



HAL
open science

OSLO: On-the-Sphere Learning for Omnidirectional images and its application to 360-degree image compression

Navid Mahmoudian Bidgoli, Roberto Azevedo, Thomas Maugey, Aline Roumy, Pascal Frossard

► To cite this version:

Navid Mahmoudian Bidgoli, Roberto Azevedo, Thomas Maugey, Aline Roumy, Pascal Frossard. OSLO: On-the-Sphere Learning for Omnidirectional images and its application to 360-degree image compression. IEEE Transactions on Image Processing, 2022, pp.1-16. hal-03833562

HAL Id: hal-03833562

<https://inria.hal.science/hal-03833562>

Submitted on 28 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

OSLO: On-the-Sphere Learning for Omnidirectional images and its application to 360-degree image compression

Navid Mahmoudian Bidgoli, Roberto G. de A. Azevedo, Thomas Maugey, Aline Roumy, Pascal Frossard

Abstract—State-of-the-art 2D image compression schemes rely on the power of convolutional neural networks (CNNs). Although CNNs offer promising perspectives for 2D image compression, extending such models to omnidirectional images is not straightforward. First, omnidirectional images have specific spatial and statistical properties that can not be fully captured by current CNN models. Second, basic mathematical operations composing a CNN architecture, *e.g.*, translation and sampling, are not well-defined on the sphere. In this paper, we study the learning of representation models for omnidirectional images and propose to use the properties of HEALPix uniform sampling of the sphere to redefine the mathematical tools used in deep learning models for omnidirectional images. In particular, we: i) propose the definition of a new convolution operation on the sphere that keeps the high expressiveness and the low complexity of a classical 2D convolution; ii) adapt standard CNN techniques such as stride, iterative aggregation, and pixel shuffling to the spherical domain; and then iii) apply our new framework to the task of omnidirectional image compression. Our experiments show that our proposed on-the-sphere solution leads to a better compression gain that can save 13.7% of the bit rate compared to similar learned models applied to equirectangular images. Also, compared to learning models based on graph convolutional networks, our solution supports more expressive filters that can preserve high frequencies and provide a better perceptual quality of the compressed images. Such results demonstrate the efficiency of the proposed framework, which opens new research venues for other omnidirectional vision tasks to be effectively implemented on the sphere manifold.

Index Terms—360-degree images, spherical convolution, end-to-end compression, anisotropic filter

I. INTRODUCTION

360-degree, or *omnidirectional*, visual contents are high-resolution *spherical* images/videos that capture a 3D scene with a full field of view. By allowing users to adjust their viewing orientation freely, omnidirectional contents support an immersive and interactive experience. Due to their unique features, omnidirectional images have become increasingly popular in different applications such as Virtual Reality (VR) [1],

This work was supported by the Cominlabs excellence laboratory with funding from the French National Research Agency (ANR-10-LABX-07-01), InterCom and MOVE projects.

Navid Mahmoudian Bidgoli, Thomas Maugey, Aline Roumy are with Inria, Univ Rennes, CNRS, IRISA, France (e-mails: {navid.mahmoudian-bidgoli, thomas.maugey, aline.roumy}@inria.fr).

Roberto G. de A. Azevedo is with ETH Zurich.

Pascal Frossard is with the Signal Processing Laboratory (LTS4), École Polytechnique Fédérale de Lausanne (EPFL).

This work was developed while Roberto G. de A. Azevedo was with the Signal Processing Laboratory (LTS4), École Polytechnique Fédérale de Lausanne (EPFL).

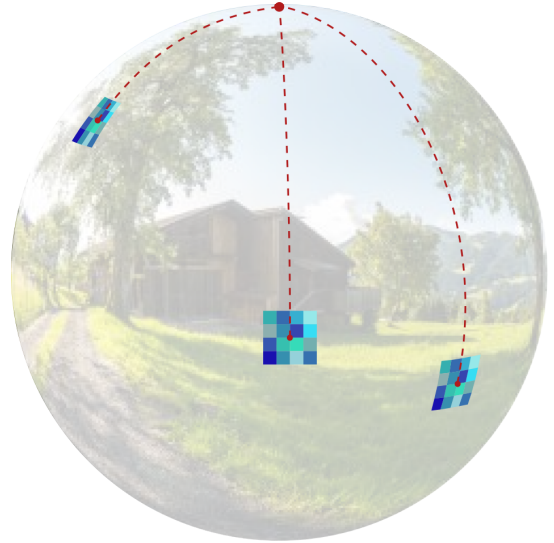


Fig. 1. In a *coherent* filtering, the kernel is translated by keeping orientation with respect to the north pole, interpixel correlation and surface area.

[2], augmented reality (AR), and autonomous driving [3], which at the same time outlines the crucial need for effective processing tools. Many recent imaging techniques are based on Convolutional Neural Networks (CNNs), which have achieved great success for planar images in various computer vision and image processing applications such as image classification or compression. However, their extension to omnidirectional images raises several issues.

A spherical image is a function defined over a sphere and the development of CNN processing tools over this 2D manifold is not straightforward. More precisely, the difficulty consists in constructing a convolution and its resulting CNN that satisfy the three following key properties: (i) *rotation equivariance* (meaning that filtering and rotation commute), (ii) *expressive filter* (the filter can have any size and any impulse response), and (iii) *computational efficiency* (complexity grows linearly with the number of pixels N_{pix} , and computation is local). Indeed, existing solutions maintain two out of these three characteristics. For instance, in a first set of methods, the sphere is mapped on one or multiple planes, and 2D-processing is applied on the plane(s) (equirectangular (ERP) [4], the cube map [5], [6], and the rhombic dodecahedron [7] projections). In these approaches, the convolution is not equivariant to rotation as the projection induces distortion, that depends on

the position on the sphere. A second approach consists in performing CNN on a graph [8]. Here, the filters are restricted since only one coefficient per neighborhood can be learned, and not one per neighbor. In the last category of methods (multiple graph [9]), and spherical harmonics based methods [10]–[13]), the two first properties are maintained but the complexity is at least $N_{pix}^{3/2}$, see also Table I.

In this paper, by effectively integrating the geometry of the problem, we propose a novel learning approach for spherical images, called OSLO, that satisfies all above mentioned properties. First, to achieve rotation equivariance (i), we perform *consistent* filtering. This means that, by translating the filter over the sphere surface, the filter (a) keeps the same orientation towards the north pole, (b) processes pixels with the same interpixel correlation, and (c) processes the same area of the sphere. These properties, illustrated in Fig. 1, result from the exploitation of a sampling that is both uniform and oriented with respect to the north pole. It also takes advantage of the fact that omnidirectional images are usually level or registered¹ [14], [15]. Second, the image being registered, and the sampling being oriented, one coefficient per neighbor can be learned. Moreover, we propose an iterative construction of this convolution with a dedicated aggregation such that the kernel can be learned for every pixel neighborhood size. This leads to a filter with high expressiveness (ii). Finally, since filtering consists in translating the filter over the sphere, complexity grows linearly with the number of pixels, and computation is local (iii). We also extend some of the important operators from classical 2D CNN to the sphere, such as stride, pooling, pixel shuffling, so that each elementary module of a 2D CNN can be defined on the sphere.

We finally illustrate the benefits of the proposed spherical framework, by extending two well-known, learning-based, 2D image compression architectures [16], [17] onto the sphere. More precisely, we transform these two complete architectures directly on the sphere using all the elementary modules proposed in our OSLO solution. We then compare the compression performance of such architectures with the initial architectures from [16], [17], when applied to ERP images. The results show that our approach leads to significant compression gain with -18.3% and -13.7% of bit rate saving for the two architectures considered. Moreover, in the reconstructed images, the high frequency details are preserved at high bitrates, showing the benefits of the filter expressiveness. All this demonstrates the efficiency of the proposed convolution and CNN modules, and their potential impact for the extension of any architecture (performing any processing task) on the sphere.

The rest of the paper is organized as follows. Section II explains in detail the challenges in designing a spherical CNN, and notably the difficulty to construct filters that satisfy all necessary properties for a CNN to be both efficient and tractable. We then present in Section III our proposed ‘On the Sphere Learning for Omnidirectional’ (OSLO) framework and describe an efficient implementation. In Section IV, we present

an application of this new framework to omnidirectional image compression, and validate the benefit of our solution via extensive experiments.

