UvA-DARE (Digital Academic Repository)

ChebLieNet: Invariant Spectral Graph NNs Turned Equivariant by Riemannian Geometry on Lie Groups

Aguettaz, H.; Bekkers, E.J.; Defferrard, M.

Link to publication

ChebLieNet: Invariant Spectral Graph NNs Turned Equivariant by Riemannian Geometry on Lie Groups

Hugo Aguettaz
EPFL, Lausanne, Switzerland
`hugo.aguettaz@epfl.ch`

Erik J Bekkers
UvA, Amsterdam, Netherlands
`e.j.bekkers@uva.nl`

Michaël Defferrard
EPFL, Lausanne, Switzerland
`michael.defferrard@epfl.ch`

Abstract

We introduce ChebLieNet, a group-equivariant method on (anisotropic) manifolds. Surfing on the success of graph- and group-based neural networks, we take advantage of the recent developments in the geometric deep learning field to derive a new approach to exploit any anisotropies in data. Via discrete approximations of Lie groups, we develop a graph neural network made of anisotropic convolutional layers (Chebyshev convolutions), spatial pooling and unpooling layers, and global pooling layers. Group equivariance is achieved via equivariant and invariant operators on graphs with anisotropic left-invariant Riemannian distance-based affinities encoded on the edges. Thanks to its simple form, the Riemannian metric can model any anisotropies, both in the spatial and orientation domains. This control on anisotropies of the Riemannian metrics allows to balance equivariance (anisotropic metric) against invariance (isotropic metric) of the graph convolution layers. Hence we open the doors to a better understanding of anisotropic properties. Furthermore, we empirically prove the existence of (data-dependent) sweet spots for anisotropic parameters on CIFAR10. This crucial result is evidence of the benefice we could get by exploiting anisotropic properties in data. We also evaluate the scalability of this approach on STL10 (image data) and ClimateNet (spherical data), showing its remarkable adaptability to diverse tasks.

1   Introduction

Deep learning is a class of machine learning algorithms inspired by the human brain's network of neurons [Goodfellow et al., 2016]. These algorithms use a hierarchical structure of neural layers to extract higher-level features from the raw input progressively. In the past few years, the growing computational power of modern GPU-based computers and the availability of large training datasets in the field of machine learning have made it possible to successfully train neural networks with many layers and degrees of freedom. Consequently, deep learning has revolutionized many machine learning tasks in recent years, ranging from image and video processing to speech recognition and natural language understanding.

Many neuroscientific research results served as focal points in the development of deep learning algorithms. When Hubel and Wiesel [1962] studied the visual cortex in the brain, they made three important discoveries. First, they observed a one-to-one correspondence between spatial locations in the retina and neurons in the brain that fired as a response to line-like visual stimuli. Second, the activity of the neurons changed depending on the orientation of the line, uncovering a neat organization based on local orientations. Last, the neurons sometimes fired only when the line was

moving in a particular direction. Later, Bosking et al. [1997] showed that neurons that are aligned fire together, indicating the presence of a type of long-range interactions. All these results motivated the development of a mathematical framework for modeling visual perception based on sub-Riemannian geometry on the space of positions and orientations, which is typically modeled with the Lie group SE(2) [Petitot, 2003, Citti and Sarti, 2006, Duits et al., 2014]. Apart from the neurophysiological inspiration, group equivariance has also been proven to be an excellent inductive bias [Cohen and Welling, 2016] not only in computer vision (as the translation equivariance property of CNNs as shown) but also in physics [Finzi et al., 2020] and molecular data analysis [Fuchs et al., 2021, Jumper et al., 2020]. In this work, we propose to build group equivariant graph neural networks via the same principle that underlie the sub-Riemannian, neurogeometrical modeling of the visual cortex.

Our work connects the observations by Hubel and Wiesel [1962] and Bosking et al. [1997] on two levels. First, the organization of visual data based on their location and orientation [Hubel and Wiesel, 1962] is modeled by Lie group convolutions [Bekkers, 2019], in which feature maps encode response for every position and every orientation. Second, long-range interactions between aligned neurons [Bosking et al., 1997] are modeled by building graphs with affinity matrices based on (approximate) sub-Riemannian distances on the Lie groups, inspired by sub-Riemannian image analysis methods such as [Franken and Duits, 2009, Bekkers et al., 2015, Favali et al., 2016, Mashtakov et al., 2017, Boscain et al., 2018, Duits et al., 2018, Baspinar et al., 2021].

Defferrard et al. [2020] showed how to construct powerful graph NNs that are faithful to the manifolds on which they are defined. Nevertheless, the layers themselves are based on rotationally invariant (Laplacian) convolutions. In order to exploit directional cues in the data, group convolutions are desirable [Cohen et al., 2018, Kondor and Trivedi, 2018, Cohen and Welling, 2016, Bekkers, 2019]. However, since Laplacian operators are intrinsically isotropic, there is no point applying them to the lifted feature maps on the group unless we construct anisotropic metrics on the groups. Therefore, we adopt the Lie group viewpoint by Sanguinetti et al. [2015] to define anisotropic Riemannian metrics based on left-invariant vector fields on the group. Once an anisotropic Riemannian graph is constructed, any spectral method can directly be applied to this graph. The resulting graph neural networks will then, by construction, be equivariant and capable of utilizing directional cues in data.

Before going further into the details, we summarize our main contributions:

- We introduce ChebLieNet, an equivariant graph Laplacian-based neural network based on Lie groups equipped with an anisotropic Riemannian metric.
- The Riemannian geometry is automatically derived from a standard base space (e.g. $\mathbb{R}^2$ or the sphere), which makes our approach flexible and effective in building group equivariant graph neural networks for a variety of data structures (e.g. 2D and spherical data).
- We demonstrate the equivariance property of ChebLieNet, both in theory and in practice. This property guarantees that the neural network's predictions are robust against given transformations, which is not necessarily the case with methods based on data augmentation.
- We show that the use of directional information via anisotropic Riemannian spaces could benefit many tasks.
- We show the flexibility of the method by considering two different problems; we validate on classification problems with 2D image data and a segmentation problem on spherical data via the construction of a sub-Riemannian geometry on $SE(2)$ and $SO(3)$ respectively.

2 Related works

2.1 Group equivariant convolutional neural networks

Deep convolutional neural networks [LeCun et al., 1995] have proven to be compelling models for pattern recognition tasks on images, video, and audio data. Although a robust theory of neural network design is currently lacking, a large amount of empirical evidence supports the notion that both convolutional weight sharing, depth, and width are essential for good predictive performance. Such properties are enabled through the equivariance property of convolutions (convolving a shifted image is the same as translating its result).

Lenc and Vedaldi [2015] showed that the AlexNet CNN Krizhevsky et al. [2012] trained on ImageNet learns representations equivariant to flips, scalings, and rotations spontaneously. This supports the

idea that equivariance is an excellent inductive bias for deep convolutional networks. In the last few years, a joint effort has been made to build group equivariant networks. By the introduction of group convolutions in deep learning, Cohen and Welling [2016] generalize the translation equivariance property to larger groups of symmetries, including rotations and reflections. Kondor and Trivedi [2018] gave a rigorous, theoretical treatment of convolution and equivariance in neural networks concerning any compact group's action. One of the main contributions of that work was to show that, given some natural constraints, the convolutional structure is not just a sufficient but also a necessary condition for equivariance to a compact group's action. In a similar spirit, in [Bekkers, 2019] it is shown that any bounded linear operator is equivariant to Lie groups if and only if it is a group convolution. In our work, we propose to build group equivariant neural networks via left-invariant Laplace operators on Lie groups, which indeed can be seen as group convolutions with kernels that are the fundamental solutions of the Laplace operator. The result is a Lie group equivariant Chebyshev-type neural network [Defferrard et al., 2016] that we will refer to as ChebLieNet.

2.2 Graph neural networks

Using the term geometric deep learning, Bronstein et al. [2017, 2021] give an overview of deep learning methods in the non-Euclidean domain, including graphs and manifolds. They present different examples of geometric deep learning problems and available solutions, fundamental difficulties, applications, and future research directions in this nascent field.

One of the main challenges when working with graph data it to deal with the inter-dependencies between points. Indeed, the derivations of most standard machine learning models firmly base on an independence assumption. For this reason, transferring existing methods on a graph appears doomed to failure, and it seems necessary to build models acting directly on graphs. Due to its success on Euclidean data, the development of a convolution-like operator on graphs has been largely studied. Because the notion of space is not naturally defined on a graph, we lack a straightforward generalization of the convolutional operator from grid data to graphs [Scarselli et al., 2008, Bruna et al., 2013, Henaff et al., 2015, Defferrard et al., 2016, Kipf and Welling, 2016, Masci et al., 2015, Boscaini et al., 2016, Monti et al., 2017].

Spectral approaches have a solid mathematical foundation in graph signal processing. Rather than using the traditional spatial definition of the convolution, it proposes to see this operation from a spectral perspective. Based on the convolution theorem, it defines the convolution operator from the graph spectral domain via the eigendecomposition of the graph Laplacian (see App. A.3).

Definition 2.1 (Spectral graph convolution) *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \boldsymbol{W})$ be a graph with Laplacian $\hat{\boldsymbol{\Delta}}$ and let f and g be two functions defined on $\mathcal{V}$. We define the $\mathcal{G}$-convolution $*_\mathcal{G}$ of f and g as:*

$$f *_\mathcal{G} g = \boldsymbol{\Phi}(\hat{\boldsymbol{g}} \odot \hat{\boldsymbol{f}}) = \boldsymbol{\Phi}(\boldsymbol{\Phi}^\top \boldsymbol{g} \odot \boldsymbol{\Phi}^\top \boldsymbol{f}), \tag{1}$$

with eigenvectors $\boldsymbol{\Phi}$ obtained through the unique eigendecomposition $\hat{\boldsymbol{\Delta}} = \boldsymbol{\Phi}\boldsymbol{\Lambda}\boldsymbol{\Phi}^T$.

While this definition alleviates the difficulty of deriving a convolution operator in the spatial domain, other difficulties arise. First of all, because the Laplacian of a graph is an intrinsic operator, it is domain-dependent, and the spectral-convolution is too. It implies that a model built on this framework cannot be easily transferred from a graph to another as expressed in a different "language". Nevertheless, this is not a problem for us since we are focusing on fixed manifold graphs. Next, there is no guarantee that filters represented in the spectral domain are spatially localized. Henaff et al. [2015] successfully bypassed this problem by defining smooth spectral filter coefficients, arguing that if spectral filters are smooth, they are spatially localized. Last but not least, the Laplacian's eigendecomposition makes the method expensive in terms of memory and time. Indeed, the forward and inverse graph Fourier transforms (via $\boldsymbol{\Phi}^T$ and $\boldsymbol{\Phi}$) incur expensive multiplications as no FFT-like algorithm exists on general graphs. Defferrard et al. [2016] alleviated the cost of explicitly computing the graph Laplacian using spatially-localized filters with Chebyshev polynomials.

