# PDE-Based Group Equivariant Convolutional Neural Networks 

Bart M. N. Smets ${ }^{1}$ (D) Jim Portegies ${ }^{2}$ •Erik J. Bekkers ${ }^{3}$ •Remco Duits ${ }^{1}$

Received: 9 July 2021 / Accepted: 3 July 2022 / Published online: 27 July 2022
© The Author(s) 2022


#### Abstract

We present a PDE-based framework that generalizes Group equivariant Convolutional Neural Networks (G-CNNs). In this framework, a network layer is seen as a set of PDE-solvers where geometrically meaningful PDE-coefficients become the layer's trainable weights. Formulating our PDEs on homogeneous spaces allows these networks to be designed with built-in symmetries such as rotation in addition to the standard translation equivariance of CNNs. Having all the desired symmetries included in the design obviates the need to include them by means of costly techniques such as data augmentation. We will discuss our PDE-based G-CNNs (PDE-G-CNNs) in a general homogeneous space setting while also going into the specifics of our primary case of interest: roto-translation equivariance. We solve the PDE of interest by a combination of linear group convolutions and nonlinear morphological group convolutions with analytic kernel approximations that we underpin with formal theorems. Our kernel approximations allow for fast GPU-implementation of the PDE-solvers; we release our implementation with this article in the form of the LieTorch extension to PyTorch, available at https://gitlab.com/bsmetsjr/ lietorch. Just like for linear convolution, a morphological convolution is specified by a kernel that we train in our PDE-GCNNs. In PDE-G-CNNs, we do not use non-linearities such as max/min-pooling and ReLUs as they are already subsumed by morphological convolutions. We present a set of experiments to demonstrate the strength of the proposed PDE-G-CNNs in increasing the performance of deep learning-based imaging applications with far fewer parameters than traditional CNNs.


Keywords PDE • Group equivariance • Deep learning • Morphological scale-space

## 1 Introduction

In this work, we introduce PDE-based Group CNNs. The key idea is to replace the typical trifecta of convolution, pooling and ReLUs found in CNNs with a Hamilton-Jacobi type evolution PDE, or more accurately a solver for a HamiltonJacobi type PDE. This substitution is illustrated in Fig. 1 where we retain (channel-wise) affine combinations as the means of composing feature maps.

[^0]The PDE we propose to use in this setting comes from the geometric image analysis world [1-12]. It was chosen based on the fact that it exhibits similar behavior on images as traditional CNNs do through convolution, pooling and ReLUs. Additionally it can be formulated on Lie groups to yield equivariant processing, which makes our PDE approach compatible with Group CNNs [13-25]. Finally, an approximate solver for our PDE can be efficiently implemented on modern highly parallel hardware, making the choice a practical one as well.

Our solver uses the operator splitting method to solve the PDE in a sequence of steps, each step corresponding to a term of the PDE. The sequence of steps for our PDE is illustrated in Fig. 2. The morphological convolutions that are used to solve for the nonlinear terms of the PDE are a key aspect of our design. Normally, morphological convolutions are considered on $\mathbb{R}^{d}$ [26,27], but when extended to Lie groups such as $S E(d)$ they have many benefits in applications (e.g., crossing-preserving flow [28] or tracking [29,30]). Using morphological convolutions allows our net-


Fig. 1 In a PDE-based CNN, we replace the traditional convolution, pooling and ReLU operations by a PDE solver. The inputs of a given layer serve as initial conditions for a set of evolution PDEs; the outputs consist of affine combinations of the solutions of those PDEs at a fixed point in time. The parameters of the PDE become the trainable weights (alongside the affine parameters) over which we optimize


Fig. 2 Our Hamilton-Jacobi type PDE of choice contains a convection, diffusion, dilation and erosion term (CDDE for short). Through operator splitting, we solve for these terms separately by using resampling (for convection), linear convolution (for diffusion) and morphological convolution (for dilation and erosion)
work to have trainable non-linearities instead of the fixed non-linearities in (G-)CNNs.