II. REPRESENTATION LEARNING ON THE SPHERE

A. Problem settings

The unit 2-sphere \mathcal{S}^2 is a 2D manifold defined as the set of 3D points in Euclidean space \mathbb{R}^3 with a norm of 1. A spherical image f is a continuous function $f : \mathcal{S}^2 \rightarrow \mathbb{R}^3$ representing red, green, and blue color channels². For each point of the sphere $x \in \mathcal{S}^2$, the image value $f(x)$ corresponds to the color value of one light field’s sample converging to the center of the sphere O through x . In other words, it depicts the color of one visible point of the 3D scene lying on the line (Ox) . It is expected that $f(x)$ is correlated with $f(x')$ when $\|x - x'\|_2$ is small because they are likely to belong to the same object in the scene. In practice, the spherical image is captured with a specific hardware that depends on the application. In general, it comes from the combination of several captures that can be done with perspective, fish eye, or catadioptric cameras [18]. Then, in most applications, these images are not rendered in their spherical shape. For example, in virtual reality applications, they are partly projected on a Head Mounted Display (HMD) that is controlled by a user. In such a context, it is preferred that the images are level or registered. In other words, the north pole of these images should point towards the upright direction such that buildings, humans, doors, etc. have a natural upright position. A non-level 360° image completely breaks the sense of realism for virtual reality users and leads to an unpleasant immersive experience and even severe user sickness [19]. In case the 360° images are not upright, they can be leveled using [14], [15]. Another consequence of the rendering modes is that the images should be of a very high resolution (up to 10K) so that the HMD visualization remains of a sufficiently good quality.

Like 2D images, spherical images are employed in various contexts for which efficient processing tools are required: compression, object recognition, classification, super-resolution, etc. Such tasks may benefit from the recent advances in machine learning, and therefore, there is nowadays a crucial need for defining proper representations for spherical data. However, due to their special domain of definition, huge size, and specific geometry distortion, spherical images are not straightforwardly processed by traditional CNNs. Specific spherical CNN architectures have to be constructed while respecting a certain number of requirements summarized in Table I. First, the convolutional filters have to be *rotation equivariant*, i.e., filtering and rotation commute. This enables the filters involved in the CNN architecture to be *applied consistently all over the sphere* (see Fig. 1). Meaning that the learned filters should keep the same orientation with respect to the north pole and surface area on the sphere, while modeling the same inter-pixel correlation. Second, to guarantee a good

¹The north pole of these images usually point towards the upright direction such that buildings, humans, doors, etc. have a natural upright position.

²Here, we focus mainly on color (RGB) images defined on the sphere since that is the most common use case, but all of our definitions are also valid to any multi-channel images, i.e., when $f : \mathcal{S}^2 \rightarrow \mathbb{R}^{ch}$, with ch being the number of channels.

TABLE I
CHARACTERISTICS OF SPHERICAL CNNs

Domain of definition:		Representation method				
		2D learning mapped 2D grid	spectral learning all	graph-based representation learning all	multi-graph representation learning all	proposed OSLO uniform sampling
(i) filter consistency	(a) same orientation	Yes	Yes	Yes	Yes	Yes
	(b) same interpixel correlation	No	Yes	Yes	Yes	Yes
	(c) same surface on the sphere	No	Yes	Yes	Yes	Yes
(ii) filter expressiveness	(a) anisotropic	Yes	Yes	No	Yes	Yes
	(b) large domain	Yes	Yes	Yes	Yes	Yes
(iii) computational efficiency	(a) complexity	$\mathcal{O}(N_{pix})$	at least $\mathcal{O}(N_{pix}^{3/2})$	$\mathcal{O}(N_{pix})$	$\mathcal{O}(N_{pix}^{3/2})$	$\mathcal{O}(N_{pix})$
	(b) local computation	Yes	No	Yes	Yes	Yes

processing performance, the filters used at the different layers should be *as expressive as those of the common 2D CNN architectures*. Meaning that they should be able to perform convolution on a *large neighborhood*, and have *anisotropic* responses. Finally, the filters computation should be *efficient*. In other words, the computation of the convolution should be linear with the number of pixels (as in a standard 2D CNN), and a local computation should be feasible (not involving the whole image at each filter).

In the next section, we will discuss the challenges of achieving these goals for spherical data.

B. Challenges

On-the-sphere operations such as convolution and translation admit several definitions. Thus, filtering a spherical image by a kernel can be performed in several manners, each of them raising specific issues.

A first definition represents the sphere with $SO(3)$, *i.e.*, the group of rotations in \mathbb{R}^3 , and computes a filtered output of a spherical image for every rotation in $SO(3)$, *i.e.*, moving the kernel on the sphere with 3 degrees of freedom. While exhaustive, this filter representation has the drawback of an increased dimension (from 2 to 3) during the filtering process.

Another definition restricts the displacement of the kernel to 2 degrees of freedom. In such a case, the spherical convolution is defined as the inner product between the image f and a localized kernel h translated at different positions $x \in \mathcal{S}^2$:

$$(f \star h)(x) = \langle f, \mathcal{T}_x h \rangle, \quad (1)$$

where \mathcal{T}_x is a translation operator. We can see that the filtering of a spherical image by a kernel remains a spherical image defined on \mathcal{S}^2 . However, there are an infinite number of ways of defining \mathcal{T}_x , depending, for example, on the direction of the polar axis. This ambiguity may be problematic in some applications, in particular when the spherical data does not have any a priori fixed orientation. However, this definition can be advantageous when dealing with registered (upright) omnidirectional images since it ensures that the kernel is always oriented similarly and consistently with the south/north axis at every position (as illustrated in Fig. 1).

Even with only 2 degrees of freedom, the definition of the translation \mathcal{T}_x is challenging. It can even become an open problem when dealing with discrete pixel positions. On top of the necessary interpolation artifacts, the pixelization of the sphere impacts the structure and properties of the discrete domain of definition. On the one hand, as 2D translations

are well defined on a Cartesian grid, one can map the sphere onto one or multiple 2D planes. This unfortunately results in heterogeneous (and thus inconsistent) pixel distributions on the sphere, which have to be compensated at the filtering stage. This is not straightforward, and may lead to complex operations or necessary simplifications of the filters, penalizing at the same time their expressiveness. On the other hand, when the sphere is pseudo-uniformly sampled, the pixel's set becomes non-Cartesian, and therefore \mathcal{T}_x is undefined. In a nutshell, we can see that the main challenge raised by spherical learning is to find the adequacy between the sphere sampling, the signal processing tools that it enables, and the compensation of their limitations.

C. Current Solutions

We now review in detail the solutions proposed in the literature to learn representations for omnidirectional images. Their properties in terms of processing accuracy and efficiency are summarized in Table I and discussed in detail in the following.

1) *2D learning*: Mapping spherical images on 2D planar grid enables to directly apply the well-known 2D CNN tools on such a planar representation. For instance, authors in [20] use the popular equirectangular projection (ERP) that samples the sphere with fixed longitude and latitude steps. ERP sampling, however, results in highly heterogeneous pixel distribution on the sphere. As a consequence, they cannot guarantee that the kernel keeps the same surface nor describes a constant inter-pixel correlation over the sphere (see Table I). To circumvent this issue, the authors adapt the size of the kernel at different elevations by learning different variable-size kernels for each row of the ERP. However, this dramatically increases the model size and does not fully utilize CNN's weight-sharing property.

Instead of ERP, the works in [21], [22] propose to use cube map projection (CMP). The image is mapped to the six faces of a cube, and each face is processed with conventional 2D CNN, considering the face as the image plane of a perspective camera. While the CMP pixel distribution on the sphere is more uniform than ERP's one, the discontinuities at the borders of the cube faces require to be handled carefully that needs post-processing. Also, defining a kernel that keeps its orientation constant with respect to the north pole is impossible since either the north pole lies in the middle of the top face, or worst at one of the corners of the cube.

Other mappings have been investigated in [23]–[25]. By increasing the number of faces of a polyhedron on which the

sphere is projected, the pixel’s distribution on the sphere is closer to a uniform pixelization. However, the border discontinuity problem is increased, and the orientation of the filter with respect to the north is not always guaranteed. Despite their convenient link with existing 2D tools, mapping-based solutions inherently lead to too many drawbacks, (*e.g.*, non-uniform sampling, discontinuities and local orientation loss) to be convincing.