Definition 2.2 (Chebyshev convolutional layer) *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \boldsymbol{W})$ be a graph with rescaled Laplacian[1] $\tilde{\boldsymbol{\Delta}}$, $\boldsymbol{x} \in \mathbb{R}^{|\mathcal{V}| \times d_i}$ be an input features' vector and $\Theta_j \in \mathbb{R}^{d_i \times d_o}$ learnable filters. The*

---

[1]Because Chebyshev polynomials are defined in the range $[-1, 1]$, it is necessary to rescale the graph Laplacian with $\tilde{\boldsymbol{\Delta}} = 2\lambda_{\max}^{-1}\hat{\boldsymbol{\Delta}} - \boldsymbol{I}$ where $\lambda_{\max}$ is the largest eigenvalue of $\hat{\boldsymbol{\Delta}}$.

output features' vector $\boldsymbol{y} \in \mathbb{R}^{|\mathcal{V}| \times d_o}$ is computed as:

$$\boldsymbol{y} = \sum_{j=0}^{R-1} \boldsymbol{z}_j \boldsymbol{\Theta}_j \quad \textit{with} \quad \boldsymbol{z}_0 = \boldsymbol{x}, \quad \boldsymbol{z}_1 = \tilde{\boldsymbol{\Delta}} \boldsymbol{x} \quad \textit{and} \quad \boldsymbol{z}_j = 2\tilde{\boldsymbol{\Delta}} \boldsymbol{z}_{j-1} - \boldsymbol{z}_{j-2}. \quad \forall j \geq 2. \quad (2)$$

Kipf and Welling [2016] simplified this formulation a bit by considering the construction of single-parametric filters that are linear with relation to $\tilde{\boldsymbol{\Delta}}$. They further approximate $\lambda_{\max} \simeq 2$ as they expect that neural network parameters will adapt to this change in scale during training.

3 Method

Our method can be seen as an extension of the original ChebNet [Defferrard et al., 2016, Perraudin et al., 2019]. Instead of directly working on a homogeneous base space, we first extend it to a higher dimensional space (Lie group). The goal of this extension is to convert the previously invariant spectral convolutional layers into equivariant layers.[2]

3.1 Anisotropic manifold graph

In order to define the anisotropic manifold graphs we have to consider two types of manifolds. The base manifold $\mathcal{M}$ and a Lie group G that acts transitively on $\mathcal{M}$. The latter implies that $\mathcal{M}$ is a homogeneous space of G, which means that any two points $m_1, m_2 \in \mathcal{M}$ can be mapped to each other via the action of a group element $g \in G$ via $m_2 = g \cdot m_1$. E.g., the plane $\mathcal{M} = \mathbb{R}^2$ is a homogeneous space of the special Euclidean motion group $G = SE(2)$ as any two points can be mapped to each other through a rotation and a translation. Such groups G, which have $\mathcal{M}$ as a homogeneous space, can always be split in two parts via the semi-direct product $G = \mathcal{M} \rtimes H$, with H a sub-group of G that leaves some reference point $m_0 \in \mathcal{M}$ invariant, i.e., $\forall_{h \in H} : m_0 = h \cdot m_0$. E.g., rotations leave the zero vector in $\mathcal{M} = \mathbb{R}^2$ invariant, and thus $H = SO(2)$ in the $SE(2)$ case. Conversely, any homogenous space can be modeled with a group quotient $\mathcal{M} = G/H$.

We define an anisotropic manifold graph to be a discretization of a Lie group G of which $\mathcal{M}$ is a homogeneous space. It consists of a finite set of vertices corresponding to a random sampling of group elements, and a finite set of similarity-based edges that are constructed via a left-invariant Riemannian metric on G. In our work we consider two anisotropic manifold graphs: one associated with the base manifold $\mathcal{M} = \mathbb{R}^2$ which we extend with an additional orientation/rotation dimension $H = SO(2)$ to come to the Lie group $G = SE(2) = \mathbb{R}^2 \rtimes SO(2)$, and the other associated with the sphere $\mathcal{M} = S^2$ which we similarly "lift" to the Lie group $G = SO(3)$ by adding an additional rotation dimension. Considering the similarity between the two cases (the sphere locally looks like $\mathbb{R}^2$) we will refer to $\mathcal{M}$ as the "spatial" part, and H as the "orientation" part of the group.

Uniform sampling of the vertices. The first step to construct an anisotropic manifold graph is to sample elements on the group uniformly or as uniformly as possible if the manifold does not permit a uniform grid. We split the grid construction in two parts, a grid on $\mathcal{M}$ which is sampled with $|\mathcal{V}_s|$ points and a grid on H that is sampled with $|\mathcal{V}_o|$ points, leading to a total of $|\mathcal{V}| = |\mathcal{V}_s||\mathcal{V}_o|$ vertices.

Left-invariant anisotropic Riemannian distance. Once vertices have been uniformly sampled on the group manifold, a similarity measure between vertices is computed. This measure is based on a Riemannian distance between points in G. The only thing one needs in our algorithm is the implementation of the logarithmic map on the Lie group (see e.g. [Bekkers, 2019]), and a diagonal Riemannian metric tensor (see e.g. [Sanguinetti et al., 2015] and [Mashtakov et al., 2017] for the $SE(2)$ and $SO(3)$ case respectively). In the following we provide the essential idea and intuition behind the construction of the similarity measure and provide a more extensive treatment in App. B.

In Riemannian geometry on Lie groups it is common to express tangent vectors of curves in a basis of left-invariant vector fields as it allows to measure their lengths with a single Riemannian metric tensor that is shared over the entire group. This works as follows. Consider curve $\gamma : [0, 1] \rightarrow G$ with its tangent vectors $\dot{\gamma}(t) = \sum_{i=1}^{d} u^i(t) \mathcal{A}_i|_{\gamma(t)}$ expressed in a basis/moving frame of reference

---

[2]Because spectral graph NNs are able to capture the geometry of the space, which in this work we equip with anisotropic metrics, any spectral method could be made equivariant using our method.

4

$\{\mathcal{A}_i|_{\gamma(t)}\}_{i=1}^d$, in which $\mathcal{A}_i$ are left-invariant vector fields. The length of these tangent vectors is then measured by a Riemannian metric tensor that we denote with $\|\dot{\gamma}(t)\|_{\mathbf{R}}^2 := \mathbf{u}(t)^T \mathbf{R} \mathbf{u}(t)$, with $\mathbf{R}$ a symmetric positive definite matrix defined relative to the basis $\{\mathcal{A}_i|_{\gamma(t)}\}_{i=1}^d$, and with $\mathbf{u}(t) = (u_0(t), u_1(t), \dots)^T$. The $\mathcal{A}_i$ are left-invariant vector fields and the notation $\mathcal{A}_i|_g$ means the vector in the vector field $\mathcal{A}_i$ at location g. The vector fields are constructed by choosing a vector A_i in the tangent space at origin (the Lie algebra) which then defines a complete vector field on G via the push-forward of left-multiplication. In less technical terms this means that if we pick a direction vector at the origin, and we move it to another point in, e.g. $G = SE(2)$, via a roto-translation, this vector will move and rotate along. By defining everything in terms of these left-invariant vector fields, every tangent space $T_g(G)$ at each $g \in G$ can be identified with the tangent space at the origin, and a single Riemannian metric tensor $\mathbf{R}$ can be shared over the entire space. Moreover, the induced Riemannian distance $d(g, h)$ between any two points $g, h \in G$ is then by construction left-invariant, i.e., $\forall_{g,h,i \in G} : d(g \cdot h, g \cdot i) = d(h, i)$.

Expressing tangent vectors in such left-invariant vector fields allows us to reason in terms of the generators of the group and even design learnable equivariant differential operators for the construction of deep NNs [Smets et al., 2020]. Consider the $G = SE(2)$ case. As a basis we pick the 3 generators of the group: a forward motion represented by a vector A_1 pointing in the forward direction within the plane, a side-ways motion represented by a perpendicular planar vector A_2, and a rotation/change of orientation represented by a vector A_3 that points vertically in along the H-dimension. We then work with diagonal Riemannian metric tensors $\mathbf{R} = \mathrm{diag}(1, \epsilon^{-2}, \xi^2)$, which penalize each type of motion (represented by the vector components) differently. When $\epsilon \to 0$ one arrives at the *sub-Riemannian geometry* which forms the basis for the mathematical modeling of visual perception. It quantifies a notion of alignment through the sub-Riemannian distance; the length of a distance-minimizing geodesic that connects two local orientations that lie in the extend of each other will be much smaller that that of a geodesic connecting two local orientations parallel to each other. An analogy can be found with the example of a car in a parking lot where it can move forward/backward (A_1) and change orientation (A_3) [Reeds and Shepp, 1990]. It will be easier to move it to the more aligned spot directly ahead then it will to the spot next to the car, as sideways motion (A_2) is impossible.

Parameters ϵ and ξ will respectively be referred to as spatial and orientation anisotropy parameters. With $\epsilon = 1$ the metric is isotropic and there will be no distinction between different orientations. When $\epsilon < 1$, ξ determines the flexibilty/curvature of the geodesics as it balances spatial motion against angular motion. In a sense it defines how easily one connects local orientations that are not optimally aligned. In Figure 1 this behavior is visualized by running a diffusion process on the anisotropic manifold graph. In the anisotropic case ($\epsilon < 1$) diffusion is faster along the forward direction within a θ-plane. From a graph NN perspective this suggests that information is propagated more quickly between vertices that are aligned, nevertheless, Chow's theorem (see e.g. [Montgomery, 2006]) guarantees that any point pair in the (sub-)Riemannian manfiold can interact with one another.

