as long as the data samples cover the same area.



Figure C.15: Convolution kernel (also called filter) localized in the center and left corner of a graph built from $1/12^{\text{th}}$ of the sphere at $N_{side} = 16$. A filter *h* is localized on a pixel *i* as $\mathcal{T}_i h = h(\mathbf{L})\delta_i$ (see Section 2.4, equation (2)). The filter is not isotropic anymore when localized on the corner as the graph representation of a manifold assumes reflective border conditions.

Appendix D. Projection onto a 2D plane

In order to compare the DeepSphere architecture with a 2D CNN, we needed to project the spherical maps onto a plane. A radial projection is computationally so expensive that, because of data augmentation, it is significantly more expensive than the actual training of the NN. Hence, we decided to leverage the geometrical properties of HEALPix to create another projection. Remember that HEALPix models the sphere as 12 pixels that are then subdivided into 4 recursively. Eventually, each of these 12 bases resembles a square grid. See Figure C.15 to visualize it. As our classification problem involves a twelfth of the sphere at most, we can use this natural structure to simply assign each HEALPix pixel to a pixel of a 2D grid of size $N_{sides} \times N_{sides}$. This operation is computationally cheap and probably spreads the projection error relatively evenly across the resulting image. This projection has actually been independently proposed (while this paper was under review) for a spherical CNN on HEALPix for cosmology [74].

Appendix E. Augmentation of the dataset for the SVM classifier

For the trained network to be robust to noise, we employ the following data augmentation technique: a random realization of Gaussian noise is added to each sample before it is fed to the NN. Besides robustness to noise, the network is less likely to over-fit the training set as it never sees the same sample twice. For a fair comparison between DeepSphere and the baselines, we have to ensure that the SVM classifier has access to the same amount of training data as Deep-Sphere. That is potentially an infinite number of samples. Hence, we fit different classifiers using various training set sizes until we experimentally observe that increasing the amount of training data does not improve performance. An example convergence of the training and validation errors is shown in Figure E.16.



Figure E.16: Error w.r.t. the number of training samples for histogram features and the linear SVM (setting: order o = 2, relative noise level 1.5). The error clearly converged such that adding more samples will not improve the classifier. The validation error might be slightly below the training error as it was used to select the hyper-parameters.

Appendix F. Convergence Mass Maps

In the following we describe the weak lensing mapmaking pipeline UFalcon (ultra fast lightcone) developed in [71] and the way it was used to generate data for this work. The data consist of two-dimensional mass maps on the full sphere, which represent the dimensionless, weighted and projected mass along the light of sight. The matter content of the universe consists mostly of dark matter. This kind of projected mass maps are typically measured using the weak gravitational lensing technique. Gravitational lensing occurs when images of distant galaxies are distorted by the matter clumps between the galaxy and the observer, which act as lenses. Measurements of the small, spatially coherent distortions of galaxy images can be used to infer the matter density between the observer and the background galaxies. [see 70, for a review of gravitational lensing].

In order to obtain a convergence map characterising the gravitational lensing of distance sources, one has to start with simulations of the dark matter distribution. This is achieved through cosmological N-Body simulations, which evolve a chosen number of dark matter particles $N_{\text{part}}^{\text{sim}}$ within a specific simulation volume V_{sim} under the influence of gravity across cosmic time-scales. The simulation outputs the positions of the particles at user-specific time-steps. These time-steps correspond to redshifts *z*, by which the spectrum is shifted the observer would measure. A redshift of z = 0 correspond to the present time and higher redshifts to earlier times.