2) *Spectral learning*: Similar to classical 2D convolution, one can define convolution on the sphere as a product in the Fourier domain. Such an idea is interesting since the spherical harmonics are a good Fourier basis to represent spherical signals. More precisely, as summarized in Table I, they enable defining consistent and expressive filtering all over the sphere. However, these methods are complex, and the simplifications proposed to reduce the computation time often imply a loss in expressiveness. For instance, Esteves et al. [10] propose a spherical convolution with reasonable complexity $\mathcal{O}(N_{pix} \log^2 N_{pix})$, but it only generates isotropic filters. Instead, Cohen et al. [11] propose using spherical correlation defined in $SO(3)$. Such an approach enables the convolution parameters to take into account filter direction. Nevertheless, the complexity of this approach ($\mathcal{O}(N_{pix}^2)$) rapidly becomes a limitation when handling high-resolution spherical images. Recently, new spectral spherical CNNs with spin-weighted functions have been proposed [12]. The proposed convolution is not restricted to isotropic filters and remains 2D and defined on the sphere. However, the $\mathcal{O}(N_{pix}^{3/2})$ complexity still limits its use to high-resolution images. Similarly, Roddy et al. [13] have proposed a sifting convolution defined as the inner product between the image f and a localized kernel h translated at different positions in the spherical harmonics domain (with 2 degrees of freedom). This approach has strong similarities with dictionary-based ones [26]. Indeed, atoms are defined locally and translated on the sphere to form a dictionary used for image representation. While such approaches match the consistency and expressiveness expectations summarized in Table I, their supra-linear complexity remains a major limitation for high-resolution spherical images. Moreover, local estimation of the filtering process is not possible since spherical harmonics have to be estimated on the whole sphere.

3) *Graph-based representation learning*: Recently, graph signal processing tools have demonstrated their efficiency when dealing with non-euclidean topology. By modeling the sphere as a graph and embedding the inter-pixel distance in the edge weights, a filter can be applied consistently all over the spherical data, independently of the chosen sampling. For example, Khasanova and Frossard [27] propose a graph-based convolution for ERP in which the edge weights between neighboring nodes (pixels) are equal to the inverse of their distances. Alternatively, other works use a quasi-uniform sampling of the sphere and define graph-based CNNs directly on the sphere [8], [28]. Perraudin et al. [8] use Hierarchical Equal Area isoLatitude Pixelation (HEALPix) [29], and Yang et al. [28] use Geodesic ICOSahedral Pixelation (GICOPix). To keep a linear complexity filtering, the convolution is approximated using a Chebyshev polynomial formulation, thus avoiding the graph spectral transform computation. Despite

their great adaptivity to the topology and a good filtering consistency over the sphere, the main drawback of such methods is the low expressiveness of the learned filters (see Table I). Due to the lack of direction in the graph, only isotropic filters can be learned, which is a major issue in imaging applications.

4) *Multi-graph representation learning*: To overcome the aforementioned limitations of graph-based approaches, Khasanova and Frossard [9] use multiple directed graphs, in which each graph represents a specific direction and location of neighboring pixels relative to the central pixels. Such an approach allows adaptation to any topology while preserving orientation information and filter expressiveness. However, their graph construction is complex and requires a large storage space. Indeed, to create an $n \times n$ convolution kernel, n^2 graphs are constructed *at each pixel position*. More precisely, the graph k is constructed by projecting the k^{th} area of an $n \times n$ tangent rectangle on the sphere surface and linking all pixels in the projected area to the central pixel. This implies a high complexity to estimate all the n^2 graphs at each kernel position and requires a large storage size (for each layer of the neural network).

D. Discussion

As we can see in Table I, none of the current learning approaches is able to satisfy all the initial requirements and have at the same time a linear complexity, *i.e.*, $\mathcal{O}(N_{pix})$. While it is not possible to guarantee consistent filtering with the mapping-based solution, using graph-based approaches to adjust to the topology reduces the filter expressiveness. Spectral and multi-graph approaches support more expressive filters that can be applied consistently. However, they come with a significant increase in complexity, resulting in issues to handle high-resolution 360° images. Alternatives to the aforementioned categories of methods have also been explored by defining tools directly on the uniformly sampled sphere (*e.g.*, [30], [31]). Their approaches however remain pretty complex and too specific to the target applications.

III. ON-THE-SPHERE LEARNING FOR OMNIDIRECTIONAL IMAGE

A. OSLO Overview

As discussed earlier (Table I), existing spherical learning methods fail to construct expressive and consistent filters while being computationally tractable. Thus, we propose a new framework, On-the-Sphere Learning for Omnidirectional Images (OSLO), to build effective representations of spherical data that attain the same filter expressiveness, consistency, and complexity of 2D CNNs. More precisely, we rely on the HEALpix sampling’s properties to redefine all the elementary modules composing a CNN architecture *on-the-sphere*. In addition to the convolution and pooling operations that appear in most deep network models, we define 2D CNN’s operators such as stride, skip connection, and patching *on-the-sphere*. In this section, we present the design of the different operators, which can eventually be assembled to create different architectures trained on diverse tasks. One example of using OSLO




Fig. 2. Desired sampling's properties : (i) It should be *regular* such that each pixel has same number of neighbors, and we should be able to discriminate between neighboring pixels by assigning different directions to each neighbor (as shown in (a)); (ii) The sampling should be *rigid* (as shown in (b)) such that the distance between a specific neighbor (1 to 8) and the center pixel should be almost constant all over the sphere (e.g., $d_1 = d'_1$). Additionally, its relative orientation with respect to the center pixel should be constant over the sphere, (e.g., $\theta_1 = \theta'_1$). This orientation is computed with respect to the north pole.

for neural compression of omnidirectional images is presented in Section IV.

B. Sampling

To be able to have fast convolution and avoid using spectral domain, we propose implementing the spherical convolution in the pixel domain, which corresponds to translating a kernel as defined in (1). One of the key elements in defining an effective representation learning solution is the choice of a proper sampling method that enables the design of expressive filters with a fast implementation. The filter output is computed with the same weights, irrespective of the position of the filter on the sphere. As a consequence, the complexity scales linearly with the number of pixels. In order to construct this computationally efficient convolution, the chosen sampling must fulfill some properties summarized in Fig. 2. First, to have an expressive filter, the sampling grid must ensure *regularity* of the neighborhood such that the same number of neighbors always surrounds each pixel, and those can be identified (based on their relative position with respect to the central point). Second, the sampling grid should ensure that the relative distance and orientation between a point and its neighbor is constant over the sphere; thus, the *rigidity* of the neighborhood and, therefore, the consistency of the filter shape is guaranteed. Several sampling methods exist, most of them relying on a sphere tessellation. Often, only one of the regularity or rigidity constraints is achieved. For example, being pseudo-uniform as GICOPIX [28] is not sufficient since the pixels have additionally to be organized such that the translation of the kernel remains coherent over the sphere. In other words, a specific neighbor that has a specific position relative to the center pixel must be easily identifiable throughout the sphere.

We now show that HEALPix [29] is a reliable sampling process for our convolution that fulfills the above requirements for effective and expressive feature representation. In HEALPix, the tessellation process begins by partitioning the spherical surface into 12 equal-area regions (base resolution). To have




Fig. 3. The distribution of neighboring pixels around the central pixel. Left column: ERP. Right column: HEALPix sampling. Top row: Pixel positions on the sphere when ERP and HEALPix have almost the same number of pixels with 3072 and 39×78 pixels. Bottom row: Pixel distributions on the tangent plane when ERP and HEALPix have almost the same number of pixels with 12582912 and 2508×5016 pixels. Marker \times shows the central pixel, black circles represent the mean position of each neighbor. Each color corresponds to a neighbor pixel in a specific direction (SW, W, NW, N, NE, E, SE, and S) and the corresponding area's size is set based on the standard deviation of pixel distribution in the plane tangent to the sphere at the central pixel.

finer pixelization, each region is recursively divided into 2×2 equal-area sub-pixels until the desired resolution is reached. By construction, each pixel in HEALPix has eight adjacent neighbors (structured in a diamond pattern), except for 24 pixels that have seven neighbors (for any resolution higher than the base resolution). The effect of these exceptional points is negligible for high-resolution 360° images (a HEALPix resolution close to conventional 4K resolution has less than 0.0002% pixels with 7 neighbors). Finally, a fundamental property of HEALPix, often overlooked in the literature, is that the orientation of neighboring points with the central point is almost fixed all over the sphere. Each neighbor can be identified by its relative direction to the central pixel: SW, W, NW, N, NE, E, SE, and S. Therefore, HEALPix fulfills the *regularity* property that is essential to define an expressive filter.

To demonstrate the *rigidity* property of HEALPix, we calculate the relative distance and angle (d and θ of Fig. 2) for each position on the sphere. More precisely, for a given central pixel, we define a tangent plane in which the y axis points towards the north pole. Then, each neighboring pixel is projected onto the tangent plane. The relative distance and angle to the central pixel are then computed. By doing that for every pixel position on the sphere, we are able to show the neighbor's distribution. The distribution of neighboring points

TABLE II
RELATIVE STANDARD DEVIATION (IN %) OF DISTANCES WITH RESPECT TO THE CENTRAL PIXEL FOR EACH NEIGHBOR ALL OVER THE SPHERE. BOTH SAMPLINGS HAVE THE SAME NUMBER OF PIXELS (HEALPIX SAMPLING WITH 12582912 PIXELS AND ERP OF SIZE 2508×5016)

	SW	W	NW	N	NE	E	SE	S	Mean
ERP	12.01	48.34	12.01	0	12.01	48.34	12.01	0	18.09
HEALPix	11.21	10.28	11.21	13.90	11.21	10.28	11.21	13.9	11.65

for both HEALPix and ERP are shown in Fig. 3.