The exact computation of the (sub-)Riemannian distances is challenging and can generally not be done in closed form, but can be done numerically via method such as [Bekkers et al., 2015, Sanguinetti et al., 2015, Mashtakov et al., 2017]. In order to keep our graph construction algorithm efficient though, we will approximate the Riemannian distances via an efficient analytic formula based on those in [Bekkers et al., 2018] that only involves the Lie group's logarithmic map $\log : G \to T_e(G)$ and the Riemannian metric tensor $\mathbf{R}$. We then approximate the distance between points $g, h \in G$ by

$$d(g, h) = d(e, g^{-1} \cdot h) \simeq \| \log(g^{-1} \cdot h) \|_{\mathbf{R}}. \tag{3}$$

Similarity measure. Encoding a similarity measure in the edges of a graph requires defining a weighting scheme. It is common to use a Gaussian kernel and set the weights via

$$w(v_i, v_j) = \begin{cases} \exp\left(-\frac{d^2(v_i, v_j)}{4t}\right) & \text{if } e(v_i, v_j) \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases}. \tag{4}$$

The choice for kernel bandwidth t is essentially arbitrary, but good heuristics exist. Perraudin et al. [2019] set it to half the average squared distance between connected vertices. Defferrard et al. [2020], however, showed that this heuristic has the tendency to overestimate it and preferred to choose it as the minimizer of the mean equivariance error. Following this overestimation observation, we fix the kernel bandwidth as 20% of the average squared Riemannian distance between connected vertices. As such, the weights diversely cover values in the whole range $[0, 1]$. The most similar vertices are connected with close-to-one weighted edges whereas the lowest connections are close to zero.

Quality of the approximation. In theory, we would like our approximation to be as precise as possible. In practice, a high-resolution approximation leads to computational issues in time and memory. Hence, tuning of the graph parameters becomes a trade-off between theoretical consistency and practical feasibility. First of all, the graph resolution (or the number of vertices we sample) is directly related to the quality of the approximation. While the spatial resolution $|\mathcal{V}|_s$ is usually determined by the data (up to up- and down-samplings), the orientation resolution $|\mathcal{V}|_o$ is a design choice. An important remark is to notice that a large orientation resolution does not necessarily help if two different orientations are not distinguishable because of a poor spatial resolution [Weiler et al., 2018, Bekkers, 2019]. Secondly, the connectivity of the graph is also a crucial parameter. A fully connected graph is theoretically the best approximation. Nevertheless, for computational reasons, we use K-NN graphs[3] to sparsify the graph Laplacians.

(a) Base space $\mathcal{M}$ with an isotropic Riemannian metric

(b) Lie group extension G with an isotropic Riemannian metric

(c) Lie group extension G with an anisotropic Riemannian metric
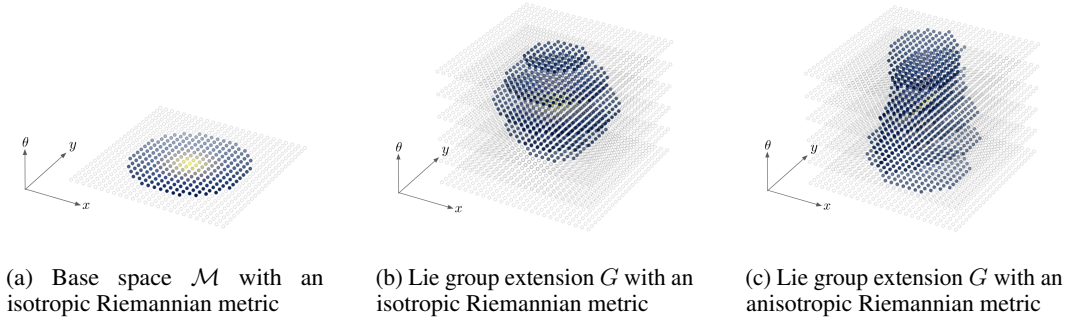
Figure 1: Isotropic diffusion applied to an impulse signal on Riemannian manifolds on $\mathcal{M} = \mathbb{R}^2$ and $G = SE(2)$.

Theoretical group equivariance of the graph Laplacian. Due to the success of machine learning algorithms based on graph Laplacian, the theoretical convergence of the graph Laplacian to its continuous analogue has been largely studied [Hein et al., 2005, Singer, 2006]. Belkin and Niyogi [2006] noticed that in many graph-based algorithms, a central role is played by the graph Laplacian's eigenvectors. Thus, they focused on proving convergence in eigenmaps as it is sufficient in this case. They proved that if the graph's vertices are sampled uniformly from an unknown submanifold $\mathcal{M} \in \mathbb{R}^d$, then the eigenvectors of a suitably constructed graph Laplacian converges to the eigenfunctions of the Laplace-Beltrami operator on $\mathcal{M}$. Consequently, as the latter operator is left-invariant, as we show in theorem A.1, the graph Laplacian is asymptotically [4] group equivariant.

Empirical group equivariance of the graph Laplacian. We empirically confirm the group equivariance property of the graph Laplacian applied to our anisotropic manifold graphs. By checking $\boldsymbol{P}^\top \tilde{\boldsymbol{\Delta}} \boldsymbol{P} = \tilde{\boldsymbol{\Delta}}$ where $\boldsymbol{P}$ is a permutation matrix, we can verify that the graph Laplacian is invariant under a given permutation of vertices corresponding to a group transformation (e.g. a rotation of the graph). Moreover, we can also compare the eigenmaps of a graph Laplacian and its continuous counterpart if it is well-known. For a further discussion about this, see App. E.

3.2 ChebLieNet

Chebyshev convolutional layer. As introduced in Defferrard et al. [2016], a Chebyshev convolutional layer is a spectral layer based on a continuous kernel parametrization with graph Laplacians. This parameterization makes such layers highly suitable for our method, as they intrinsically capture the Riemannian geometry of the graphs on G. Moreover, the Chebyshev convolutions on the anisotropic manifold graphs are equivariant by construction because the graph Laplacians are equivariant operators (see Figure 2).

---

[3]Note that in our implementation, a K-NN graphs does not mean that each vertex has K neighbors but at most K neighbors. Indeed, if the graph domain has boundaries, using exactly K neighbors for each vertex could lead to asymmetries that may introduce biases and harm the permutation invariances in the graph.

[4]The asymptotic case corresponds to $|\mathcal{V}| \to \infty$ and a Gaussian weight kernel with kernel bandwidth $t \to 0$.

Figure 2: Rotation equivariance of a randomly initialized $SE(2)$ Chebyshev convolutional layer. From left to right shows different rotations of an input (top row) and the activations for different slices of $\theta \in [0, \pi]$ in the graph (bottom 6 rows). A rotation of an input image followed by Chebyshev convolution is equivalent to first convolution followed by a planar rotation in each θ slice and a roll in the θ-axis.

Spatial pooling and unpooling layers. Graph pooling is a central component in a myriad of graph neural network architectures. Producing coarsened graphs from a finer graph have two main advantages: first, it reduces the computational cost, and second, it could improve performance by reducing the overfitting effect and adding a multiscale perspective. As an inheritance from traditional CNNs, most approaches formulate graph pooling as a cluster assignment problem, extending local patches' idea in regular grids to graphs [Dhillon et al., 2007, Ying et al., 2018, Khasahmadi et al., 2020, Mesquita et al., 2020]. We propose similar operations on the base space (spatial domain) and involving two steps (see Figure 3). First, each sample is assigned to a cluster that will correspond to the output sample; this is the down- (resp. up-) sampling phase. With a well designed method, this change of data-resolution can be made equivariant to any group transformation.[5] Then, each cluster is reduced (resp. expanded) according to a given scheme (e.g. maximum, average or random); this is the reduction (resp. expansion) phase. When the reduction and expansion steps are permutation-invariant operations, such layers are automatically invariant under any transformation in the group.


(a) R2RandPool

(b) S2MaxPool

(c) R2RandUnpool

(d) S2AvgUnpool

Figure 3: Spatial pooling and unpooling layers on the 2D grid and the sphere.

Global pooling (projection) layer and point-wise operations. When the neural network does not need to be equivariant but invariant (e.g. classification task), it is common to rely on a global pooling layer (or simply projection layer). This layer reduces the d-dimensional signal on the graph's vertices to a d-dimensional vector of features derived from information on the whole graph. As a permutation-invariant operation, such a layer does not break the equivariance property of the neural network. Finally, point-wise operations do not affect the equivariance of a neural network.

---

[5]Altough down- and up-samplings are naturally defined on the Euclidean grid, this task is more complicated on the sphere. However, using an icosahedron decomposition of the sphere, we make it more natural as down- and up-sampling consists of decreasing or increasing the subdivision level.

4 Experiments

In this section, we show the benefits of working on the anisotropic manifold graphs compared to the base manifold graphs. We believe that further improvements could be achieved through tuning and hyper-parameter optimization of the models [Yu and Zhu, 2020], using high-capacity networks, or via a more advanced training process, but this is not the goal of our work. We here intent to illustrate the adaptability of our approach to different tasks such as classification and segmentation in 2D images or spherical data. In the first couple of experiments, we motive the use of anisotropic spaces. By varying the anisotropies, we show the existence of sweet spots, both for the spatial anisotropy parameter ϵ and the orientation anisotropy parameter ξ. In the second couple of experiments, we show that even if we add a new orientation dimension, our method remains scalable using a proper implementation.

Our implementation is fully PyTorch [Paszke et al., 2019] and available at `https://github.com/haguettaz/ChebLieNet`. We perform all the experiments on a single GeForce GTX 1080 Ti gpu and track them with the Weights & Biases library [Biewald, 2020]. The details of the experiments are given in the App. C.

4.1 Why using tunable anisotropic kernels?

As introduced in Section 3.1, the anisotropies are tunable via the parameters ϵ and ξ of the Riemannian metric, respectively responsible for the spatial and orientation anisotropies. As the ξ parameter should depend on the spatial and orientation resolutions, we use the following parameterisation: $\xi^2 = \alpha \frac{|\mathcal{V}_o|}{|\mathcal{V}_s|}$. Setting $\alpha = 1$ yields a 40/60 ratio of neighbors within versus outside the orientation plane. We ran different experiments with a Wide Residual architecture [Zagoruyko and Komodakis, 2016] on CIFAR10 [Krizhevsky et al., 2009], varying the spatial and orientation anisotropic parameters.


(a) Test-accuracy against orientation anisotropies

(b) Test-accuracy against spatial anisotropies

Figure 4: Empirical proof of existence of sweet spots for data-dependent anisotropic parameters.

Orientation anisotropy. The orientation anisotropy ξ controls how strongly orientation layers are connected. At the limit $\xi \to \infty$, orientation layers are decoupled. It is like test-time augmentation with rotations: running a CNN working with one anisotropic Laplacian (e.g., only vertically aligned filters) and testing the network for different input rotations before averaging the output. The other extreme $\xi \to 0$ keeps all layers equally close to each other, and features are essentially identified with just a spatial coordinate. This would then correspond to a WideResNet with isotropic Chebyshev convolutions. For reasonable values of ξ, interactions between orientation layers take place. Figure 4a is evidence of the existence of a sweet spot for this parameter in the range of reasonable values. At the moment, we expect with no certainty that this parameter could be set *a priori* of the data, only considering the data resolution. As a rule of thumb, we set ξ such that each vertex has approximately 40% of its neighbors in the same orientation layer and 60% on others.