Since the convergence map characterises the dark matter lenses at different redshifts, one has to first construct a lightcone. The lightcone represents the volume of the universe around the observer, where the radius represents the distance the light from distant galaxies at different times and locations has to travel to the observer. Therefore, the particles outputted at different redshifts (time-steps) of the simulation are concentrically arranged in shells of constant redshift around the observer (each redshift *z* corresponds to a comoving distance away from the observer $\chi(z)$), which is located at redshift z = 0 (see Figure F.17). The shells are then weighted by a weak lensing specific weight W_b and projected onto the sphere for each pixel θ_{pix} , resulting in a convergence map

$$\kappa(\theta_{\rm pix}) \sim \sum_{b} W_{b} \frac{V_{\rm sim}}{N_{\rm part}^{\rm sim}} \frac{n_{p}(\theta_{\rm pix}, \Delta \chi_{b})}{\mathcal{D}^{2}(z_{b})}, \qquad (F.1)$$

where n_p represents the number of particles in a specific pixel on the shell $\Delta \chi_b$. The convergence map depends on the cosmological parameters.



Figure F.17: Left: Dark matter density field simulated using L-PICOLA for specific cosmological parameters. Right: Sketch of the lightcone constructed by arranging particles at different redshifts in concentric shells around the observer located at z = 0.

The data generated for the present analysis is based on the fast approximate N-Body simulation code L-PICOLA [72]. For our purposes, we chose a simulation volume $V_{\rm sim} = (4200 \, h^{-1} \, \text{Mpc})^3$, a total number of particles $N_{\rm part}^{\rm sim} = 1024^3$ placed on a mesh of size $N_{\rm mesh}^{\rm sim} = 2048^3$. The simulation was then performed 30 times with different seeds for each of the two Λ CDM cosmological models described in section Section 4.1 covering a redshift range from z = 0.1 to z = 0.8. The L-PICOLA were performed in lightcone mode, where the simulation output is already arranged as a lightcone. Here we chose a redshift-shell thickness of $\Delta z = 0.01$.