At first sight, it is clear that HEALPix pseudo-uniform sampling provides a *rigid* structure, allowing each neighbor to keep its relative distance and direction with respect to the center point, all over the sphere. When compared to ERP, for the majority of neighbors (6 out of 8 neighbors), HEALPix has much lower discrepancies. North and south neighbors present greater discrepancies than ERP. However, since the correlation in 360° images is mostly horizontal [32], the north and south neighbors have less impact than other pixels. Furthermore, the discrepancy in HEALPix does not affect the distance to the central node in comparison to the ERP. The relative standard deviation (ratio of standard deviation to the mean) of the distances between each neighbor and central pixel is given in Table II. This is important for the design of a convolution filter because the signal correlation depends on the distance of central pixels to their neighbors [8], [27]. In other words, the impact of kernel weights on neighboring pixels is affected by their distance to the central node, and the filter response is effective if this distance is uniform throughout the sphere.

Since HEALPix fulfills the regularity and rigidity properties, we propose to use that sampling method to develop computationally efficient operators on the sphere. In particular, we can now define a new linear complexity convolution operator, and show that it leads to expressive and consistent filters (properties (i) and (ii) in Table I).

C. Spherical convolution

Since we can distinguish the orientation of neighbors relative to the center and the neighborhood structure is fixed all over the sphere, it is possible to assign different learning weights to different directions, as shown in Fig. 4a. Let $\mathcal{N}_i(k)$ denote the index of k -th neighbor of vertex i ($k = 1 \dots 8$ representing {SW, W, NW, N, NE, E, SE, S}), and let L_{in} and L_{out} be the number of input and output features in the convolution, respectively. For each output feature l , $1 \leq l \leq L_{out}$, we define the convolution operation as:

$$x_i^l = \langle \Theta_0, \mathbf{x}_i \rangle + \sum_{k=1}^8 \langle \Theta_k, \mathbf{x}_{\mathcal{N}_i(k)} \rangle \cdot w_{\mathcal{N}_i(k),i}, \quad (2)$$

where Θ_k is the learnable weights of the CNN filters, \mathbf{x}_i represents the input signal/features at vertex i , and $\langle \cdot, \cdot \rangle$ denotes the inner product operator. Both Θ_k and \mathbf{x}_i are L_{in} -length vectors. To handle the 24 exceptions of HEALPix pixels that do not have 8 neighbors, we define:

$$w_{\mathcal{N}_i(k),i} = \begin{cases} 0, & \text{if the neighbor } \mathcal{N}_i(k) \text{ is missing} \\ 1, & \text{otherwise} \end{cases}$$

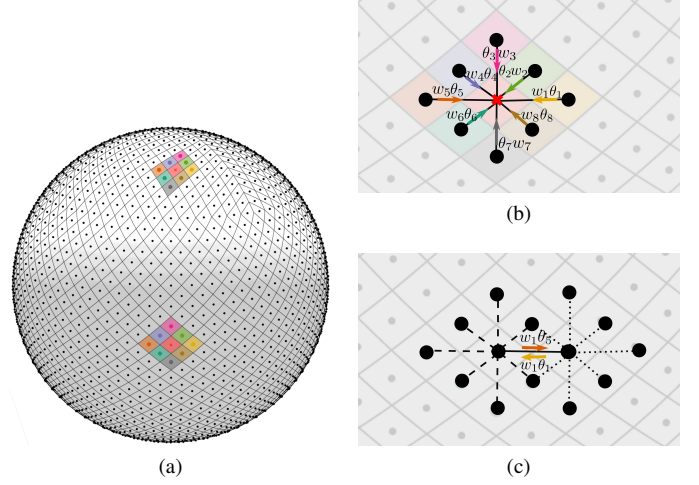


Fig. 4. Our spherical convolution. (a) Kernel on the sphere at different positions. Each color corresponds to a different learnable kernel weight corresponding to a different orientation. (b) Convolution operation can be seen as a message-passing scheme. (c) The same pair of pixels contribute differently depending on the position of central pixel.

Interestingly, the proposed filter is anisotropic and consistent all over the sphere (properties (i.a) and (ii) in Table I). Note also that our proposed convolution can be interpreted as an update rule of a message-passing algorithm applied to graphs [33]. Such an interpretation is of great importance since it allows us to define convolution with larger filter sizes and therefore further gains in terms of filter expressiveness. The proposed convolution, however, is a generalization of the classical update rule in graph message-passing since the weighting applied to a message depends on the direction of the edge, as illustrated in Fig. 4c.

D. Increasing the local support to n -hop neighborhood

Extending the support of the filters is of great importance in order to increase the filter expressiveness (*i.e.*, property (i) in Table I). This requires generalizing the filtering operation (2) to n -hop neighborhoods and, therefore, being able to identify the nodes in these larger neighborhoods. However, direct identification of the neighbors and their relative positions to the central pixel is a complex operation and requires a large amount of storage space that grows quadratically with n . More precisely, for n -hop neighbors we need to store the index of $(2n+1)^2 - 1$ neighbors per pixel and per layer with different resolutions. In this section, we propose a fast and efficient iterative convolution that permits to automatically compute the convolution for any n -hop neighborhood. Our solution relies on the aggregation of n 1-hop neighborhood convolutions. This allows us to approximate the heavy convolution of a large neighborhood by iterative operation of lighter convolutions that only require access to direct neighbors. Although the expressiveness, in terms of the number of learned parameters, may be decreased with such a strategy, the resulting filter remains directional and quite complex, especially for small n . Moreover, in practice, this strategy has obtained promising results in the context of 2D CNN and graph-based CNN [34], [35].

Theoretically, a cascade of n neighborhood aggregation makes use of a subtree structure of depth n rooted at every pixel. Although deeper versions of neighborhood aggregation have access to more information, they may not provide better performance as the local information is diluted by too much averaging through successive neighborhood aggregation. Additionally, with deep neighborhood aggregation, vanishing gradient problem is encountered in the backpropagation of the training phase. For example, it has been observed that the best performance with graph convolutional networks is achieved with a cascade of two 1-hop convolutions [35]. To overcome this limitation, skip connections are utilized in computer vision and graph learning [35] where the output of each iteration is sent to a further layer. Finally, a pixel-wise aggregation mechanism combines the features of intermediate layers independently. Formally, let $\mathbf{z}_i^{(l)}, l = 1 \dots n$ denote the hidden feature of pixel i at layer l that is sent to the last layer as shown in Fig. 5. A skip strategy is applied to reach an overall filter support size of n . More precisely, each intermediate layer aggregates messages from the neighbors of the previous layer. In the last layer, an aggregation layer combines the features of each pixel from all the intermediate layers.

In this paper, we follow a similar idea and investigate the best aggregation strategy for our spherical convolution. We consider 3 different strategies for the last aggregation layer:

- **Concatenation** aggregation uses the feature maps of all preceding convolutional block outputs by concatenating them together [35]:

$$\mathbf{z}_i^{(1)} \parallel \dots \parallel \mathbf{z}_i^{(n)}, \quad (3)$$

where \parallel represents the operation of concatenation.

- **Max** aggregation selects the most informative layer for each pixel [35]

$$\max(\mathbf{z}_i^{(1)}, \dots, \mathbf{z}_i^{(n)}). \quad (4)$$

- **Addition** adopts summation to connect all preceding feature maps, as in residual networks [34],

$$\sum_{l=1}^n \mathbf{z}_i^{(l)}. \quad (5)$$

The choice of the aggregation method may eventually depend on the target task. In the compression scenario investigated in the next Section, we have proven that the addition strategy is the most performing approach.

Another advantage of the recursive implementation of n -hop neighborhood filtering lies in the memory usage. Indeed, with HEALPix, the access to the neighbor indices of a pixel is not direct and needs to be stored. With this implementation, only the 1-hop neighborhood indices need to be stored.

E. Stride in convolution

In classical CNN architectures, the stride controls how the filter moves over the input signal during convolution. Stride enables to reduce the amount of computation and perform both convolution and downsampling operations simultaneously. However, defining stride on a sphere is not




Fig. 5. Proposed strategy to increase the local support (here 3-hop) in spherical convolution. Each $\text{conv}_l, l = 1, 2, 3$ performs neighborhood aggregation for each pixel as given in (2). Finally, the last layer aggregation module combines pixel-wise the generated feature maps of the intermediate layers.

straightforward since a simple downsampling rule such as “one pixel over two” does not exist.