The theoretical contribution of this paper consists of providing good analytical approximations to the kernels that go in the linear and morphological convolutions that solve our PDE. On $\mathbb{R}^{n}$, the formulation of these kernels is reasonably straightforward, but in order to achieve group equivariance we need to generalize them on homogeneous spaces.

Instead of training kernel weights, our goal is training the coefficients of the PDE. The coefficients of our PDE have the benefit of yielding geometrically meaningful parameters from a image analysis point of view. Additionally, we will need (much) less PDE parameters than kernel weights to achieve a given level of performance in image segmentation and classification tasks, arguably the greatest benefit of our approach.

This paper is a substantially extended journal version of [31] presented at the SSVM 2021 conference.

### 1.1 Structure of the Article

The structure of the article is as follows. We first place our work in its mathematical and deep learning context in Sect. 2. Then, we introduce the needed theoretical preliminaries from Lie group theory in Sect. 3 where we also define the space of positions and orientations $\mathbb{M}_{d}$ that will allow us to construct roto-translation equivariant networks.

In Sect. 4, we give the overall architecture of a PDE-GCNN and the ancillary tools that are needed to support a PDE-G-CNN. We propose an equivariant PDE that models commonly used operations in CNNs.

In Sect. 5, we detail how our PDE of interest can be solved using a process called operator splitting. Additionally, we give tangible approximations to the fundamental solutions (kernels) of the PDEs that are both easy to compute and sufficiently accurate for practical applications. We use them extensively in the PDE-G-CNNs GPU-implementations in PyTorch that can be downloaded from the GIT-repository: https://gitlab.com/bsmetsjr/lietorch.

Section 6 is dedicated to showing how common CNN operations such as convolution, max-pooling, ReLUs and skip connections can be interpreted in terms of PDEs.

We end our paper with some experiments showing the strength of PDE-G-CNNs in Sect. 7 and concluding remarks in Sect. 8.

The framework we propose covers transformations and CNNs on homogeneous spaces in general and as such we develop the theory in an abstract fashion. To maintain a bridge with practical applications, we give details throughout the article on what form the abstractions take explicitly in the case of roto-translation equivariant networks acting on $\mathbb{M}_{d}$, specifically in 2D (i.e., $d=2$ ).

## 2 Context

As this article touches on disparate fields of study, we use this section to discuss context and highlight some closely related work.

### 2.1 Drawing Inspiration from PDE-Based Image Analysis

Since the Partial Differential Equations that we use are well known in the context of geometric image analysis [1-11], the layers also get an interpretation in terms of classical imageprocessing operators. This allows intuition and techniques from geometric PDE-based image analysis to be carried over to neural networks.

In geometric PDE-based image processing, it can be beneficial to include mean curvature or other geometric flows [32-35] as regularization and our framework provides a nat-
ural way for such flows to be included into neural networks. In the PDE layer from Fig. 2, we only mention diffusion as a means of regularization, but mean curvature flow could easily be integrated by replacing the diffusion sub-layer with a mean curvature flow sub-layer. This would require replacing the linear convolution for diffusion by a median filtering approximation of mean curvature flow [1].

### 2.2 The Need for Lifting Images

In geometric image analysis, it is often useful to lift images from a 2D picture to a 3D orientation score as in Fig. 3 and do further processing on the orientation scores [36]. A typical image processing task in which such a lift is beneficial is that of the segmentation of blood vessels in a medical image. Algorithms based on processing the 2D picture directly usually fail around points where two blood vessels cross, but algorithms that lift the image to an orientation score manage to decouple the blood vessels with different orientations as is illustrated in the bottom row of Fig. 3.