Bibliography

- Planck Collaboration, Planck 2015 results. XIII. Cosmological parameters, A&A594 (2016) A13.
- [2] E. Komatsu, K. M. Smith, J. Dunkley, C. L. Bennett, B. Gold, G. Hinshaw, N. Jarosik, D. Larson, M. R. Nolta, L. Page, D. N. Spergel, M. Halpern, R. S. Hill, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, G. S. Tucker, J. L. Weiland, E. Wollack, E. L. Wright, Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation, ApJS192 (2011) 18.
- [3] S. Staggs, J. Dunkley, L. Page, Recent discoveries from the cosmic microwave background: a review of recent progress, Reports on Progress in Physics 81 (4) (2018) 044901.
- [4] M. Santos, P. Bull, D. Alonso, S. Camera, P. Ferreira, G. Bernardi, R. Maartens, M. Viel, F. Villaescusa-Navarro, F. B. Abdalla, M. Jarvis, R. B. Metcalf, A. Pourtsidou, L. Wolz, Cosmology from a SKA HI intensity mapping survey, Advancing Astrophysics with the Square Kilometre Array (AASKA14) (2015) 19.
- [5] HI4PI Collaboration, N. Ben Bekhti, L. Flöer, R. Keller, J. Kerp, D. Lenz, B. Winkel, J. Bailin, M. R. Calabretta, L. Dedes, H. A. Ford, B. K. Gibson, U. Haud, S. Janowiecki, P. M. W. Kalberla, F. J. Lockman, N. M. McClure-Griffiths, T. Murphy, H. Nakanishi, D. J. Pisano, L. Staveley-Smith, HI4PI: A full-sky H I survey based on EBHIS and GASS, A&A594 (2016) A116.
- S. Alam, M. Ata, S. Bailey, F. Beutler, D. Bizyaev, J. A. [6] Blazek, A. S. Bolton, J. R. Brownstein, A. Burden, C.-H. Chuang, J. Comparat, A. J. Cuesta, K. S. Dawson, D. J. Eisenstein, S. Escoffier, H. Gil-Marín, J. N. Grieb, N. Hand, S. Ho, K. Kinemuchi, D. Kirkby, F. Kitaura, E. Malanushenko, V. Malanushenko, C. Maraston, C. K. McBride, R. C. Nichol, M. D. Olmstead, D. Oravetz, N. Padmanabhan, N. Palanque-Delabrouille, K. Pan, M. Pellejero-Ibanez, W. J. Percival, P. Petitjean, F. Prada, A. M. Price-Whelan, B. A. Reid, S. A. Rodríguez-Torres, N. A. Roe, A. J. Ross, N. P. Ross, G. Rossi, J. A. Rubiño-Martín, S. Saito, S. Salazar-Albornoz, L. Samushia, A. G. Sánchez, S. Satpathy, D. J. Schlegel, D. P. Schneider, C. G. Scóccola, H.-J. Seo, E. S. Sheldon, A. Simmons, A. Slosar, M. A. Strauss, M. E. C. Swanson, D. Thomas, J. L. Tinker, R. Tojeiro, M. V. Magaña, J. A. Vazquez, L. Verde, D. A. Wake, Y. Wang, D. H. Weinberg, M. White, W. M. Wood-Vasey, C. Yèche, I. Zehavi, Z. Zhai, G.-B. Zhao, The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample, MNRAS470 (2017) 2617-2652.
- [7] M. A. Troxel, N. MacCrann, J. Zuntz, T. F. Eifler, E. Krause, S. Dodelson, D. Gruen, J. Blazek, O. Friedrich, S. Samuroff, J. Prat, L. F. Secco, C. Davis, A. Ferté, J. DeRose, A. Alarcon, A. Amara, E. Baxter, M. R. Becker, G. M. Bernstein, S. L. Bridle, R. Cawthon, C. Chang, A. Choi, J. De Vicente, A. Drlica-Wagner, J. Elvin-Poole, J. Frieman, M. Gatti, W. G. Hartley, K. Honscheid, B. Hoyle, E. M. Huff, D. Huterer, B. Jain, M. Jarvis, T. Kacprzak, D. Kirk, N. Kokron, C. Krawiec, O. Lahav, A. R. Liddle, J. Peacock, M. M. Rau, A. Refregier, R. P. Rollins, E. Rozo, E. S. Rykoff, C. Sánchez, I. Sevilla-Noarbe, E. Sheldon, A. Stebbins, T. N. Varga, P. Vielzeuf, M. Wang, R. H. Wechsler, B. Yanny, T. M. C. Abbott, F. B. Abdalla, S. Allam, J. Annis, K. Bechtol, A. Benoit-Lévy, E. Bertin, D. Brooks, E. Buckley-Geer, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, F. J. Castander, M. Crocce, C. E. Cunha, C. B. D'Andrea, L. N. da Costa, D. L. DePoy, S. Desai, H. T. Diehl, J. P. Dietrich, P. Doel, E. Fernandez, B. Flaugher, P. Fosalba, J. García-Bellido, E. Gaztanaga, D. W. Gerdes, T. Giannantonio, D. A. Goldstein, R. A. Gruendl, J. Gschwend,

G. Gutierrez, D. J. James, T. Jeltema, M. W. G. Johnson, M. D. Johnson, S. Kent, K. Kuehn, S. Kuhlmann, N. Kuropatkin, T. S. Li, M. Lima, H. Lin, M. A. G. Maia, M. March, J. L. Marshall, P. Martini, P. Melchior, F. Menanteau, R. Miquel, J. J. Mohr, E. Neilsen, R. C. Nichol, B. Nord, D. Petravick, A. A. Plazas, A. K. Romer, A. Roodman, M. Sako, E. Sanchez, V. Scarpine, R. Schindler, M. Schubnell, M. Smith, R. C. Smith, M. Soares-Santos, F. Sobreira, E. Suchyta, M. E. C. Swanson, G. Tarle, D. Thomas, D. L. Tucker, V. Vikram, A. R. Walker, J. Weller, Y. Zhang, Dark Energy Survey Year 1 Results: Cosmological Constraints from Cosmic Shear, arxiv:1708.01538.