In our OSLO solution, we resort to the HEALPix *nested* numbering scheme [29] to define stride. This scheme arranges the pixel indices in 12 tree structures, where each tree represents one pixel of the coarse base resolution. Finer pixelization is obtained by splitting each pixel into 2×2 pixel structures. The embedded pixel indexing is obtained by labeling a child pixel through the concatenation of the parent index and two additional bits. For example, pixel index 5 (101, in binary notation) of the base HEALPix resolution (Fig. 6a) is the parent of pixels [20...23] ($101b_1b_2$ in binary notation) in the first resolution (Fig. 6b). This way, the data elements that are nearby on the sphere surface are also nearby in the tree structure of the database; hence the neighbor searches are conducted efficiently in memory.

Based on the above scheme, we define a stride of $n \times n$ on the sphere by sliding the filter at every n^2 pixels, *i.e.*, visiting pixels with a step of n^2 . Due to the HEALPix subdivision scheme, which divides each parent into 4 children, n must be a power of 2. For convolutions with hop greater than 1 (Section III-D), all intermediate layers (conv_1 to conv_2 in Fig. 5) are convolved with stride 1. Finally, only the last layer (conv_3 in Fig. 5) is convolved with stride $n \times n$, which is then combined with subsampled versions of the previous layers in the aggregation step. We adopt such a strategy of not using stride in the intermediate layers so that the message passing algorithm can successfully share the information of all the n -hop pixels until the last convolution layer. Examples of 2×2 strides for different resolutions are shown in Fig. 6. Pixel indices in red indicate the visiting pixels where the convolution filter will be placed.

F. Pooling/Unpooling

1) *Pooling*: is a multipurpose tool in CNN. It reduces the size of the feature maps and reduces overfitting by being robust to local variations. Additionally, it tends to compute orientation and position invariant features at least locally. Pooling can be easily implemented by exploiting the hierarchical structure of HEALPix sampling. Indeed, as shown in Fig. 6, the sampling is constructed from low to high resolution, where a pixel at low resolution corresponds to four equal-area sub-pixels in the next higher resolution. Therefore, for pooling, we apply the process in reverse order and merge the child pixels into one value, assigned to the parent pixel. Formally, let $\mathcal{C}(i)$ denote the set of children of the parent pixel i . The pooling function $\mathbb{R}^{N_{pix} \times D} \rightarrow \mathbb{R}^{N'_{pix} \times D}$ performs for each feature map, $\forall d \in$




Fig. 6. Hierarchical pixelization of the sphere together with the nested pixel numbering scheme. (a) HEALPix base resolution. Red lines show base-pixel boundaries. (b) HEALPix first resolution each base pixel is divided into 2×2 pixels. Blue lines indicate the boundaries of child pixels for base-resolution (c) HEALPix second resolution. The black lines indicate the boundaries of child pixels. The pixel indices where the center of convolution filter is placed on them to have a stride of 2×2 are shown in red.

$[1..D]$, and at every pixel position $\forall i \in [1..N'_{pix}]$ the following operation:

$$y_{i,d} = f(\{x_{j,d} \mid j \in \mathcal{C}(i)\}), \quad (6)$$

where D is the number of feature maps, N_{pix} and N'_{pix} are the numbers of pixels before and after pooling, respectively [8], and f is a function that operates on a set (for instance f may compute the maximum or the average over the set). Note that the resulting subsampling factor is a power of 4: $N'_{pix}/N_{pix} = 4^n$, $n \in \mathbb{N}$, since the HEALPix subdivision scheme divides each parent pixel into 4 sub-pixels.

2) *Unpooling (with sub-pixel convolution)*: is the inverse of pooling, *i.e.*, a layer for increasing input resolution, and is necessary to reconstruct the input data. Unpooling can either be achieved by transpose convolution or sub-pixel convolution. Since both constructions achieve similar performance [36], we restrict the presentation to on-the-sphere *sub-pixel convolution*. However, similar ideas can be deployed to generalize transpose convolution to the sphere.

Sub-pixel convolution technique can be interpreted as a *standard convolution in low-resolution space* followed by a *periodic shuffling operation* known as pixel shuffle layer [37]. Now let us consider the convolution output that is still in low-resolution and made of $4^n \times D$ channels with N_{pix} number of pixels, where $D, n \in \mathbb{N}$. A pixel shuffle transformation consists of reorganizing it to D channels with $4^n \times N_{pix}$ number of pixels. Formally, this function can be written as:

$$\mathbb{R}^{N_{pix} \times 4^n \cdot D} \rightarrow \mathbb{R}^{N_{pix} \cdot 4^n \times D}. \quad (7)$$

Defining such an unpooling operation for the spherical images is not straightforward. Our construction takes advantage of the hierarchical HEALPix sampling. We propose to first perform convolution as in (2) to produce $L_{out} = 4^n \cdot D$ feature maps with all exactly the same low resolution as the input feature maps, where $D, n \in \mathbb{N}$. Then, the L_{out} output feature parts are partitioned into D sets of 4^n maps each. Then, pixel shuffling is implemented. More precisely, in each set, the 4^n maps with N_{pix} -resolution are combined into a single map with N'_{pix} -resolution, where $N'_{pix} = 4^n N_{pix}$. This is achieved by assigning to a high-resolution-block of size 4^n , the pixel values of their low-resolution-parents in the 4^n maps of the set, as illustrated in Fig. 7.




Fig. 7. Sub-pixel convolution on the sphere. Convolution is applied on the low-resolution sphere to produce “ $4^n \cdot D$ ” feature maps/channels ($n, D \in \mathbb{N}$). Then, the output feature maps are partitioned into D sets of 4^n feature maps, and in each set, the 4^n channels of the same parent pixel are organized into a specific channel of child pixels.

G. Patching

A patch is a subsection of a signal, *e.g.*, a 2D crop of an image. In CNNs, due to GPU memory limitations and to speed up the training process, random patches sampled from higher resolution inputs are commonly used. Due to the local properties of convolutions, patches can be successfully used to train the network, while the full resolution input can be used during inference. Since omnidirectional images are commonly much higher resolution than traditional images, patching is a must-have feature in our OSLO framework so that we can efficiently train OSLO-based networks using current GPUs. The hierarchical property of HEALPix is used to define patches. To define a random patch of size 2^n by 2^n on the sphere sampled with HEALPix at resolution r , one needs to randomly choose a parent pixel on the sphere sampled at resolution $r - n$ and apply CNN to its corresponding child pixels. For instance, to define random patches of size 4 by 4 pixels for the second resolution of HEALPix (Fig. 6c), let’s assume pixel 5 is randomly selected at the base resolution (Fig. 6a). Then, CNN is applied to all its 16 child pixels [80...95] at the second resolution. The neighboring pixels that are outside the patch do not contribute in convolution operation, *i.e.*, for a pixel i in the patch boundary, if its adjacent pixel j is outside of the patch, $w_{j,i}$ is set to zero in Eq. (2).

Thanks to the above proposed CNN modules, OSLO supports the development of several kinds of architectures for different applications related to 360° images. In other words, by redefining, on-the-sphere, all the blocks commonly used in 2D CNN, many currently successful architectures applied in 2D images can be adapted to omnidirectional images as well. As an example, the next section highlights the benefits of the OSLO solution in the omnidirectional image compression application.

IV. APPLICATION TO END-TO-END OMNIDIRECTIONAL IMAGE COMPRESSION

A. OSLO compression model

We evaluate our novel OSLO framework in the context of end-to-end deep image compression. We use state-of-the-art representation learning frameworks that perform well for regular images, and re-define such architecture for omnidirectional images, using the on-the-sphere learning operators defined in Section III. For that purpose, we use the two well-known 2D

deep learning image compression architectures of [16], [17]. They are two auto-encoder-like architectures, and they enable extremely effective coding performance for perspective 2D images. Note that our methodology could have been deployed for any reference 2D compression architecture [38], [39].