Spatial anisotropy. The spatial anisotropy ϵ regulates the anisotropy of the space on the spatial domain. For $\epsilon = 1$, the Riemannian metric is spatially isotropic; all directions are treated equally and the resulting model would effectively be a WideResNet with isotropic Chebyshev convolutions. At the limit $\epsilon \to 0$, the main direction has a minimal cost, and the resulting space is highly spatially anisotropic. In figure 4a we observe that using anisotropic spaces instead of isotropic ones is relevant, as we almost get an 8% test-accuracy improvement. Unlike the orientation anisotropic parameter,

in our opinion, this parameter is task/data-dependent; different datasets could benefit in different degrees from the utilization of directional information through different spatial anisotropy settings.

4.2 How scalable is the method?

Scalability is often an important limitation of graph- and group-based neural networks. By adding an orientation dimension, we do not run from this rule as we necessarily increase the number of vertices of the anisotropic manifold graphs. To permit experiments on larger images, it becomes crucial to pre-compute anisotropic manifold graphs and their Laplacians. Dedicated librairies like PyKeops [Charlier et al., 2020] enable this without memory issues. Nevertheless, the graph operations (convolutions, pooling or unpooling) still scale with the size of the graph. Fortunately, PyTorch provides sparse operations that increase efficiency in terms of time and memory compared to dense operations in cases of sufficiently sparse graph Laplacians (typically a sparsity $\mathcal{S}(\tilde{\boldsymbol{\Delta}}) \geq 98.5\%$).

We evaluate our models on an image classification task on STL10 [Coates et al., 2011] and an image segmentation task on ClimateNet [Kashinath et al., 2021]. We show the adaptability of our method by using a Wide Residual architecture [Zagoruyko and Komodakis, 2016] on STL10 and a U-Net-like network [Ronneberger et al., 2015] on ClimateNet. We also demonstrate the potential of our approach and the benefits of using anisotropic spaces. Indeed, while on ClimateNet the use of anisotropies is neither beneficial nor detrimental, the difference in performance on STL10 is significant.

Table 1: Mean of test performance and training duration on ClimateNet and STL10. Errorbars are 1 standard deviation computed over 5 trials.

		ClimateNet		STL10	
ϵ		Test F1	Duration	Test accuracy	Duration
1	(invariant)	$\mathbf{85.62 \pm 0.09\%}$	$\sim 2\,\mathrm{d}$	$68.98 \pm 0.56\%$	$\sim 9\,\mathrm{h}$
0.1	(equivariant)	$85.25 \pm 0.19\%$	$\sim 7\,\mathrm{d}$	$\mathbf{74.02 \pm 1.10\%}$	$\sim 16\,\mathrm{h}$

5 Conclusion

Scope. With our method, geometric graph NNs are made equivariant to Lie groups. Via the groups $SE(2)$ and $SE(3)$, we can construct roto-translation equivariant networks for $2D$ image data and $3D$ volumetric data. Based on the group $SO(3)$, our method can deal with meteorological or cosmological data while preserving rotation equivariance. We believe that our flexible approach is ideal for further explorations on the relevance of group equivariance in tasks not considered in this work.

Limitations. The main weakness of our method is its relatively high memory requirement. Although all experiments ran on a single gpu, by adding an orientation axis, we significantly enlarge the feature maps. As a result, anisotropic graph manifolds are memory-heavier than isotropic ones and prone to a slowdown during the forward- and backward-pass. Nevertheless, with the emergence of geometric deep learning, we expect improvement in the hardware and implementation of graph-oriented operations. Another challenge is the increased number of hyper-parameters for which we only have derived rules of thumb. The graph connectivity and resolutions require a tradeoff between efficiency and quality of the manifold approximation. The anisotropic parameters require an analysis of the dataset and some intuition about the amount of anisotropy to set. With systematic hyper-parameter optimization, we can find an optimal combination, but requires more computational resources.

Potential and future research. Thanks to its easy-to-tune anisotropic properties, our model can be used to better understand anisotropic properties in data. In particular, one could explore the effect of using anisotropic spaces instead of isotropic ones on many tasks and conclude when such anisotropic information is relevant. In this vein, it could also be interesting to derive anisotropic pooling and unpooling layers based on anisotropic spaces instead of isotropic ones as it is usually done. More generally, our method is simple enough to be extended to shapes/surfaces with a Riemannian manifold structure [Cohen et al., 2019], following the framework of coordinate independent convolutions [Weiler et al., 2021]. In this work, we focused on 2D images and spherical data on, but the method is readily extendable to higher dimensional Lie groups such as the $SE(3)$ group to obtain 3D roto-translation equivariant ChebLieNets. Moreover, our method for constructing anisotropic geometries

could directly improve other successful Euclidean distance-based graph NNs such as [Satorras et al., 2021] by making them fully equivariant. Last but not least, despite graph-based algorithms being computationally sub-optimal compared to CNNs, their flexibility is a real asset. We see high potential in the exploration of graph sparsification to reduce computational complexity.

Acknowledgments and Disclosure of Funding

---

[6]`https://userinfo.surfsara.nl/`

References

Emre Baspinar, Luca Calatroni, Valentina Franceschi, and Dario Prandi. A cortical-inspired sub-riemannian model for poggendorff-type visual illusions. *Journal of Imaging*, 7(3):41, 2021.

John R Baumgardner and Paul O Frederickson. Icosahedral discretization of the two-sphere. *SIAM Journal on Numerical Analysis*, 22(6):1107–1115, 1985.

Erik J Bekkers. Retinal image analysis using sub-riemannian geometry in se (2). 2017.

Erik J Bekkers. B-spline cnns on lie groups. In *International Conference on Learning Representations*, 2019.

Erik J Bekkers, Remco Duits, Alexey P Mashtakov, and Gonzalo R Sanguinetti. A PDE Approach to Data-Driven Sub-Riemannian Geodesics in SE(2). *SIAM Journal on Imaging Sciences*, 8(4): 2740–2770, 2015. doi: 10.1137/15M1018460. URL https://doi.org/10.1137/15M1018460.

Erik J Bekkers, Da Chen, and Jorg M Portegies. Nilpotent approximations of sub-riemannian distances for fast perceptual grouping of blood vessels in 2d and 3d. *Journal of mathematical imaging and vision*, 60(6):882–899, 2018.

Mikhail Belkin and Partha Niyogi. Convergence of laplacian eigenmaps. *Advances in neural information processing systems*, 19:129–136, 2006.

Lukas Biewald. Experiment tracking with weights and biases, 2020. URL https://www.wandb.com/. Software available from wandb.com.

Ugo V Boscain, Roman Chertovskih, Jean-Paul Gauthier, Dario Prandi, and Alexey Remizov. Highly corrupted image inpainting through hypoelliptic diffusion. *Journal of Mathematical Imaging and Vision*, 60(8):1231–1245, 2018.

Davide Boscaini, Jonathan Masci, Emanuele Rodoià, and Michael Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 3197–3205, 2016.

William H Bosking, Ying Zhang, Brett Schofield, and David Fitzpatrick. Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *Journal of neuroscience*, 17 (6):2112–2127, 1997.

Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.

Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.

Benjamin Charlier, Jean Feydy, Joan Alexis Glaunès, François-David Collin, and Ghislain Durif. Kernel operations on the gpu, with autodiff, without memory overflows. *arXiv preprint arXiv:2004.11127*, 2020.

Fan RK Chung and Fan Chung Graham. *Spectral graph theory*. Number 92. American Mathematical Soc., 1997.

Giovanna Citti and Alessandro Sarti. A cortical based model of perceptual completion in the roto-translation space. *Journal of Mathematical Imaging and Vision*, 24(3):307–326, 2006.

Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011.

Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999, 2016.

Taco Cohen, Mario Geiger, and Maurice Weiler. A general theory of equivariant cnns on homogeneous spaces. *arXiv preprint arXiv:1811.02017*, 2018.

Taco S Cohen, Maurice Weiler, Berkay Kicanaoglu, and Max Welling. Gauge equivariant convolutional networks and the icosahedral cnn. *arXiv preprint arXiv:1902.04615*, 2019.

Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in neural information processing systems*, pages 3844–3852, 2016.

Michaël Defferrard, Martino Milani, Frédérick Gusset, and Nathanaël Perraudin. Deepsphere: a graph-based spherical cnn. *arXiv preprint arXiv:2012.15000*, 2020.

Inderjit S Dhillon, Yuqiang Guan, and Brian Kulis. Weighted graph cuts without eigenvectors a multilevel approach. *IEEE transactions on pattern analysis and machine intelligence*, 29(11): 1944–1957, 2007.

James R Driscoll and Dennis M Healy. Computing fourier transforms and convolutions on the 2-sphere. *Advances in applied mathematics*, 15(2):202–250, 1994.

R. Duits, S. P. L. Meesters, J.-M. Mirebeau, and J. M. Portegies. Optimal Paths for Variants of the 2d and 3d Reeds–Shepp Car with Applications in Image Analysis. *Journal of Mathematical Imaging and Vision*, February 2018. ISSN 1573-7683. doi: 10.1007/s10851-018-0795-z. URL https://doi.org/10.1007/s10851-018-0795-z.

Remco Duits, Ugo Boscain, Francesco Rossi, and Yuri Sachkov. Association fields via cuspless sub-Riemannian geodesics in SE (2). *Journal of mathematical imaging and vision*, 49(2):384–417, 2014.

Marta Favali, Samaneh Abbasi-Sureshjani, Bart ter Haar Romeny, and Alessandro Sarti. Analysis of vessel connectivities in retinal images by cortically inspired spectral clustering. *Journal of Mathematical Imaging and Vision*, 56(1):158–172, 2016.

Marc Finzi, Samuel Stanton, Pavel Izmailov, and Andrew Gordon Wilson. Generalizing convolutional neural networks for equivariance to lie groups on arbitrary continuous data. In *International Conference on Machine Learning*, pages 3165–3176. PMLR, 2020.

Erik Franken and Remco Duits. Crossing-preserving coherence-enhancing diffusion on invertible orientation scores. *International Journal of Computer Vision*, 85(3):253, 2009.