- [8] H. Hildebrandt, M. Viola, C. Heymans, S. Joudaki, K. Kuijken, C. Blake, T. Erben, B. Joachimi, D. Klaes, L. Miller, C. B. Morrison, R. Nakajima, G. Verdoes Kleijn, A. Amon, A. Choi, G. Covone, J. T. A. de Jong, A. Dvornik, I. Fenech Conti, A. Grado, J. Harnois-Déraps, R. Herbonnet, H. Hoekstra, F. Köhlinger, J. McFarland, A. Mead, J. Merten, N. Napolitano, J. A. Peacock, M. Radovich, P. Schneider, P. Simon, E. A. Valentijn, J. L. van den Busch, E. van Uitert, L. Van Waerbeke, KiDS-450: cosmological parameter constraints from tomographic weak gravitational lensing, MNRAS465 (2017) 1454–1498.
- [9] J. Schmelzle, A. Lucchi, T. Kacprzak, A. Amara, R. Sgier, A. Réfrégier, T. Hofmann, Cosmological model discrimination with Deep Learning, arxiv:1707.05167.
- [10] L. Lucie-Smith, H. V. Peiris, A. Pontzen, M. Lochner, Machine learning cosmological structure formation, MNRAS479 (2018) 3405–3414.
- [11] A. Gupta, J. M. Z. Matilla, D. Hsu, Z. Haiman, Non-Gaussian information from weak lensing data via deep learning, Phys. Rev. D97 (10) (2018) 103515.
- [12] N. Gillet, A. Mesinger, B. Greig, A. Liu, G. Ucci, Deep learning from 21-cm images of the Cosmic Dawn, arxiv:1805.02699.
- [13] S. Hassan, A. Liu, S. Kohn, J. E. Aguirre, P. La Plante, A. Lidz, Reionization Models Classifier using 21cm Map Deep Learning, in: V. Jelić, T. van der Hulst (Eds.), IAU Symposium, Vol. 333 of IAU Symposium, 2018, pp. 47–51.
- [14] M. A. Aragon-Calvo, Classifying the Large Scale Structure of the Universe with Deep Neural Networks, arXiv:1804.00816.
- [15] R. Ciuca, O. F. Hernández, M. Wolman, A Convolutional Neural Network For Cosmic String Detection in CMB Temperature Maps, arxiv:1708.08878.
- [16] S. Ravanbakhsh, J. Oliva, S. Fromenteau, L. C. Price, S. Ho, J. Schneider, B. Póczos, Estimating cosmological parameters from the dark matter distribution, in: Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16, JMLR.org, 2016, pp. 2407–2416.
- [17] Planck Collaboration, Planck 2015 results. I. Overview of products and scientific results, A&A594 (2016) A1.
- [18] B. Abolfathi, D. S. Aguado, G. Aguilar, C. Allende Prieto, A. Almeida, T. T. Ananna, F. Anders, S. F. Anderson, B. H. Andrews, B. Anguiano, et al., The Fourteenth Data Release of the Sloan Digital Sky Survey: First Spectroscopic Data from the Extended Baryon Oscillation Spectroscopic Survey and from the Second Phase of the Apache Point Observatory Galactic Evolution Experiment, ApJS235 (2018) 42.
- [19] Dark Energy Survey Collaboration, The Dark Energy Survey Data Release 1, arxiv: 1801.03181.
- [20] W. Boomsma, J. Frellsen, Spherical convolutions and their application in molecular modelling, in: Advances in Neural Information Processing Systems, 2017, pp. 3436–3446.
- [21] Y.-C. Su, K. Grauman, Learning spherical convolution for fast features from 360 imagery, in: Advances in Neural Information Processing Systems, 2017, pp. 529–539.