The model of [17] is an extension of [16] that explicitly estimates the entropy model with hyperpriors to effectively capture spatial dependencies in the latent representation. Fig. 8 shows an overview of the two architectures. They involve two types of autoencoders: (1) an image autoencoder architecture (e, d) in both [16] and [17], (2) a hyperprior network (e_s, d_s) only in [17]. The image encoder uses a parametric analysis transform e (encoder) to transform the image vector \mathbf{x} into a latent representation \mathbf{y} , which is subsequently quantized to $\hat{\mathbf{y}}$. In order to make the network end-to-end trainable, the quantization component, which is not differentiable by nature, is approximated by an additive uniform noise [40] in the training phase. The additional hyperprior network in [17] estimates the likelihood of \mathbf{y} with a Gaussian entropy model. For that, \mathbf{y} is fed into e_s (encoder of statistics) to summarize the distribution of standard deviations in $\boldsymbol{\nu}$, which is then quantized into $\hat{\boldsymbol{\nu}}$ and losslessly compressed as side information with an entropy coder. Then $\hat{\boldsymbol{\nu}}$ is fed into d_s (decoder of statistics) to estimate the spatial distribution of standard deviations $\hat{\sigma}$. Finally, the image autoencoder uses $\hat{\sigma}$ to losslessly compress $\hat{\mathbf{y}}$. The decoder also uses $\hat{\sigma}$ to successfully recover $\hat{\mathbf{y}}$. It then feeds $\hat{\mathbf{y}}$ into a parametric synthesis transform d (decoder) to obtain the reconstructed image $\hat{\mathbf{x}}$. The loss function of the model eventually maximizes the reconstruction quality and minimizes the bit rate:

$$\mathcal{L} = D(\mathbf{x}, \hat{\mathbf{x}}) + \lambda \cdot R(\hat{\mathbf{y}}, \hat{\boldsymbol{\nu}}), \quad (8)$$

where $D(\cdot, \cdot)$ indicates the distortion between the input and the reconstruction, R represents the bit rate, and λ is a scalar that balances the reconstruction quality and the bit rate. It is used to set the step of the quantization applied to the latent description.

We transform the two architectures on the sphere using the OSLO framework, as shown in Fig. 9. The 3×3 and 5×5 2D convolutional filters correspond to 1-hop and 2-hop filters on the sphere, respectively. For the 2-hop layer aggregation strategy, we use addition aggregation of (5) as it provides better performance than concatenation and max aggregations (see ablation study in Section IV-D). A 2×2 2D stride operator is equivalent to stride 4 on the sphere, as explained in Section III-E. GDN/IGDN nonlinearities are transformations that are suitable for density modeling and compression of images utilizing local divisive normalization [41]. Since the GDN/IGDN are applied pixel-wise, we use them without modifications.

B. Experimental setup

We evaluate the performance of on-the-sphere learning modules on the compression of 360° images under different conditions. We use 2170 images from the SUN360 equirectangular image database for our experiments [42]. All experiments use the same split of the dataset for learning and evaluation,




Fig. 8. Overview of the two network architectures used in our experiment [16], [17]. Q represents quantization, and AE, AD represent arithmetic encoder and arithmetic decoder, respectively.

consisting of 1737 images for training (80%), 10 images for validation (0.5%), and 423 images for testing (19.5%). To have a fair comparison, we train all CNN models (our model and baselines) using this dataset from scratch (no pre-trained model is used).

We compare our OSLO solution with different baselines. On the one hand, we consider the solutions of [16], [17] trained and tested on ERP. On the other hand, we consider the same architecture, as in Fig. 9, defined on HEALPix sampling, using DeepSphere’s architecture, *i.e.*, graph-based convolution and max pooling [8]. To increase the local support of the Deep sphere convolutional filters, the extension to greater than one hop is implemented by increasing the Chebyshev filter size as suggested in [8], [43] instead of the strategy proposed for OSLO in Section III-D. Additionally, as no upsampling strategy is introduced in DeepSphere, we use our proposed sub-pixel convolution (Section III-F2) for DeepSphere to have a fair comparison. Finally, we also consider the classical JPEG image compression algorithm on ERP for the sake of completeness.

For the evaluation, we use the testing procedure recommended by JVET for 360° image compression [44], [45]: The high-resolution SUN360 equirectangular images of size 9104×4552 are considered as the ground truth for quality evaluation. To remove the unfair bias due to the fact that the ground-truth images are available in ERP format, we downsample ERP images to a resolution of 4992×2496 and HEALPix to 12582912 pixels, such that both have almost the same number of pixels (equivalent to 4K resolution), as suggested in [45], [46]. Considering that 360° images are spherical signals by nature, the Peak signal-to-noise ratio (PSNR) does not reflect the actual 360° image quality. Thus, we use *Spherical PSNR* (S-PSNR) [46] and *Weighted to Spherically uniform PSNR* (WS-PSNR) [47] as the objective quality evaluation.

In S-PSNR the ground-truth and the decompressed images are first mapped to a sphere that is uniformly sampled with 655,362 points [46]. To ensure a fair comparison, the uniform sampling used for the S-PSNR is different from HEALPix

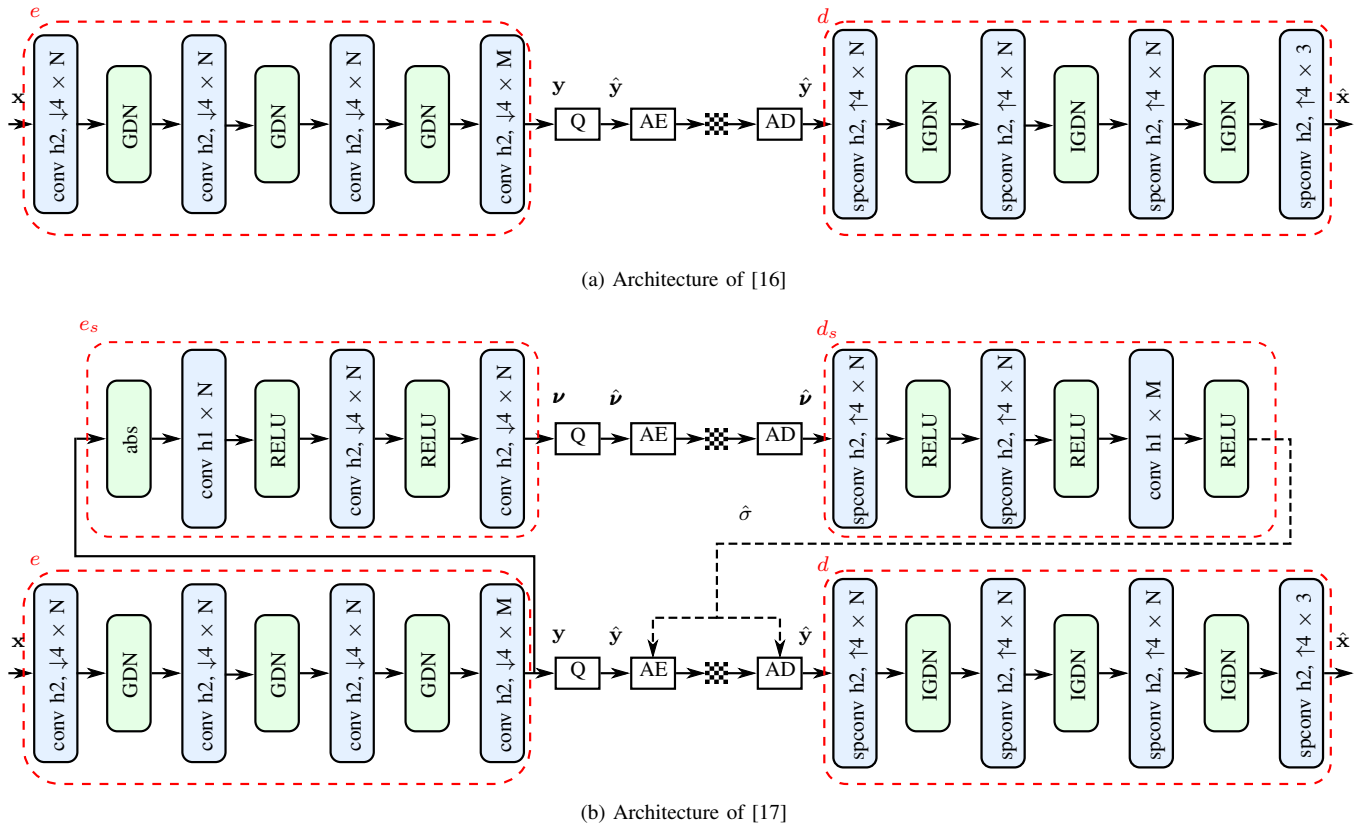


Fig. 9. Equivalent architectures of [16], [17] for spherical images. Convolution parameters are denoted as: “conv $hk, \downarrow l \times N$ ” representing a k -hop convolution (see Fig. 5), downsampled by $l \times$ number of filters. Similarly, sub-pixel convolution parameters are denoted as: “spconv $hk, \uparrow l \times N$ ” representing a k -hop convolution, followed by upsampling of l with pixel shuffle transformation (see Fig. 7) \times number of filters.

sampling. Then the mean error between the original and decompressed signals sampled with these uniformly distributed points is calculated to compute the PSNR on the sphere. Bilinear interpolation is used to calculate values in fractional pixels. The S-PSNR value can approximate the average quality of all possible views presented to the observers. Alternately, WS-PSNR does not require remapping, and the error of each pixel is weighted to obtain the equivalent spherical area in the observation space [47].