To be able to endow image processing neural networks with the added capabilities (such as decoupling orientations and guaranteeing equivariance) that result from lifting data to an extended domain, we develop our theory for the more general CNNs defined on homogeneous spaces, rather than just the prevalent CNNs defined on Euclidean space. One can then choose which homogeneous space to use based on the needs of one's application (such as needing to decouple orientations). A homogeneous space is, given subgroup $H$ of a group $G$, the manifold of left cosets, denoted by $G / H$. In the above image analysis example, the group $G$ would be the special Euclidean group $G=S E(d)$, the subgroup $H$ would be the stabilizer subgroup of a fixed reference axis, and the corresponding homogeneous space $G / H$ would be the space of positions and orientations $\mathbb{M}_{d} \equiv \mathbb{R}^{d} \times S^{d-1}$, which is the lowest dimensional homogeneous space able to decouple orientations. By considering convolutional neural networks on homogeneous spaces such as $\mathbb{M}_{d}$, these networks have access to the same benefits of decoupling structures with different orientations as was highly beneficial for geometric image processing [37-51].

Remark 1 (Generality of the architecture) Although not considered here, for other Lie groups applications (e.g., frequency scores [52,53], velocity scores, scale orientation scores [54]) the same structure applies; therefore, we keep our theory in the general setting of homogeneous spaces $G / H$. This generality was also important in non-PDE-based learning [22], but also for PDE-based learning it is again beneficial.


Fig. 3 Illustrating the process of lifting and projecting; in this case the advantage of lifting an image from $\mathbb{R}^{2}$ to the 2 D space of positions and orientations $\mathbb{M}_{2}$ derives from the disentanglement of the lines at the crossings


Fig. 4 Spatial CNNs, as used for image classification for example, are translation equivariant but not necessarily equivariant with respect to rotation, scaling and other transformations as the illustrative tags of the differently transformed apples images suggest. Building a G-CNN with the appropriately chosen group confers the network with all the equivariances appropriate for the chosen application. Our PDE-based approach is compatible with the group CNN approach [22] and so can confer the same symmetries

### 2.3 The Need for Equivariance

We require the layers of our network to be equivariant: a transformation of the input should lead to a corresponding transformation of the output, in other words: first transforming the input and then applying the network or first applying the network and then transforming the output should yield the same result. A particular example, in which the output transformation is trivial (i.e., the identity transformation), is that of invariance: in many classification tasks, such as the recognition of objects in pictures, an apple should still be recognized as an apple even if it is shifted or otherwise transformed in the picture as illustrated in Fig. 4. By guaranteeing equivariance of the network, the amount of data necessary or the need for data augmentation is reduced as the required symmetries are intrinsic to the network and need not be trained.

### 2.4 Related Work

$G-C N N s$ After the introduction of G-CNNs by Cohen and Welling [13] in the field of machine and deep learning, GCNNs became popular. This resulted in many articles on showing the benefits of G-CNNs over classical spatial CNNs. These works can be roughly categorized as

- discrete G-CNNs [13-17],
- regular continuous G-CNNs [19-21,29,55,56],
- and steerable continuous G-CNNs [22-25,57] that rely on Fourier transforms on homogeneous spaces $[50,58]$.

Both regular and steerable G-CNNs naturally arise from linear mappings between functions on homogeneous spaces that are placed under equivariance constraints [22,24,55,57]. Regular G-CNNs explicitly extend the domain and lift feature maps to a larger homogeneous space of a group, whereas steerable CNNs extend the co-domain by generating fiber bundles in which a steerable feature vector is assigned to each position in the base domain. Although steerable operators have clear benefits in terms of computational efficiency and accuracy [59,60], working with steerable representations puts constraints on nonlinear activations within the networks which limits the representation power of G-CNNs [57]. Like regular G-CNNs, the proposed PDE-G-CNNs do not suffer from this. In our proposed PDE-G-CNN framework, it is essential that we adopt the domain-extension viewpoint, as this allows to naturally and transparently construct scale space PDEs via left-invariant vector fields [12]. In general, this viewpoint entails that the domain of images is extended from the space of positions only, to a higher-dimensional homogeneous space, and originates from coherent state theory [61], orientation score theory [36], cortical perception models [39], G-CNNs [13,20] and rigid-body motion scattering [62].

The proposed PDE-G-CNNs form a new, unique class of equivariant neural networks, and we show in Sect. 6 how regular continuous G-CNNs arise as a special case of our PDE-G-CNNs.
Probabilistic-CNNs Our geometric PDEs relate to $\alpha$-stable Lévy processes [50] and cost-processes akin to [26], but then on $\mathbb{M}_{d}$ rather than $\mathbb{R}^{d}$. This relates to probabilistic equivariant numerical neural networks [63] that use anisotropic convection-diffusions on $\mathbb{R}^{d}$.