Fabian B Fuchs, Edward Wagstaff, Justas Dauparas, and Ingmar Posner. Iterative se (3)-transformers. *arXiv preprint arXiv:2102.13419*, 2021.

Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.

Krzysztof M Gorski, Eric Hivon, Anthony J Banday, Benjamin D Wandelt, Frode K Hansen, Mstvos Reinecke, and Matthia Bartelmann. Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622(2):759, 2005.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

Matthias Hein, Jean-Yves Audibert, and Ulrike Von Luxburg. From graphs to manifolds–weak and strong pointwise consistency of graph laplacians. In *International Conference on Computational Learning Theory*, pages 470–485. Springer, 2005.

Mikael Henaff, Joan Bruna, and Yann LeCun. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163*, 2015.

David H Hubel and Torsten N Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, 160(1):106–154, 1962.

Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

John Jumper, R Evans, A Pritzel, T Green, M Figurnov, K Tunyasuvunakool, O Ronneberger, R Bates, A Zidek, A Bridgland, et al. High accuracy protein structure prediction using deep learning. *Fourteenth Critical Assessment of Techniques for Protein Structure Prediction (Abstract Book)*, 22: 24, 2020.

Karthik Kashinath, Mayur Mudigonda, Sol Kim, Lukas Kapp-Schwoerer, Andre Graubner, Ege Karaismailoglu, Leo von Kleist, Thorsten Kurth, Annette Greiner, Ankur Mahesh, et al. Climatenet: an expert-labeled open dataset and deep learning architecture for enabling high-precision analyses of extreme weather. *Geoscientific Model Development*, 14(1):107–124, 2021.

Nicolas Keriven, Alberto Bietti, and Samuel Vaiter. Convergence and stability of graph convolutional networks on large random graphs. *arXiv preprint arXiv:2006.01868*, 2020.

Amir Hosein Khasahmadi, Kaveh Hassani, Parsa Moradi, Leo Lee, and Quaid Morris. Memory-based graph networks. *arXiv preprint arXiv:2002.09518*, 2020.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

Risi Kondor and Shubhendu Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In *International Conference on Machine Learning*, pages 2747–2755. PMLR, 2018.

Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.

Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010. URL `http://yann.lecun.com/exdb/mnist/`.

Yann LeCun, Yoshua Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.

Karel Lenc and Andrea Vedaldi. Understanding image representations by measuring their equivariance and equivalence. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 991–999, 2015.

Jonathan Masci, Davide Boscaini, Michael Bronstein, and Pierre Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 37–45, 2015.

Alexey Mashtakov, Remco Duits, Yu Sachkov, Erik J Bekkers, and Ivan Beschastnyi. Tracking of lines in spherical images via sub-riemannian geodesics in SO(3). *Journal of mathematical imaging and vision*, 58(2):239–264, 2017.

Diego Mesquita, Amauri H Souza, and Samuel Kaski. Rethinking pooling in graph neural networks. *arXiv preprint arXiv:2010.11418*, 2020.

Richard Montgomery. *A tour of sub-Riemannian geometries, their geodesics and applications*. Number 91. American Mathematical Soc., 2006.

Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5115–5124, 2017.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pages 8026–8037, 2019.

Nathanaël Perraudin, Michaël Defferrard, Tomasz Kacprzak, and Raphael Sgier. Deepsphere: Efficient spherical convolutional neural network with healpix sampling for cosmological applications. *Astronomy and Computing*, 27:130–146, 2019.

Jean Petitot. The neurogeometry of pinwheels as a sub-Riemannian contact structure. *Journal of Physiology-Paris*, 97(2-3):265–309, 2003.

Jorg Portegies, Gonzalo Sanguinetti, Stephan Meesters, and Remco Duits. New approximation of a scale space kernel on se (3) and applications in neuroimaging. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 40–52. Springer, 2015.

James Reeds and Lawrence Shepp. Optimal paths for a car that goes both forwards and backwards. *Pacific journal of mathematics*, 145(2):367–393, 1990.

Olinde Rodrigues. *Des lois géométriques qui régissent les déplacements d'un système solide dans l'espace: et de la variation des cordonnées provenant de ces déplacements considérés indépendamment des causes qui peuvent les produire*. 1840.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

Gonzalo Sanguinetti, Erik Bekkers, Remco Duits, Michiel HJ Janssen, Alexey Mashtakov, and Jean-Marie Mirebeau. Sub-riemannian fast marching in se (2). In *Iberoamerican Congress on Pattern Recognition*, pages 366–374. Springer, 2015.

Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks. *arXiv preprint arXiv:2102.09844*, 2021.

Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.

David I Shuman, Sunil K Narang, Pascal Frossard, Antonio Ortega, and Pierre 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):83–98, 2013.

Amit Singer. From graph to manifold laplacian: The convergence rate. *Applied and Computational Harmonic Analysis*, 21(1):128–134, 2006.

Bart Smets, Jim Portegies, Erik Bekkers, and Remco Duits. Pde-based group equivariant convolutional neural networks. *arXiv preprint arXiv:2001.09046*, 2020.

Maurice Weiler, Fred A Hamprecht, and Martin Storath. Learning Steerable Filters for Rotation Equivariant CNNs. In *International Conference on Computer Vision and Pattern Recognition*, 2018.

Maurice Weiler, Patrick Forré, Erik Verlinde, and Max Welling. Coordinate independent convolutional networks–isometry and gauge equivariant convolutions on riemannian manifolds. *arXiv preprint arXiv:2106.06020*, 2021.

Douglas Brent West et al. *Introduction to graph theory*, volume 2. Prentice hall Upper Saddle River, NJ, 1996.

Rex Ying, Jiaxuan You, Christopher Morris, Xiang Ren, William L Hamilton, and Jure Leskovec. Hierarchical graph representation learning with differentiable pooling. *arXiv preprint arXiv:1806.08804*, 2018.

Tong Yu and Hong Zhu. Hyper-parameter optimization: A review of algorithms and applications. *arXiv preprint arXiv:2003.05689*, 2020.

Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.

A Background

A.1 Riemannian geometry

Riemannian geometry is a part of differential geometry that studies Riemannian manifolds, smooth manifolds equipped with a Riemannian metric. A manifold $\mathcal{M}$ is a generalization and abstraction of the notion of a curved surface. It is a topological set that is modeled closely on Euclidean space locally but may vary widely in global properties. It means for each $p \in \mathcal{M}$, one can associate a tangent space $T_p(\mathcal{M}) \subseteq \mathbb{R}^d$, corresponding to the union of all tangent vectors of differentiable curves passing through p. Consequently, at any point $p \in \mathcal{M}$, the tangent space $T_p(\mathcal{M})$ is spanned by a basis $\{\partial_{x_i,p}\}_{i=1}^N$:

$$T_p(\mathcal{M}) = \{\dot{\gamma}(0)|\gamma : \mathbb{R} \to \mathcal{M} \in \mathcal{C}^1 \text{and } \gamma(0) = p\} = \text{span}\{\partial_{x_i,p}\}_{i=1}^N. \tag{5}$$

A Riemannian manifold is a differentiable manifold $\mathcal{M}$ with a Riemannian metric $\mathcal{R}$, a 2-tensor field, such that at each $p \in \mathcal{M}$, we have a function $\mathcal{R}|_p : T_p(\mathcal{M}) \times T_p(\mathcal{M}) \to \mathbb{R}$ which is symmetric and positive definite. At each $g \in G$, the Riemannian metric can be expressed by its local representation, a symmetric positive definite matrix $\boldsymbol{R}_p$ whose components are given by:

$$r_{ij}(p) = \mathcal{R}|_p(\partial_{x_i,p}, \partial_{x_j,p}). \tag{6}$$

In the following we consider Riemannian metric with diagonal local representation at the origin e, that is:

$$r_{ij}(e) = \begin{cases} r_i & \text{if } i = j \\ 0 & \text{otherwise} \end{cases} . \tag{7}$$

The Riemannian metric induces an inner product, such that for any $u = u^i\partial_{x_i,e}$ and $v = v^i\partial_{x_i,e}$, we have $\langle u, v \rangle_{\mathcal{R}(e)} = u^i r_i v^i$ and $||u||^2_{\mathcal{R}(e)} = u^i r_i u^i$. Between each pair of points p and q of a Riemannian manifold, we define their Riemannian distance as the length of the shortest curve connecting the two points (a.k.a. geodesic). More formally, we have:

$$d(p,q) = \inf_{\gamma \in \mathcal{C}_{p,q}^\infty} \int_0^1 ||\dot{\gamma}(\tau)||_{\mathcal{R}_{\gamma(\tau)}} d\tau, \tag{8}$$

where $\mathcal{C}_{p,q}^\infty = \{\gamma : [0,1] \to \mathcal{M}|\gamma \in \mathcal{C}^\infty, \gamma(0) = p, \gamma(1) = q\}$. An more convenient way to define the Riemannian between two points is via the exponential and logarithmic Riemannian maps (Figure 5). The Riemannian exponential map $\exp_p T_p(\mathcal{M}) \to \mathcal{M}$ is defined as:

$$\exp_p(v) = \gamma_{p,v}(1), \qquad (\text{and by extension } \exp_p(tv) = \gamma_{p,v}(t)), \tag{9}$$

where $\gamma_{p,v}(1)$ is the unique Riemannian geodesic starting at $p \in \mathcal{M}$ with initial velocity $v \in T_p(\mathcal{M})$. We can interpret the exponential map as follows. Let's choose any direction in our tangent space and follow it with a step forward. We make sure to take the shortest path and end up at a new point. This process is what we called the exponential map.[7] The inverse mapping of the exponential map, the logarithmic map $\log_p : \mathcal{M} \to T_p(\mathcal{M})$ will map an element $p \in \mathcal{M}$ to the smallest vector $v = \log_p q \in T_p(\mathcal{M})$ as measured by the Riemannian metric such that $q = \exp_p v \in \mathcal{M}$.

A.2 Group theory

A group $(G, \cdot)$ is a set G equipped with a binary operation $\cdot : G \to G$ called group product, satisfying the four group axioms (closure, associativity, identity element and inverse elements). To map the structure of the group to some mathematical object, one requires a representation. We define H as the vector space to which our mathematical object belongs and $\mathcal{B}(H)$ the space of bounded linear invertible operators $H \to H$. A representation $\mathcal{V} : G \to \mathcal{B}(H)$ maps a group element to an operator such that the identity element, the group product and the group inverse are preserved. We define the left-regular representation $\mathcal{L}_g$ on the (infinite-dimensional) vector space of functions $G \to \mathbb{R}^d$ via:

$$(\mathcal{L}_g \circ f)(h) = f(g^{-1}h), \tag{10}$$

with $f : G \to \mathbb{R}^d$ a function on the group G and g, h elements of the group G. Using this group representation, we can formally define the left-invariance (a.k.a. equivariance).