To train the 2D models, crops of randomly positioned patches of 256×256 pixels from the training images are extracted as suggested in [17]. Similarly, our proposed model is also trained by randomly cropping the sphere into a region of the same size, as explained in Section III-G. We use the Pytorch implementation of the architecture of [16], [17] given in [48] for training the 2D ERP models. We modify and adapt this library to implement our spherical equivalent. In all learning-based models, we use batches of 10 training images to perform stochastic gradient descent with the Adam algorithm [49] and learning rate of 10^{-4} . The training is performed for 1000 epochs, and the validation set is introduced at epoch 800 to dynamically reduce the learning rate by factor of 0.316 when the validation loss is no longer improving more than 0.0001 for 10 epochs.

C. Rate-Distortion results

Given a test image set denoted by \mathcal{X} , we compute for each $i \in \mathcal{X}$ and for several λ , the rate-distortion pair $(R(i, \lambda), D(i, \lambda))$ achieved by each architecture of interest. The aggregated Rate-Distortion performance of each architecture is then computed as in [17] by averaging over \mathcal{X} :

$$R_\lambda = \frac{1}{|\mathcal{X}|} \sum_{i \in \mathcal{X}} R(i, \lambda) \quad (9)$$

$$D_\lambda = \frac{1}{|\mathcal{X}|} \sum_{i \in \mathcal{X}} D(i, \lambda). \quad (10)$$

Fig. 10 present, for each architecture, the (R_λ, D_λ) for different values of λ . We can see that our OSLO-based approach significantly outperforms the baseline methods. The DeepSphere solution can not reconstruct the image properly, and even when the bitrate is increased, the performance gets saturated because the isotropic filters are not expressive enough to reconstruct the image. Comparing 2D learning architectures applied to ERP with JPEG compression reveals that, similar to the results given in [16], [17] for perspective images, these models do not lose performance when the input train/test data is ERP, and they still perform much better than JPEG. We can conclude that 2D learning approaches can keep their efficiency compared to JPEG even with ERP format. However, our proposed solution is able to yet outperform the ERP representation due to the efficient learning tools of OSLO that are directly defined on the sphere.

The compressed results between OSLO and 2D deep learning architecture applied to equirectangular projection are also visually compared. Comparison between OSLO and the DeepSphere’s architecture [8], that use the same HEALPix sampling, will be studied in Section IV-E. Figs. 11 and 12 illustrate viewports of spherical images (the portion of spherical images seen by HMD devices at a specific direction) that are compressed almost at the same rate with the same architecture as [16]. It can be seen that an image compressed with OSLO has less distortion and keeps the detail, for example, the texts are more visible.

In addition to these aggregated plots, the Bjontegaard-Delta (BD) rate [50] is also calculated separately for each test image, allowing us to measure bitrate reduction while maintaining the same quality. More precisely, for each test image RD curve the BD-rate between 2D learning method on ERP (as the reference) and our proposed OSLO is calculated. Negative values of the BD rate represent percentage saving over the reference. For the architecture of [16], OSLO gains -18.3% and -12.4% in terms of S-PSNR and WS-PSNR respectively. Regarding the architecture of [17], the average BD-rate is -13.7% in terms of S-PSNR and -8.6% in terms of WS-PSNR, demonstrating OSLO’s superiority over ERP. The frequency histograms of BD-rate gains over the test images under SPSNR and WS-PSNR metrics are shown in Fig. 13 for different architectures. The figure clearly shows that OSLO improves the compression performance significantly.

D. Ablation study: hop+n aggregation

We now evaluate the 3-layer aggregation strategies introduced in Section (III-D). To have a fair comparison, we choose the number of feature maps N and M in Fig. 9 such that all three strategies have almost the same number of learning parameters. Since all experiments use HEALPix sampling, we only consider WS-PSNR as our objective quality metric, which evaluates the performance in the sampling domain. Results are presented in Fig. 14. We observe that, for our spherical convolution, the addition aggregation (5) outperforms the concatenation (3) and max aggregations (4). This justifies the study of aggregation methods, as for graph convolution with regular structure instead, concatenation is shown to be the best aggregation method [35].

E. Filter visualization

In this experiment, we compare our learned filters to those generated from the DeepSphere solution (using graph convolution). We recall that both use the same HEALPix sampling. We randomly select filters at different layers of the network and visualize them in Fig. 15. Graph convolution results isotropic filters where each pixel at the same distance from the central one is filtered with the same weight. Such a filter shape obviously limits the expressiveness of the feature maps, which is especially important in image compression. On the contrary, OSLO provides anisotropic filters and can thus be more expressive. To evaluate the impact of this filter’s expressiveness on the reconstructed image, we show in Fig. 16, the decoded image with both DeepSphere and OSLO-based

solutions. It is clear that our solution preserves the details of the image and requires a much lower number of bits.

V. CONCLUSION

We proposed a new representation learning framework for omnidirectional images. The efficient learning operators of 2D CNN models for regular images are redefined on-the-sphere, exploiting the benefits of HEALPix sampling. The proposed OSLO solution allows reaching the expressiveness of a 2D filter while having the advantage of uniform sampling on the sphere, which preserves a low level of complexity. We have finally applied the OSLO solution to the omnidirectional image compression problem by adapting state-of-the-art end-to-end image compression architectures to work on the sphere. Our experiments have shown that our on-the-sphere strategy leads to impressive coding gains when compared with mapping-based or graph-based solutions and demonstrate the promising potential of our OSLO framework for the representation of spherical data, which goes beyond the compression application. Future work includes instantiating OSLO in other contexts for deep learning-based processing of omnidirectional images.

REFERENCES

- [1] X. Hou, S. Dey, J. Zhang, and M. Budagavi, “Predictive adaptive streaming to enable mobile 360-degree and vr experiences,” *IEEE Transactions on Multimedia*, vol. 23, pp. 716–731, 2021.
- [2] D.-Y. Nam and J.-K. Han, “An efficient algorithm for generating harmonized stereoscopic 360° vr images,” *IEEE Transactions on Circuits and Systems for Video Technology*, pp. 1–1, 2021.
- [3] L. Deng, M. Yang, H. Li, T. Li, B. Hu, and C. Wang, “Restricted deformable convolution-based road scene semantic segmentation using surround view cameras,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 10, pp. 4350–4362, 2020.
- [4] J. Snyder, *Flattening the Earth: Two Thousand Years of Map Projections*. University of Chicago Press, 1993.
- [5] K.-T. Ng, S.-C. Chan, and H.-Y. Shum, “Data compression and transmission aspects of panoramic videos,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 15, no. 1, pp. 82–95, Jan 2005.
- [6] E. Kuzyakov and D. Pio, “Next-generation video encoding techniques for 360 video and vr,” 2016. [Online]. Available: <https://code.fb.com/virtual-reality/next-generation-video-encoding-techniques-for-360-video-and-vr/>
- [7] C. Fu, L. Wan, T. Wong, and C. Leung, “The rhombic dodecahedron map: An efficient scheme for encoding panoramic video,” *IEEE Transactions on Multimedia*, vol. 11, no. 4, pp. 634–644, June 2009.
- [8] N. Perraudin, M. Defferrard, T. Kacprzak, and R. Sgier, “DeepSphere: Efficient spherical convolutional neural network with HEALPix sampling for cosmological applications,” *Astronomy and Computing*, vol. 27, pp. 130–146, Apr. 2019.
- [9] R. Khasanova and P. Frossard, “Geometry aware convolutional filters for omnidirectional images representation,” in *International Conference on Machine Learning (ICML)*. PMLR, 2019, pp. 3351–3359.
- [10] C. Esteves, C. Allen-Blanchette, A. Makadia, and K. Daniilidis, “Learning SO(3) equivariant representations with spherical CNNs,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [11] T. S. Cohen, M. Geiger, J. Koehler, and M. Welling, “Spherical CNNs,” in *International Conference on Learning Representations (ICLR)*, 2018.
- [12] C. Esteves, A. Makadia, and K. Daniilidis, “Spin-weighted spherical CNNs,” in *Advances in Neural Information Processing Systems (NIPS)*, 2020.
- [13] P. J. Roddy and J. D. McEwen, “Sifting convolution on the sphere,” *IEEE Signal Processing Letters*, vol. 28, pp. 304–308, 2021.
- [14] M. Uyttendaele, “Optimizing 360 photos at scale,” 2017. [Online]. Available: <https://engineering.fb.com/2017/08/31/ml-applications/optimizing-360-photos-at-scale/>





Fig. 10. Rate-distortion curves aggregated over the test images. The solid, dashed, dotted, and dash-dotted lines represent OSLO implementation, 2D deep learning architectures on Equirectangular projections, Deepsphere architectures, and Equirectangular JPEG compression. Red color refers to the architecture of [17], blue color indicates [16] architecture, and green color indicates JPEG compression. (a) S-PSNR. (b) WS-PSNR.