In contrast with these networks, the PDE-G-CNNs that we propose allow for simultaneous spatial and angular diffusion on $\mathbb{M}_{d}$. Furthermore, we include nonlinear Bellman processes [26] for max pooling over Riemannian balls.
KerCNNs An approach to introducing horizontal connectivity in CNNs that does not require a Lie group structure
was proposed by Montobbio et al. $[64,65]$ in the form of KerCNNs. In this biologically inspired metric model, a diffusion process is used to achieve intra-layer connectivity.

While our approach does require a Lie group structure, it is not restricted to diffusion and also includes dilation/ erosion.

Neural Networks and Differential Equations The connection between neural networks and differential equations became widely known in 2017, when Weinan [66] explicitly explained the connection between neural networks and dynamical systems, especially in the context of the ultradeep ResNet [67]. This point of view was further expanded by Lu et al. [68], showing how many ultradeep neural networks can be viewed as discretizations of ordinary differential equations. The somewhat opposite point of view was taken by Chen et al. [69], who introduced a new type of neural network which no longer has discrete layers, them being replaced by a field parameterized by a continuous time variable. Weinan E also indicated a relationship between CNNs and PDEs, or rather with evolution equations involving a nonlocal operator. Implicitly, the connection between neural networks and differential equations was also explored by the early works of Chen et al. [70] who learn parameters in a reactiondiffusion equation. This connection between neural networks and PDEs was then made explicit and more extensive by Long et al. who made it possible to learn a much wider class of PDEs [71] with their PDE-Net. More recent work in PDE inspired neural networks includes [72,73].

Basing neural network computations on PDEs formulated on manifolds also makes the processing independent with respect to the choice of coordinates on the manifold in the fashion of Weiler et al. [74].

More recent work in this direction includes integrating equivariant partial differential operators in steerable CNNs [75], drawing a strong analogy between deep learning and physics.

A useful aspect of the connection between neural networks and differential equations is the observation that the stability of the differential equation can give into the stability and generalization ability of the neural network [76]. Moreover, there are intriguing analogies with numerical PDE-approximations and specific network architectures (e.g. ResNets), as can be seen in the comprehensive overview article by Alt et al. [77].

The main contribution of our work in the field of PDErelated neural networks is that we implement and analyze geometric PDEs on homogeneous spaces, to obtain general group equivariant PDE-based CNNs whose implementations just require linear and morphological convolutions with new analytic approximations of scale space kernels.

## 3 Equivariance: Groups and Homogeneous Spaces

We want to design the PDE-G-CNN, and its layers, in such a way that they are equivariant. Equivariance essentially means that one can either transform the input and then feed it through the network, or first feed it through the network and then transform the output, and both give the same result. We will give a precise definition after introducing general notation.

### 3.1 The General Case

A layer in a neural network (or indeed the whole network) can be viewed as an operator from a space of real-valued functions defined on a space $X$ to a space of real-valued functions defined on a space $Y$. It may be helpful to think of these function spaces as spaces of images.

We assume that the possible transformations form a connected Lie group $G$. Think for instance of a group of translations which shift the domain into different directions. The Lie group being connected excludes transformations such as reflections, which we want to avoid for the sake of simplicity. We further assume that the Lie group $G$ acts smoothly on both spaces $X$ and $Y$, which means that there are smooth maps $\rho_{X}: G \times X \rightarrow X$ and $\rho_{Y}: G \times Y \rightarrow Y$ such that for all $g, h \in G$,
$\rho_{X}(g h, x)=\rho_{X}\left(g, \rho_{X}(h, x)\right)$
and
$\rho_{Y}(g h, x)=\rho_{Y}\left(g, \rho_{Y}(h, x)\right)$,
making $\rho_{X}$ and $\rho_{Y}$ group actions on their respective spaces.
Additionally we will assume that the group $G$ acts transitively on the spaces, meaning that for any two elements of these spaces there exists a transformation in $G$ that maps one to the other. This has as the consequence that $X$ and $Y$ can be seen as homogeneous spaces [78]. In particular, this means that after selecting a reference element $x_{0} \in X$ we can make the following isomorphism:
$X \equiv G / \operatorname{Stab}_{G}\left(x_{0}\right)$
using the mapping
$x \mapsto\left\{g \in G \mid \rho_{X}\left(g, x_{0}\right)=x\right\}$,
which is a bijection due to transitivity and the fact that