[7] It comes from the fact that all these tiny steps magically resemble the series expansion of the exponential function.

Figure 5: Riemannian exponential map $\exp_p : T_p(\mathcal{M}) \to \mathcal{M}$ and Riemannian logarithmic map $\log_p : \mathcal{M} \to T_p(\mathcal{M})$.

Definition A.1 (Equivariance) *An operator $\Phi : \mathcal{X} \to \mathcal{Y}$ from one vector space to the other is equivariant (or left-invariant under group transformation) if it satisfies the following property:*

$$\mathcal{L}'_g \circ \Phi = \Phi \circ \mathcal{L}_g, \qquad \forall g \in G. \tag{11}$$

Equivariance can be realized in many ways, and in particular, the group representations $\mathcal{L}_g$ and $\mathcal{L}'_g$ need not be the same, as they act on different spaces $\mathcal{X}$ and $\mathcal{Y}$. Note that the familiar concept of invariance is a special kind of equivariance where $\mathcal{L}'_g$ is the identity transformation for all element g of the group G.

A Lie group is a continuous group whose group elements are parameterized by a finite-dimensional differentiable manifold. In essence, this means that a Lie group is a group to which we can apply differential geometry. From now on, we also assume the group manifold is equipped with a Riemannian metric, that is it is Riemannian manifold. To each element g of G, we can associate a tangent space $T_g(G)$, which is spanned by a basis of left-invariant vectors denoted by $\{\mathcal{A}_i|_g\}_{i=1}^n$. We write $T_g(G) = \mathrm{span}\{\mathcal{A}_1|_g, \ldots, \mathcal{A}_n|_g\}$ and we can express tangent vectors $\dot{\gamma}(t)$ of $T_g(G)$ in this basis via $\dot{\gamma}(t) = c^i(t)\mathcal{A}_i|_{\gamma(t)}$.

Moreover, because this frame of basis vector is left-invariant, the coefficients $c_i(t)$ remain unchanged if the curve is moved by applying left group product. The tangent space at the origin $T_e(G)$ is spanned by a basis $\{A_i\}_{i=1}^n$ where $A_i = \mathcal{A}_i|_e$. We make a subtle difference in notation between $\mathcal{A}$ and A, where $\mathcal{A}$ represents a whole vector field and $\mathcal{A}|_g$ represents the vector at location g. The straight A is used to indicate a vector in the Lie algebra, the tangent space at the origin. Via the push-forward $(L_g)_*$, we can generate a whole vector space by picking a vector in the Lie algebra and transporting tangent vectors from $\dot{\gamma}(0) \in T_e(G)$ to $g \cdot \dot{\gamma}(0) \in T_g(G)$ using:

$$\mathcal{A}_i|_g = (L_g)_* A_i, \qquad \forall g \in G. \tag{12}$$

The left-invariant frame of basis vectors could also be interpreted as the directional derivative of functions defined on the group G. At any element g, the value of the directional derivative of a function f defined on G can be computed using the push-forward operation:

$$\mathcal{A}_i|_g f = (L_g)_* \mathcal{A}_i|_e f := \mathcal{A}_i|_e (f \circ L_{g^{-1}}), \tag{13}$$

so $\mathcal{A}_i|_g f$ represents the directional derivative along the vector field at location g, which can be defined by translating the function back to the origin via $L_{g^{-1}}$ and compute the derivative at the origin along the direction specified by $\mathcal{A}_i|_e$. In the above $L_g h := gh$.

17

The Laplace-Beltrami operator, the generalization to Riemannian manifold of the Laplace operator[8] is defined below.

Definition A.2 (Laplace-Beltrami operator) *Let g be an element of the Lie group $(\mathcal{M}, \cdot)$. The Laplace-Beltrami operator on the Riemannian manifold $\mathcal{M}$ is defined as*

$$\Delta_{\mathcal{R}(g)} = \mathrm{div}(\nabla_{\mathcal{R}(g)}) \tag{14}$$

Using the left-invariant frame of basis vectors as directional differential operators, we can define the gradient, expressed as a vector relative to the basis $\mathcal{A}_i|_g$, of a function $f : G \to G$:

$$\nabla_{\mathcal{R}(g)} f = \mathbf{R}_g^{-1}(\mathcal{A}_1|_g f, \ldots, \mathcal{A}_n|_g f)^\top, \tag{15}$$

where $\mathbf{R}_g$ is the Riemannian metric tensor defined relative to the basis $\mathcal{A}_i$, that corrects for local scaling and shrinking of the manifold as measured by the metric tensor. The divergence of a vector field $F : G \to \mathbb{R}^n$ is the operator:

$$\mathrm{div}(F) = \sum_{i=1}^{n} \mathcal{A}_i|_g F_i. \tag{16}$$

The Laplace-Beltrami operator depends on a Riemannian metric tensor, which describes how lengths of vectors should be measured in different directions, and in the Laplace-Beltrami operator, it rescales the derivates accordingly. While the usual Laplacian is isotropic (derivatives are threated the same in each direction), the Laplace-Beltrami operator can be anisotropic due to the Riemannian metric that is used.

Theorem A.1 (Left-invariance of the Laplace-Beltrami operator) *The Laplace-Beltrami operator $\Delta_{\mathcal{R}(g)}$ is left-invariant and satisfies:*

$$\Delta_{\mathcal{R}(g)} = (L_g)_* \Delta_{\mathcal{R}(e)}. \tag{17}$$

A Lie algebra $\mathfrak{g}$ is a vector space (here the tangent space at the identity element $T_e(G)$) that is endowed with a binary operator called the Lie bracket or commutator $[\cdot, \cdot] : T_e(G) \times T_e(G) \to T_e(G)$ that is bilinear, alternative and satisfies the Jacobi identity. Conceptually, the Lie bracket $[A, B]$ of two vector fields A, B is the derivative of B along the flow generated by A.[9] In the following, we take $\{A_i\}_{i=1}^n$ with $A_i = \mathcal{A}_i|_e$ as Lie Algebra and get left-invariant vector fields via the push-forward operation.

The Lie group exponential and logarithmic maps define the mappings between the group and the tangent space. The exponential map on a Lie group can be thought of as picking a vector A in the Lie algebra, construct a left-invariant vector field $\mathcal{A}$ via the push-forward operator, and follow this vector field by taking infinitesimal steps along the direction indicated by the vector field. Integrating along the vector field defined by A for unit time brings to some point $g \in G$. So $g = \exp A$, where $\exp : \mathfrak{g} \to G$.[10]

---

[8]Another equivalent definition in local coordinates exists too. It explicitly takes into account the Riemannian metric tensor field $\mathcal{R}$ previously defined. Here we use a more compact definition but using the subscript $\mathcal{R}(g)$ to indicate that the Laplace-Beltrami operator depends on the location on the Riemannian manifold and the Riemannian metric tensor field.

[9]Two vector fields are commutative (zero Lie bracket) if and only if its flows are too, in the sense that there is no difference starting at one point p, traveling a time t_a over the flow of A and then a time t_b over the flow of B, or, instead, traveling first t_b over the flow of B and then t_a over the flow of A.

[10]Note that Riemannian and Lie group exponential maps are different. For the general Riemannian exponential map, the vector field along which we compute the path integral is defined by the Riemannian metric. In the Lie group exponential map, the vector field is defined by the push forward of left-multiplication. The curves defined by the path integrals in the Riemannian exponential are known to be geodesics. The exponential curves in the Lie group case are "straight curves" with respect to a moving frame of reference but are not necessarily geodesics. The Riemannian and Lie group exponential maps only coincide when the Riemannian metric is both left and right invariant (see Section 4.5 in Bekkers [2017]).

A.3 Graph theory

Graphs are generic data representation forms that are useful for describing the geometric structure of data domains [West et al., 1996]. More formally, we define a graph $\mathcal{G}$ as a structure modeling a finite set of interactions called edges $\mathcal{E}$ between a finite set of objects called vertices $\mathcal{V}$. In the following, we denote by $|\mathcal{V}|$ the number of vertices and by $|\mathcal{E}|$ the number of edges. We denote v_i the i-th vertex and $e(v_i, v_j)$ the potential edge from v_i to v_j. We call the neighborhood of the vertex v_i the set of vertices connected to v_i by an edge and denote it by $\mathcal{N}(v_i)$. More generally, we write $\mathcal{N}^k(v_i)$ for the k-hops neighborhood of vertex v_i, that is, the set of vertices connected to v_i with a path of at most k edges. In some cases, it can be useful to add weights on graph edges. In general, the weights can take any value. Nevertheless, in this thesis, we assume weights in the range $[0, 1)$ and measuring the similarity rather than the distance. When the edges' weights are not naturally defined by an application, a common way to define them is to apply a kernel $\mathfrak{K} : \mathbb{R}_+ \to [0, 1)$ on the distance between connected vertices. For theoretical convergence results, we use a Gaussian weighting scheme:

$$w(v_i, v_j) = \begin{cases} \exp\left(-\frac{d^2(v_i, v_j)}{4t}\right) & \text{if } e(v_i, v_j) \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases}, \tag{18}$$

where $d(v_i, v_j)$ denotes the (Riemannian) distance between vertices v_i and v_j and t is a positive real number called bandwidth of the Gaussian kernel. From now on, we furthermore assume that the graphs are undirected and without self-loop.

The field of signal processing on graphs merges algebraic and spectral graph theoretic concepts with computational harmonic analysis to process such signals on graphs [Shuman et al., 2013]. A signal on a graph is a function $f : \mathcal{V} \to \mathbb{R}^d$, mapping each vertex of the graph to a d-dimensional real valued vector. In matrix form, this signal is a $|\mathcal{V}| \times d$ real valued matrix $\boldsymbol{f}$ whose rows are given by $\boldsymbol{f}_i = f(v_i) \in \mathbb{R}^d$.

An essential object in graph signal processing is the Laplacian operator. Under some specific conditions that we will state later, it can be interpreted as a discrete version of the Laplace-Beltrami operator.