- [15] B. Davidson, M. S. Alvi, and J. F. Henriques, “360° camera alignment via segmentation,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [16] J. Ballé, V. Laparra, and E. P. Simoncelli, “End-to-end optimized image compression,” in *International Conference on Learning Representations (ICLR)*, 2017.
- [17] J. Ballé, D. Minnen, S. Singh, S. J. Hwang, and N. Johnston, “Variational image compression with a scale hyperprior,” in *International Conference on Learning Representations (ICLR)*, 2018.
- [18] Y. Yagi, “Omnidirectional sensing and its applications,” *IEICE transactions on information and systems*, vol. 82, no. 3, pp. 568–579, 1999.
- [19] N. Padmanaban, T. Ruban, V. Sitzmann, A. M. Norcia, and G. Wetzstein, “Towards a machine-learning approach for sickness prediction in 360 stereoscopic videos,” *IEEE transactions on visualization and computer graphics*, vol. 24, no. 4, pp. 1594–1603, 2018.
- [20] Y.-C. Su and K. Grauman, “Learning spherical convolution for fast features from 360° imagery,” in *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [21] R. Monroy, S. Lutz, T. Chalasani, and A. Smolic, “Salnet360: Saliency maps for omni-directional images with cnn,” *Signal Processing: Image Communication*, vol. 69, pp. 26–34, 2018.
- [22] M. Ruder, A. Dosovitskiy, and T. Brox, “Artistic style transfer for videos and spherical images,” *International Journal of Computer Vision*, vol. 126, no. 11, pp. 1199–1219, 2018.
- [23] Y. Lee, J. Jeong, J. Yun, W. Cho, and K.-J. Yoon, “Spherephd: Applying cnns on a spherical polyhedron representation of 360deg images,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 9181–9189.
- [24] C. Zhang, S. Liwicki, W. Smith, and R. Cipolla, “Orientation-aware semantic segmentation on icosahedron spheres,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 3533–3541.
- [25] T. Cohen, M. Weiler, B. Kicanaoglu, and M. Welling, “Gauge equivariant convolutional networks and the icosahedral cnn,” in *International Conference on Machine Learning (ICML)*. PMLR, 2019, pp. 1321–1330.
- [26] I. Tomic, P. Frossard, and P. Vandergheynst, “Progressive coding of 3-d objects based on overcomplete decompositions,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 16, no. 11, pp. 1338–1349, 2006.
- [27] R. Khasanova and P. Frossard, “Graph-based classification of omnidirectional images,” in *Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCVW)*, 2017, pp. 869–878.
- [28] Q. Yang, C. Li, W. Dai, J. Zou, G.-J. Qi, and H. Xiong, “Rotation equivariant graph convolutional network for spherical image classification,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 4302–4311, ISSN: 2575-7075.
- [29] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, “HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere,” *The Astrophysical Journal*, vol. 622, no. 2, pp. 759–771, apr 2005.
- [30] C. M. Jiang, J. Huang, K. Kashinath, P. Marcus, M. Niessner *et al.*, “Spherical cnns on unstructured grids,” in *International Conference on Learning Representations (ICLR)*, 2018.
- [31] N. Krachmalnicoff and M. Tomasi, “Convolutional neural networks on the HEALPix sphere: a pixel-based algorithm and its application to CMB data analysis,” *Astronomy & Astrophysics*, vol. 628, p. A129.
- [32] T. Maugey, O. Le Meur, and Z. Liu, “Saliency-based navigation in omnidirectional image,” in *IEEE 19th International Workshop on Multimedia Signal Processing (MMSp)*, 2017, pp. 1–6.
- [33] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, “Neural message passing for quantum chemistry,” in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. PMLR, 06–11 Aug 2017, pp. 1263–1272.
- [34] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [35] K. Xu, C. Li, Y. Tian, T. Sonobe, K.-i. Kawarabayashi, and S. Jegelka, “Representation learning on graphs with jumping knowledge networks,” in *Proceedings of the 35th International Conference on Machine Learning (ICML)*, vol. 80, pp. 5453–5462.
- [36] W. Shi, J. Caballero, L. Theis, F. Huszar, A. Aitken, C. Ledig, and Z. Wang, “Is the deconvolution layer the same as a convolutional layer?” 2016.



(a) 1980.89 KB



(b) 1920.86 KB

Fig. 11. A viewport at a specific direction of spherical images compressed almost at the same rate. A zoom version of each colored rectangular region is shown on the right. (a) Equirectangular image compressed with [16] in 1980.89 KB. (b) The same image compressed with OSLO in 1920.86 KB.

- [37] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, “Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 1874–1883.
- [38] Y. Hu, W. Yang, Z. Ma, and J. Liu, “Learning end-to-end lossy image compression: A benchmark,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021.
- [39] Z. Cheng, H. Sun, M. Takeuchi, and J. Katto, “Learned image compression with discretized gaussian mixture likelihoods and attention modules,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 7936–7945.
- [40] J. Ballé, V. Laparra, and E. P. Simoncelli, “End-to-end optimization of nonlinear transform codes for perceptual quality,” in *2016 Picture Coding Symposium (PCS)*, 2016, pp. 1–5.
- [41] —, “Density modeling of images using a generalized normalization transformation,” in *International Conference on Learning Representations (ICLR)*, 2016.
- [42] J. Xiao, K. A. Ehinger, A. Oliva, and A. Torralba, “Recognizing scene viewpoint using panoramic place representation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 2695–2702.
- [43] M. Defferrard, X. Bresson, and P. Vandergheynst, “Convolutional neural networks on graphs with fast localized spectral filtering,” in *Advances in Neural Information Processing Systems (NIPS)*. Red Hook, NY, USA: Curran Associates Inc., 2016, p. 3844–3852.
- [44] Z. Chen, Y. Li, and Y. Zhang, “Recent advances in omnidirectional video coding for virtual reality: Projection and evaluation,” *Signal Processing*, vol. 146, pp. 66–78, 2018.
- [45] J. Boyce, E. Alshina, A. Abbas, and Y. Ye, “JVET common test conditions and evaluation procedures for 360° video,” joint Video Exploration Team (JVET) of ITU-T SG 16 WP 3 and ISO/IEC JTC 1/SC 29/WG 11, Macao, CNm 2017.
- [46] M. Yu, H. Lakshman, and B. Girod, “A framework to evaluate omnidirectional video coding schemes,” in *IEEE International Symposium on Mixed and Augmented Reality*, 2015, pp. 31–36.
- [47] Y. Sun, A. Lu, and L. Yu, “Weighted-to-spherically-uniform quality evaluation for omnidirectional video,” *IEEE Signal Processing Letters*, vol. 24, no. 9, pp. 1408–1412, 2017.
- [48] J. Bégin, F. Racapé, S. Feltman, and A. Pushparaja, “Compressai: a pytorch library and evaluation platform for end-to-end compression research,” *arXiv preprint arXiv:2011.03029*, 2020.
- [49] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations (ICLR)*, 2015.
- [50] G. Bjontegaard, “Calculation of average psnr differences between rd-curves,” *VCEG-M33*, 2001.





Fig. 12. A viewport at a specific direction of spherical images compressed almost at the same rate. A zoom version of each colored rectangular region is shown on the right. (a) Equirectangular image compressed with [16] in 1766.94 KB. (b) The same image compressed with OSLO in 1613.62 KB.




Fig. 13. Frequency histogram for BD-rate gains of OSLO over 2D learning in ERP. Top row represents SPSNR and WSPSNR BD-rate gains for the architecture of [16]. Bottom row represents the results for [17].




Fig. 15. Random selection of learned filters at different layers of e and s . The weight corresponding to the central pixels is not shown in the filter for better visibility. Left column represents filters learned with DeepSphere that results in isotropic filters. Right column represents the filters learned with OSLO solution. Two filters on the same row are taken from the exact same position in the architecture.




Fig. 14. Rate-distortion curves of different aggregation strategies to increase the local support of convolution.



Fig. 16. Decompressed images in Mollweide projection. The zoomed versions of red, green, and blue rectangular regions are shown in the bottom row. (a) DeepSphere result stored in 2408.43 KB with 29.54 dB and 29.51 dB for SPSNR and WSPSNR respectively. (b) OSLO result stored in 936.65 KB with 40.91 dB and 41.65 dB for SPSNR and WSPSNR respectively.