$$
\operatorname{Stab}_{G}\left(x_{0}\right):=\left\{g \in G \mid \rho_{X}\left(g, x_{0}\right)=x_{0}\right\}
$$

is a closed subgroup of $G$. Because of this, we will represent a homogeneous space as the quotient $G / H$ for some choice of closed subgroup $H=\operatorname{Stab}_{G}\left(x_{0}\right)$ since all homogeneous spaces are isomorphic to such a quotient by the above construction.

In this article, we will restrict ourselves to those homogeneous spaces that correspond to those quotients $G / H$ where the subgroup $H$ is compact and connected. Restricting ourselves to compact and connected subgroups simplifies many constructions and still covers several interesting cases such as the rigid body motion groups $S E(d)$.

The elements of the quotient $G / H$ consist of subsets of $G$ which we will denote by the letter $p$, these subsets are know as left cosets of $H$ since every one of them consists of the set $p=g H$ for some $g \in G$, the left cosets are a partition of $G$ under the equivalence relation
$g_{1} \sim g_{2} \Longleftrightarrow g_{1}^{-1} g_{2} \in H . \Longleftrightarrow g_{1} H=g_{2} H$.
Under this notation, the group $G$ consists of the disjoint union
$G=\coprod_{p \in G / H} p$.
The left coset that is associated with the reference element $x_{0} \in X$ is $H$, and for that reason we also alias it by $p_{0}:=H$ when we want to think of it as an atomic entity rather than a set in its own right.

We will denote quotient map from $G$ to $G / H$ with $\pi$ :
$\pi(g):=g p_{0}:=g H$.
Remark 2 (Principal homogeneous space) Observe that by choosing $H=\{e\}$ we get $G / H \equiv G$, i.e., the Lie group is a homogeneous space of itself. This is called the principal homogeneous space. In that case, the group action is equivalent to the group composition. The numerical experiments we perform in this paper are on the principal homogeneous space $\mathbb{R}^{2} \times S^{1}$ of $S E(2)$.

We will denote the group action/left-multiplication by an element $g \in G$ by the operator $L_{g}: G / H \rightarrow G / H$ given by
$L_{g} p:=g p$ for all $p \in G / H$.
In addition, we denote the left-regular representation of $G$ on functions $f$ defined on $G / H$ by $\mathcal{L}_{g}$ defined by

$$
\begin{equation*}
\left(\mathcal{L}_{g} f\right)(p):=f\left(g^{-1} p\right) \tag{6}
\end{equation*}
$$

A neural network layer is itself an operator (from functions on $G / H_{X}$ to functions on $G / H_{Y}$ ), and we require the operator
to be equivariant with respect to the actions on these function spaces.

Definition 1 (Equivariance) Let $G$ be a Lie group with homogeneous spaces $G / H_{X}$ and $G / H_{Y}$. Let $\Phi$ be an operator from functions (of some function class) on $G / H_{X}$ to functions on $G / H_{Y}$, then we say that $\Phi$ is equivariant with respect to $G$ if for all functions $f$ (of that class) we have that:
$\forall g \in G:\left(\Phi \circ \mathcal{L}_{g}\right) f=\left(\mathcal{L}_{g} \circ \Phi\right) f$,
or in words: the neural network commutes with transformations.

Most of the time we will have $H_{X}=H_{Y}$ in our proposed neural networks; only the initial lifting layer and the final projection layer will be between different homogeneous spaces, as we will see later on.