Definition A.3 (Symmetric normalized Laplacian) *Let $\mathcal{G}$ be an undirected weighted graph without self-loops. The symmetric normalized Laplacian $\boldsymbol{\Delta}$ is the $|\mathcal{V}| \times |\mathcal{V}|$ real valued matrix whose components are given by:*

$$\Delta_{i,j} = \begin{cases} 1 & \text{if } i = j \text{ and } \deg(v_i) > 0 \\ -\frac{w(v_i, v_j)}{\sqrt{\deg(v_i)\deg(v_j)}} & \text{if } i \neq j \text{ and } \deg(v_i) > 0 \\ 0 & \text{otherwise} \end{cases}. \tag{19}$$

Assuming a function f defined on the graph vertices $\mathcal{V}$, by inspection on the components of $\boldsymbol{\Delta}^k \boldsymbol{f}$, we remark that the Laplacian acts as a k-op neighborhood operator. We can prove that $\boldsymbol{\Delta}$ is a symmetric positive definite matrix. Hence it admits a unique eigendecomposition of the form $\boldsymbol{\Delta} = \boldsymbol{\Phi}\boldsymbol{\Lambda}\boldsymbol{\Phi}^\top$ where the j-th column of Φ correponds to the eigenvector ϕ_j associated with real positive eigenvalue λ_j. In these settings, the eigenvalues are in the range $[0, 2]$ [Chung and Graham, 1997]. By analogy with the Euclidean case where the Laplacian's eigenfunctions correspond to the Fourier basis, we can construct a graph Fourier basis from the eigendecomposition of the graph Laplacian, and define the graph Fourier transform and its inverse.[11]

Definition A.4 (Graph Fourier transform) *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph with Laplacian $\boldsymbol{\Delta}$ and let $f : \mathcal{V} \to \mathbb{R}$ be a signal defined on the graph's vertices. The graph Fourier transform $\hat{f}$ of f is given by:*

$$\hat{f}(\lambda_j) = \hat{f}_j = \{\boldsymbol{\Phi}^\top \boldsymbol{f}\}_j = \sum_i \phi_{ij} f_i = \sum_{v_i} \phi(v_i, \lambda_j) f(v_i), \tag{20}$$

and its inverse transform by:

$$f(v_i) = f_i = \{\boldsymbol{\Phi}\hat{\boldsymbol{f}}\}_i = \sum_j \phi_{ij} \hat{f}_j = \sum_{\lambda_j} \phi(v_i, \lambda_j) \hat{f}(\lambda_j). \tag{21}$$

---

[11]The existence of an inverse transform is a direct consequence of the orthonormality of the eigenvectors.

B Lie groups

At the moment, all the group's elements we are interested in are in 1-1 correspondence with elements of the 3-dimensional general linear group $GL(3)$, the group of 3×3 real matrices. We will use this representation since working with matrices is more convenient. Indeed, in this case the group product, group inverse, group exponential and group logarithm respectively coincide with the matrix product, matrix inverse, matrix exponential and matrix logarithm.

For the roto-translation group $SE(2)$, the spatial part corresponds to the planar translations and the orientation part to the rotation angles. Group elements $g \in SE(2)$ are given in matrix formulation by:

$$g = (x, y, \theta) \leftrightarrow \boldsymbol{G_g} = \begin{pmatrix} \cos\theta & -\sin\theta & x \\ \sin\theta & \cos\theta & y \\ 0 & 0 & 1 \end{pmatrix}, \tag{22}$$

hence the group is 3 dimensional with two parameters x, y coming from the spatial space $\mathbb{R}^2$ and one orientation parameter coming from $[-\pi, \pi)$.

The Lie algebra $\mathfrak{se}(2)$ is the set of 3×3 matrices:

$$\boldsymbol{A_1} = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \boldsymbol{A_2} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \quad \text{and} \quad \boldsymbol{A_3} = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \tag{23}$$

Naturally, it is impossible to uniformly sample elements on this group because the $\mathbb{R}^2$ space is infinite. Adding boundaries to the euclidean space, we get the $[0, 1)^2$ space, which is not homogeneous anymore, but on which it is possible to uniformly sample $|\mathcal{V}_s|$ elements. Because we are considering anisotropic Laplace-Beltrami operators which are symmetric under reflections, it is sufficient to sample the rotation angles θ in the range $[-\pi/2, \pi/2)$.

As defined by its Lie algebra, the matrix logarithm related to an element of the $SE(2)$ group has the form:

$$\log_e \boldsymbol{G_g} = \begin{pmatrix} 0 & -c_3 & c_1 \\ c_3 & 0 & c_2 \\ 0 & 0 & 0 \end{pmatrix} \tag{24}$$

It results that the logarithmic map $\log_e : SE(2) \to \mathfrak{se}(2)$ admits a closed form expression $\log_e g = (c_1, c_2, c_3)^\top$ with[12]:

$$c_1 = \frac{\theta}{2}\left(y + x \cot\frac{\theta}{2} \right), \quad c_2 = \frac{\theta}{2}\left(-x + y \cot\frac{\theta}{2} \right), \quad \text{and} \quad c_3 = \theta. \tag{25}$$

The group $SO(3)$ of all rotations about the origin of 3-dimensional Euclidean space can be split into a "spatial" part which is the sphere, and a rotation part, which is a rotation around a particular reference axis. Using ZYZ representation, group elements $g \in SO(3)$ are given in matrix formulation by:

$$g = (\alpha, \beta, \gamma) \leftrightarrow \boldsymbol{G_g} = \boldsymbol{R}_{\gamma,z} \boldsymbol{R}_{\beta,y} \boldsymbol{R}_{\alpha,z}, \tag{26}$$

where $\alpha \in [-\pi, \pi]$, $\beta \in [-\pi/2, \pi/2]$ and $\gamma \in [-\pi, \pi]$. We view the sphere S^2 as the spatial part, just like we view our Earth as locally flat. S^2 is parametrized with Euler angles β and γ, which are independent of α. The rotation part is then parametrized by α.

The Lie algebra $\mathfrak{so}(3)$ is the set of antisymmetric 3×3 matrices:

$$\boldsymbol{A_1} = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \quad \boldsymbol{A_2} = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \text{and} \quad \boldsymbol{A_3} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix}. \tag{27}$$

To uniformly sample on the group $SO(3)$ is quite challenging because a perfectly uniform sampling on the sphere does not exist. Fortunately, this task has been largely studied, and many algorithms

---

[12]In the isotropic case, corresponding to the 2-d Euclidean space, one can check that using the exact same expression of the logarithmic map with $\theta = 0$ gives the Euclidean distance.

have been proposed: equiangular[13] [Driscoll and Healy, 1994], HEALPix [Gorski et al., 2005], and icosahedral [Baumgardner and Frederickson, 1985] samplings. As before, due to the same symmetry argument, it is sufficient to restrict the orientation range to $[-\pi/2, \pi/2]$ and to sample $|\mathcal{V}_o|$ elements on this restricted range.

The matrix logarithm related to an element of the $SO(3)$ group has the form:

$$\log_e \boldsymbol{G}_g = \begin{pmatrix} 0 & -c_3 & c_2 \\ c_3 & 0 & -c_1 \\ -c_2 & c_1 & 0 \end{pmatrix}. \tag{28}$$

Using the Rodrigues' rotation formula [Rodrigues, 1840], we derive a closed form expression for the logarithmic map $\log_e : SO(3) \to \mathfrak{so}(3)$, with notation $\log_e g = (c_1, c_2, c_3)^\top$. We get[14]:

$$c_1 = \frac{\theta}{2\sin\theta}(\boldsymbol{G}_{g,3,2} - \boldsymbol{G}_{g,2,3}), \quad c_2 = \frac{\theta}{2\sin\theta}(\boldsymbol{G}_{g,2,1} - \boldsymbol{G}_{g,1,2}), \text{ and } \quad c_3 = \frac{\theta}{2\sin\theta}(\boldsymbol{G}_{g,1,3} - \boldsymbol{G}_{g,3,1}). \tag{29}$$

Due to the π periodicity of the orientation axis, the π-periodic Riemannian distance between two elements of the groups is the minimal distance between:

- the original distance without offset;
- the original distance with a negative offset $-\pi$ on the orientation axis;
- the original distance with a positive offset $+\pi$ on the orientation axis.

C Experiment details

Empirical evidence showed that deep and wide neural networks are keys to good performances. To help going deeper we use residual convolutional layers [He et al., 2016] and batch normalization [Ioffe and Szegedy, 2015]. To avoid tuning of the learning rate, we use ADAM optimizers [Kingma and Ba, 2014] for all our experiments. Last, we initialise our models' parameters via the Kaiming method [He et al., 2015].

Stability test on MNIST. In these experiments, we used a Wide Residual ChebNet with three convolutional residual layers with rectified linear units and kernel of size 4. As we do not pool at all, we only define one $SE(2)$ 16-NN graph with $28 \times 28 \times 6 = 4704$ vertices. We use extreme spatial anisotropy with $\epsilon^2 = 0.1$ and reach the 40/60 ratio by setting ξ accordingly. The residual layers are followed by a projective global max pooling layer and a fully connected with LogSoftMax output layer.

Orientation anisotropic test on CIFAR10. In these experiments, we used a Wide Residual Cheb-Net with three residual layers with rectified linear units and kernel of size 4. We use two R2RandPool layers. Hence we define three $SE(2)$ 16-NN graphs with $32 \times 32 \times 6 = 6144$, $16 \times 16 \times 6 = 1536$, and $8 \times 8 \times 6 = 384$ vertices. We use extreme spatial anisotropy with $\epsilon^2 = 0.1$. The residual layers are followed by a projective global max pooling layer and a fully connected with LogSoftMax output layer.

Spatial anisotropic test on CIFAR10. In these experiments, we used a Wide Residual ChebNet with 3 residual layers with rectified linear units and kernel of size 4. We use two R2RandPool layers. Hence we define three $SE(2)$ 16-NN graphs with $32 \times 32 \times 6 = 6144$, $16 \times 16 \times 6 = 1536$, and $8 \times 8 \times 6 = 384$ vertices. We use moderate orientation anisotropies to reach the 40/60 ratio on each graphs by setting ξ accordingly. The isotropic case is constructed using $\epsilon^2 = \xi^2 = 1$ on three $\mathbb{R}^2$

---

[13]Equiangular is far from uniform, but it has a sampling theorem.

[14]To compute the logarithmic map in the isotropic case corresponding to S^2, we must use a slightly modified version such that $c_3 = 0$, which via the exponential map generate torsion free exponential curves. The logarithmic of any rotation matrix defined by $(-\gamma, \beta, \gamma)$ yields this results [Bekkers, 2019, Portegies et al., 2015].

8-NN graphs with $32 \times 32 \times 1 = 784$, $16 \times 16 \times 1 = 256$, and $8 \times 8 \times 1 = 64$ vertices. The residual layers are followed by a projective global max pooling layer and a fully connected with LogSoftMax output layer.