- [22] B. Coors, A. P. Condurache, A. Geiger, Spherenet: Learning spherical representations for detection and classification in omnidirectional images, in: European Conference on Computer Vision, 2018.
- [23] J. Fluri, T. Kacprzak, A. Lucchi, A. Refregier, A. Amara, T. Hofmann, Cosmological constraints from noisy convergence maps through deep learning, arxiv:1807.08732.
- [24] T. S. Cohen, M. Geiger, J. Koehler, M. Welling, Spherical cnns, arXiv:1801.10130.
- [25] C. Esteves, C. Allen-Blanchette, A. Makadia, K. Daniilidis, Learning so(3) equivariant representations with spherical cnns, arXiv:1711.06721.
- [26] M. Defferrard, X. Bresson, P. Vandergheynst, Convolutional neural networks on graphs with fast localized spectral filtering, in: Advances in Neural Information Processing Systems, 2016, pp. 3844–3852.
- [27] K. M. Gorski, E. Hivon, A. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, M. Bartelmann, Healpix: a framework for high-resolution discretization and fast analysis of data distributed on the sphere, The Astrophysical Journal 622 (2) (2005) 759.
- [28] C. Chang, A. Pujol, B. Mawdsley, D. Bacon, J. Elvin-Poole, P. Melchior, A. Kovács, B. Jain, B. Leistedt, T. Giannantonio, A. Alarcon, E. Baxter, K. Bechtol, M. R. Becker, A. Benoit-Lévy, G. M. Bernstein, C. Bonnett, M. T. Busha, A. C. Rosell, F. J. Castander, R. Cawthon, L. N. da Costa, C. Davis, J. De Vicente, J. DeRose, A. Drlica-Wagner, P. Fosalba, M. Gatti, E. Gaztanaga, D. Gruen, J. Gschwend, W. G. Hartley, B. Hoyle, E. M. Huff, M. Jarvis, N. Jeffrey, T. Kacprzak, H. Lin, N. Mac-Crann, M. A. G. Maia, R. L. C. Ogando, J. Prat, M. M. Rau, R. P. Rollins, A. Roodman, E. Rozo, E. S. Rvkoff, S. Samuroff, C. Sánchez, I. Sevilla-Noarbe, E. Sheldon, M. A. Troxel, T. N. Varga, P. Vielzeuf, V. Vikram, R. H. Wechsler, J. Zuntz, T. M. C. Abbott, F. B. Abdalla, S. Allam, J. Annis, E. Bertin, D. Brooks, E. Buckley-Geer, D. L. Burke, M. C. Kind, J. Carretero, M. Crocce, C. E. Cunha, C. B. D'Andrea, S. Desai, H. T. Diehl, J. P. Dietrich, P. Doel, J. Estrada, A. F. Neto, E. Fernandez, B. Flaugher, J. Frieman, J. García-Bellido, R. A. Gruendl, G. Gutierrez, K. Honscheid, D. J. James, T. Jeltema, M. W. G. Johnson, M. D. Johnson, S. Kent, D. Kirk, E. Krause, K. Kuehn, S. Kuhlmann, O. Lahav, T. S. Li, M. Lima, M. March, P. Martini, F. Menanteau, R. Miquel, J. J. Mohr, E. Neilsen, R. C. Nichol, D. Petravick, A. A. Plazas, A. K. Romer, M. Sako, E. Sanchez, V. Scarpine, M. Schubnell, M. Smith, R. C. Smith, M. Soares-Santos, F. Sobreira, E. Suchyta, G. Tarle, D. Thomas, D. L. Tucker, A. R. Walker, W. Wester, Y. Zhang, Dark Energy Survey Year 1 results: curved-sky weak lensing mass map, MNRAS475 (2018) 3165-3190.
- [29] M. Bartelmann, P. Schneider, Weak gravitational lensing, Phys. Rep.340 (2001) 291-472. arXiv:astro-ph/9912508, doi:10.1016/S0370-1573(00)00082-X.
- [30] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, X. Zheng, TensorFlow: Large-scale machine learning on heterogeneous systems, software available from tensorflow.org (2015).