### 3.2 Vector and Metric Tensor Fields

The particular operators that we will base our framework on are vector and tensor fields, if these basic building blocks are equivariant then our processing will be equivariant. We explain what left invariance means for these objects next.

For $g \in G$ and $p \in G / H$, let $T_{p}(G / H)$ be the tangent space at point $p$, then the pushforward

$$
\left(L_{g}\right)_{*}: T_{p}(G / H) \rightarrow T_{g p}(G / H)
$$

of the group action $L_{g}$ is defined by the condition that for all smooth functions $f$ on $G / H$ and all $v \in T_{p}(G / H)$ we have that

$$
\begin{equation*}
\left(\left(L_{g}\right)_{*} v\right) f:=\boldsymbol{v}\left(f \circ L_{g}\right) \tag{8}
\end{equation*}
$$

Remark 3 (Tangent vectors as differential operators) Other than the usual geometric interpretation of tangent vectors as being the velocity vectors $\dot{\gamma}(t)$ tangent to some differentiable curve $\gamma: \mathbb{R} \rightarrow G / H$, we will simultaneously use them as differential operators acting on functions as we did in (8). This algebraic viewpoint defines the action of the tangent vector $\dot{\gamma}(t)$ on a differentiable function $f$ as
$\dot{\gamma}(t) f:=\left.\frac{\partial}{\partial s} f(\gamma(s))\right|_{s=t}$.
In the flat setting of $G=\left(\mathbb{R}^{d},+\right)$, where the tangent spaces are isomorphic to the base manifold $\mathbb{R}^{d}$, when we have a tangent vector $\boldsymbol{c} \in \mathbb{R}^{d}$ its application to a function is the familiar directional derivative:
$\boldsymbol{c} f=\boldsymbol{c} \cdot \nabla f=d f(\boldsymbol{c})$.

See [79, §2.1.1] for details on this double interpretation.
Vector fields that have the special property that the push forward $\left(L_{g}\right)_{*}$ maps them to themselves in the sense that
$\forall g \in G, \forall p \in G / H: v(p) f=\boldsymbol{v}(g p)\left[\mathcal{L}_{g} f\right]$,
for all differentiable functions $f$ and where $v: p \mapsto$ $T_{p}(G / H)$ is a vector field are referred to as $G$-invariant.

Definition 2 ( $G$-invariant vector field on a homogeneous space) A vector field $v$ on $G / H$ is invariant with respect to $G$ if it satisfies
$\forall g \in G, \forall p \in G / H: \boldsymbol{v}(g p)=\left(L_{g}\right)_{*} \boldsymbol{v}(p)$.
It is straightforward to check that (9) and (10) are equivalent and that these imply the following.

Corollary 3 (Properties of $G$-invariant vector fields) On $a$ homogeneous space $G / H$, a $G$-invariant vector field $v$ has the following properties:

1. it is fully determined by its value $\left.\boldsymbol{v}\right|_{H} \in T_{H}(G / H)$ in $H$,
2. $\forall h \in H:\left.\left(L_{h}\right)_{*} \boldsymbol{v}\right|_{H}=\left.\boldsymbol{v}\right|_{H}$.

We also introduce $G$-invariant metric tensor fields.
Definition 4 ( $G$-invariant metric tensor field on $G / H$ ) A $(0,2)$-tensor field $\mathcal{G}$ on $G / H$ is $G$-invariant if and only if

$$
\begin{align*}
& \forall g \in G, \forall p \in G / H, \forall \boldsymbol{v}, \boldsymbol{w} \in T_{p}(G / H): \\
& \left.\mathcal{G}\right|_{p}(\boldsymbol{v}, \boldsymbol{w})=\left.\mathcal{G}\right|_{g p}\left(\left(L_{g}\right)_{*} \boldsymbol{v},\left(L_{g}\right)_{*} \boldsymbol{w}\right) \tag{11}
\end{align*}
$$

Recall that $L_{g} p:=g p$ and so the push-forward $\left(L_{g}\right)_{*}$ maps tangent vector from $T_{p}(G / H)$ to $T_{g p}(G / H)$. Again it follows immediately from this definition that a $G$-invariant metric tensor field has similar properties as a $G$-invariant vector field.