Scalability test on STL10. In these experiments, we used a Wide Residual ChebNet with three residual layers with rectified linear units and kernel of size 4. We use two R2RandPool layers. Hence we define three $SE(2)$ 16-NN graphs with $96 \times 96 \times 6 = 55296$, $48 \times 48 \times 6 = 13824$, and $8 \times 8 \times 6 = 3456$ vertices. We use extreme spatial anisotropy with $\epsilon^2 = 0.1$ and reach the 40/60 ratio on each graphs by setting ξ accordingly. The isotropic case is constructed using $\epsilon = \xi = 1$ on three $\mathbb{R}^2$ 8-NN graphs with $96 \times 96 \times 1 = 9216$, $48 \times 48 \times 1 = 2304$, and $24 \times 24 \times 1 = 576$ vertices. The residual layers are followed by a projective global max pooling layer and a fully connected with LogSoftMax output layer. The residual layers are followed by a projective global max pooling layer and a fully connected with LogSoftMax output layer.

Scalability test on ClimateNet. In these experiments, we used a U-ChebNet with three residual layers with rectified linear units and kernel of size 3. We use five S2MaxPool (encoding) and five S2AvgUnpool (decoding) layers. Hence we define six $SO(3)$ 16-NN graphs with $10242 \times 6 = 61452$, $2562 \times 6 = 15372$, $642 \times 6 = 3852$, $162 \times 6 = 972$, $42 \times 6 = 252$, and $12 \times 6 = 72$ vertices. We use extreme spatial anisotropy with $\epsilon^2 = 0.1$ and reach the 40/60 ratio on each graphs by setting ξ accordingly. The residual layers are followed by a projective global max pooling layer and a fully connected with LogSoftMax output layer.

D Stability under random perturbations

In sensitive domains, the stability of a neural network to random perturbation is a desired property. We aim at demonstrate the high stability of our approach, as random perturbations are added to graphs during the training. We introduce two methods, both consisting in randomly pruning the original graph to construct a random sub-graph (see Figure 6). At test time, random perturbations are removed in order to evaluate a perturbation-free model.


(a) One constructs a random sub-graph by randomly sampling a rate $\kappa_{\mathcal{E}}$ of edges of the complete graph based on their weight. As a direct consequence of the work of Keriven et al. [2020], the Laplacian of a quasi-sparse graph with randomly sampled edges remains consistent. Then, the unsampled edges are pruned by setting their weight to zero. One must be aware that this process could lead to isolated vertices or cluster of vertices. It is not necessarily a problem, but it is important to be mindful of this potential effect.

(b) This method based on vertices sampling consists of creating a random sub-graph by randomly sampling a rate $\kappa_{\mathcal{V}}$ of vertices from the complete graph. The edges between one or mode unsampled vertices are discarded too. This method is more drastic and should be used carefully, as such a vertices pruning leads to a violation of the uniform distribution of vertices.

Figure 6: Edge- and vertex-based sampling methods.

We run a bunch of experiments with a Wide Residual architecture on MNIST [LeCun and Cortes, 2010], varying the rates of edges or vertices to sample. The objective of these experiments is twofold. First, we would like to test the stability of the model under random perturbations. Second, this experiment is also a good manner to demonstrate the equivariance property of our method. In this purpose, we train our model on the original training set of MNIST, adding some perturbations. At test time, we evaluate the perturbation-free model on the test set with and without random rotations.

(a) Test-accuracy against rate of edges to sample   (b) Test-accuracy against rate of vertices to sample

Figure 7: Stability and group equivariance under random perturbations.

Equivariance. With these experiments, we empirically demonstrated the rotation-equivariance of a ChebLieNet. An extrinsic equivariance error can explain the slight difference in performance between the original test set and the randomly rotated one. This error is not related to the model but the evaluation's method and can be semantic (e.g. is it always possible to differentiate a six from a nine without orientation information) or numerical (e.g., how to rotate a low-resolution image by 10 degrees without alteration).

Robustness. While in this specific case, our dropout-like method does not help to improve our model on the test set, we expect it could help to reduce the over-fitting on other tasks. In addition, one could notice that even at middle-low sampling rates, the models remain stable. First of all, it supports the idea that there is no point using a fully connected graph; more edges capture the geometry better but with diminishing returns. Secondly, the model is robustly able to catch the semantic information in an image, even with partial alterations.

E   Empirical convergence in eigenmaps of the graph Laplacians

As it has been proven by Belkin and Niyogi [2006], assuming a suitably constructed graph, the graph Laplacian converges in eigenmaps to its continuous counterpart, the Laplace-Beltrami operator. Hence, a good sanity check is to compare the eigenmaps of the discrete and continuous Laplace operators.

Firstly, we recall the eigenvalues of symmetric normalized graph Laplacians satisfies $0 \leq \lambda_k \leq 2$ for all k. As the eigenvalues can be interpreted as frequency components, the eigenvector associated to an high eigenvalue would correspond to an high frequency signal on the graph. To finish with, we propose a detailed analysis of the eigenmaps of our spaces of interest in terms of the Fourier basis and the spherical harmonics.

Isotropic 2-dimensional grid. The eigenmaps of the $[0, 1)^2$ space are shown in figure 8. It is a well-known fact that the 2-dimensional grid with periodic boundaries can be spanned by the Fourier basis. Consequently, the eigenvectors of this space corresponds to the sine and cosine trigonometric functions. In this case, the eigenvalues have multiplicity two (excepted the first one associated with a constant eigenvector), one for the sine and the other for the cosine. When the periodic condition at the boundaries are relaxed, the space is not homogeneous anymore, and the symmetries of the space change a little bit. The $[0, 1)^2$ space is such a space where periodic condition does not hold. Hence, the eigenmaps of this space are not equal to the Fourier basis, but very close.

Anisotropic 2-dimensional grid. The eigenmaps of the $[0, 1)^2 \times [-\pi/2, \pi/2)$ space are shown in figure 10. We can extend the discussion we add in the isotropic case to justify the fact that the eigenmaps of the $[0, 1)^2 \times [-\pi/2, \pi/2)$ is close to the Fourier basis. Moreover, we an additional orientation dimension, the rotation of the $[0, 1)^2$ induces a torsion all along the orientation axis because of the anisotropic Riemannian metric.

Isotropic 2-dimensional sphere. The eigenmaps of the S^2 space are shown in figure 9. It is a well-known fact that the 2-dimensional sphere is spanned by the spherical harmonics. The eigenvalues

associated to this eigenvectors have increasing multiplicities of the form $2m + 1$ where m is the order of the spherical harmonic.

Anisotropic 2-dimensional sphere. The eigenmaps of the $S^2 \times [-\pi/2, \pi/2)$ space are shown in figure 11. We can extend the discussion we add in the isotropic case to justify the fact that the eigenmaps of the $S^2 \times [-\pi/2, \pi/2)$ are close to the spherical harmonics. Nevertheless, the role of the orientation dimension is not easily interpretable. Indeed the notion of orientation is arduous to understand, as the kernel's orientation depends on the path, and not on this orientation location.

(a) Eigenvalues of the $[0,1]^2$ space, from λ_0 to λ_{49}.

(b) Eigenvectors of the $[0,1]^2$ space, from ϕ_0 (left) to ϕ_9 (right).

Figure 8: Eigenmaps of the $[0,1]^2$ space.

(a) Eigenvalues of the S^2 space, from λ_0 to λ_{49}.

(b) Eigenvectors of the S^2 space, from ϕ_0 (left) to ϕ_9 (right).

Figure 9: Eigenmaps of the S^2 space.

(a) Eigenvalues of the $[0, 1]^2 \times [-\pi/2, \pi/2)$ space, from λ_0 to λ_{49}.
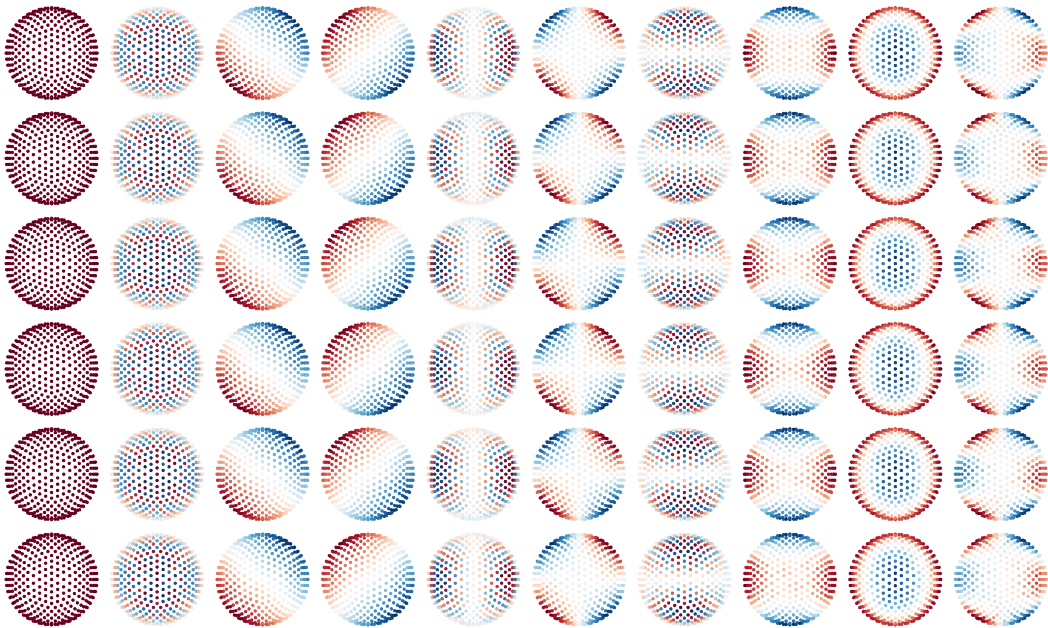

(b) Eigenvectors of the $[0, 1]^2 \times [-\pi/2, \pi/2)$ space, from ϕ_0 (left) to ϕ_9 (right).

Figure 10: Eigenmaps of the $[0, 1]^2 \times [-\pi/2, \pi/2)$ space.

(a) Eigenvalues of the $S^2 \times [-\pi/2, \pi/2)$ space, from λ_0 to λ_{49}.

(b) Eigenvectors of the $S^2 \times [-\pi/2, \pi/2)$ space, from ϕ_0 (left) to ϕ_9 (right).

Figure 11: Eigenmaps of the $S^2 \times [-\pi/2, \pi/2)$ space.