URL https://www.tensorflow.org/

[31] M. Defferrard, L. Martin, R. Pena, N. Perraudin, Pygsp: Graph signal processing in python. doi:10.5281/zenodo.1003157. URL https://github.com/epfl-lts2/pygsp/

- [32] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, Proceedings of the IEEE 86 (11) (1998) 2278–2324.
- [33] S. Ioffe, C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: International Conference on Machine Learning, 2015, pp. 448–456.
- [34] M. Belkin, P. Niyogi, Convergence of laplacian eigenmaps, in: Advances in Neural Information Processing Systems, 2007, pp. 129–136.
- [35] P. Frossard, R. Khasanova, Graph-based classification of omnidirectional images, in: 2017 IEEE International Conference on Computer Vision Workshops (ICCVW), 2017, pp. 860–869.
- [36] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, P. Vandergheynst, The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains, IEEE Signal Processing Magazine 30 (3) (2013) 83–98.
- [37] N. Perraudin, P. Vandergheynst, Stationary signal processing on graphs., IEEE Trans. Signal Processing 65 (13) (2017) 3462– 3477.
- [38] L. Le Magoarou, R. Gribonval, N. Tremblay, approximate fast graph fourier transforms via multi-layer sparse approximations, IEEE transactions on Signal and Information Processing over Networks 4 (2) (2018) 407–420.
- [39] M. Reinecke, D. S. Seljebotn, Libsharp–spherical harmonic transforms revisited, Astronomy & Astrophysics 554 (2013) A112.
- [40] K. T. Inoue, P. Cabella, E. Komatsu, Harmonic inpainting of the cosmic microwave background sky: Formulation and error estimate, Phys. Rev. D77 (12) (2008) 123539.
- [41] S. Amsel, J. Berger, R. H. Brandenberger, Detecting cosmic strings in the CMB with the Canny algorithm, J. Cosmology Astropart. Phys.4 (2008) 015.
- [42] J. Long, E. Shelhamer, T. Darrell, Fully convolutional networks for semantic segmentation, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, pp. 3431–3440.
- [43] M. Lin, Q. Chen, S. Yan, Network in network, arXiv:1312.4400.
- [44] J. T. Springenberg, A. Dosovitskiy, T. Brox, M. Riedmiller, Striving for simplicity: The all convolutional net, arXiv preprint arXiv:1412.6806.
- [45] D. K. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, R. P. Adams, Convolutional networks on graphs for learning molecular fingerprints, in: Advances in neural information processing systems, 2015, pp. 2224–2232.
- [46] Y. Li, D. Tarlow, M. Brockschmidt, R. Zemel, Gated graph sequence neural networks, in: International Conference on Learning Representation, 2016.
- [47] C. Ronchi, R. Iacono, P. S. Paolucci, The "cubed sphere": a new method for the solution of partial differential equations in spherical geometry, Journal of Computational Physics 124 (1) (1996) 93–114.
- [48] J. Xiao, K. A. Ehinger, A. Oliva, A. Torralba, Recognizing scene viewpoint using panoramic place representation, in: Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, IEEE, 2012, pp. 2695–2702.
- [49] Y. Zhang, S. Song, P. Tan, J. Xiao, Panocontext: A whole-room 3d context model for panoramic scene understanding, in: European Conference on Computer Vision, Springer, 2014, pp. 668– 686.
- [50] J. Masci, D. Boscaini, M. Bronstein, P. Vandergheynst, Geodesic convolutional neural networks on riemannian manifolds, in: Proceedings of the IEEE international conference on computer vision workshops, 2015, pp. 37–45.

- [51] D. Boscaini, J. Masci, E. Rodolà, M. Bronstein, Learning shape correspondence with anisotropic convolutional neural networks, in: Advances in Neural Information Processing Systems, 2016, pp. 3189–3197.
- [52] F. Monti, D. Boscaini, J. Masci, E. Rodola, J. Svoboda, M. M. Bronstein, Geometric deep learning on graphs and manifolds using mixture model cnns, in: Proc. CVPR, Vol. 1, 2017, p. 3.
- [53] M. J. Mohlenkamp, A fast transform for spherical harmonics, Journal of Fourier analysis and applications 5 (2-3) (1999) 159– 184.
- [54] V. Rokhlin, M. Tygert, Fast algorithms for spherical harmonic expansions, SIAM Journal on Scientific Computing 27 (6) (2006) 1903–1928.
- [55] R. Kondor, Z. Lin, S. Trivedi, Clebsch-gordan nets: a fully fourier space spherical convolutional neural network, arXiv:1806.09231.
- [56] R. Kondor, S. Trivedi, On the generalization of equivariance and convolution in neural networks to the action of compact groups, arXiv:1802.03690.
- [57] S. Mallat, Group invariant scattering, Communications on Pure and Applied Mathematics 65 (10) (2012) 1331–1398.
- [58] N. Kalchbrenner, E. Grefenstette, P. Blunsom, A convolutional neural network for modelling sentences, in: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Vol. 1, 2014, pp. 655–665.
- [59] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, P. Vandergheynst, Geometric deep learning: going beyond euclidean data, IEEE Signal Processing Magazine 34 (4) (2017) 18–42.
- [60] W. L. Hamilton, R. Ying, J. Leskovec, Representation learning on graphs: Methods and applications, IEEE Data Engineering Bulletin, arXiv:1709.05584.
- [61] S. I. Ktena, S. Parisot, E. Ferrante, M. Rajchl, M. Lee, B. Glocker, D. Rueckert, Metric learning with spectral graph convolutions on brain connectivity networks, NeuroImage 169 (2018) 431–442.
- [62] S. Parisot, S. I. Ktena, E. Ferrante, M. Lee, R. G. Moreno, B. Glocker, D. Rueckert, Spectral graph convolutions for population-based disease prediction, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2017, pp. 177–185.
- [63] P. Hop, B. Allgood, J. Yu, Geometric deep learning autonomously learns chemical features that outperform those engineered by domain experts, Molecular pharmaceutics.
- [64] X. Qi, R. Liao, J. Jia, S. Fidler, R. Urtasun, 3d graph neural networks for rgbd semantic segmentation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 5199–5208.
- [65] P. Baqué, E. Remelli, F. Fleuret, P. Fua, Geodesic convolutional shape optimization, in: International Conference on Machine Learning, 2018.
- [66] Y. Seo, M. Defferrard, P. Vandergheynst, X. Bresson, Structured sequence modeling with graph convolutional recurrent networks, arXiv:1612.07659.
- [67] Y. Li, R. Yu, C. Shahabi, Y. Liu, Diffusion convolutional recurrent neural network: Data-driven traffic forecasting.
- [68] F. Monti, M. Bronstein, X. Bresson, Geometric matrix completion with recurrent multi-graph neural networks, in: Advances in Neural Information Processing Systems, 2017, pp. 3697–3707.
- [69] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, et al., Relational inductive biases, deep learning, and graph networks, arXiv:1806.01261.
- [70] M. Bartelmann, Gravitational lensing, Classical and Quantum Gravity 27 (23) (2010) 233001.
- [71] R. Sgier, A. Réfrégier, A. Amara, A. Nicola, Fast Generation of

Covariance Matrices for Weak Lensing, arxiv:1801.05745.

- [72] C. Howlett, M. Manera, W. J. Percival, L-PICOLA: A parallel code for fast dark matter simulation, Astronomy and Computing 12 (2015) 109–126.
- [73] K. Patton, J. Blazek, K. Honscheid, E. Huff, P. Melchior, A. J. Ross, E. Suchyta, Cosmological constraints from the convergence 1-point probability distribution, MNRAS472 (2017) 439– 446.
- [74] N. Krachmalnicoff, M. Tomasi, Convolutional neural networks on the healpix sphere: a pixel-based algorithm and its application to cmb data analysis, arXiv preprint arXiv:1902.04083.
- [75] S. He, S. Ravanbakhsh, S. Ho, Analysis of cosmic microwave background with deep learning.
- [76] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv:1412.6980.
- [77] X. Glorot, Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in: Proceedings of the thirteenth international conference on artificial intelligence and statistics, 2010, pp. 249–256.
- [78] D. Ribli, B. Ármin Pataki, I. Csabai, Learning from deep learning: better cosmological parameter inference from weak lensing maps, arxiv:1806.05995.
- [79] N. Perraudin, B. Ricaud, D. I. Shuman, P. Vandergheynst, Global and local uncertainty principles for signals on graphs, APSIPA Transactions on Signal and Information Processing 7.