Corollary 5 (Properties of $G$-invariant metric tensor fields) On a homogeneous space $G / H$, a $G$-invariant metric tensor field $\mathcal{G}$ has the following properties:

1. it is fully determined by its metric tensor $\left.\mathcal{G}\right|_{p_{0}}$ at $p_{0}=H$, 2. $\forall h \in H, \forall \boldsymbol{v}, \boldsymbol{w} \in T_{p_{0}}(G / H):\left.\mathcal{G}\right|_{p_{0}}(\boldsymbol{v}, \boldsymbol{w})=$ $\left.\mathcal{G}\right|_{p_{0}}\left(\left(L_{h}\right)_{*} \boldsymbol{v},\left(L_{h}\right)_{*} \boldsymbol{w}\right)$.

Or in words, the metric (tensor) has to be symmetric with respect to the subgroup $H$.

A (positive definite) metric tensor field yields a Riemannian metric in the usual manner, as we recall next.

Definition 6 (Metric on $G / H$ ) Let $p_{1}, p_{2} \in G / H$ then:

$$
\begin{aligned}
& d_{\mathcal{G}}\left(p_{1}, p_{2}\right):=d_{G / H, \mathcal{G}}\left(p_{1}, p_{2}\right):= \\
& \inf _{\substack{\beta \in \operatorname{Lip}([0,1], G / H) \\
\beta(0)=p_{1}, \beta(1)=p_{2}}} \int_{0}^{1} \sqrt{\left.\mathcal{G}\right|_{\beta(t)}(\dot{\beta}(t), \dot{\beta}(t))} \mathrm{d} t .
\end{aligned}
$$

As metrics and their smoothness play a role in our construction, we need to take into account where that smoothness fails.

Definition 7 The cut locus $\operatorname{cut}(p) \subset G / H$ or $\operatorname{cut}(g) \subset G$ is the set of points, respectively, group elements where the distance map from $p$ resp. $g$ is not smooth (excluding the point $p$ and group element $g$ themselves).

As long as we stay away from the cut locus, the infimum from Definition 6 gives a unique geodesic.

Being derived from a $G$-invariant tensor field gives the metric $d_{\mathcal{G}}$ the same symmetries.

Proposition 8 ( $G$-invariance of the metric on $G / H$ ) Let $p_{1}, p_{2} \in G / H$ away from each other's cut locus, then we have:
$\forall g \in G: d_{\mathcal{G}}\left(p_{1}, p_{2}\right)=d_{\mathcal{G}}\left(g p_{1}, g p_{2}\right)$.

Proof We observe that we can make a bijection from the set of Lipschitz curves between $p_{1}$ and $p_{2}$ and between $g p_{1}$ and $g p_{2}$ simply by left multiplication by $g$ one way and $g^{-1}$ the other way. Due to (11), multiplying a curve with a group element preserves its length; hence, if $\gamma:[0,1] \rightarrow G / H$ is the geodesic from $p_{1}$ to $p_{2}$, then $g \gamma$ is the geodesic from $g p_{1}$ to $g p_{2}$, both having the same length.

A metric tensor field on the homogeneous space has a natural counterpart on the group.

Definition 9 (Pseudometric tensor field on $G$ ) A $G$-invariant metric tensor field $\mathcal{G}$ on $G / H$ induces a (pullback) pseudometric tensor field $\tilde{\mathcal{G}}$ on $G$ that is left invariant:
$\tilde{\mathcal{G}}:=\pi^{*} \mathcal{G}$,
where $\pi^{*}$ is the pullback of the quotient map $\pi$ from (4). This is equivalent to saying that for all $\boldsymbol{v}, \boldsymbol{w} \in T_{g} G$ :

$$
\left.\tilde{\mathcal{G}}\right|_{g}(\boldsymbol{v}, \boldsymbol{w}):=\left.\mathcal{G}\right|_{\pi(g)}\left(\pi_{*} \boldsymbol{v}, \pi_{*} \boldsymbol{w}\right)
$$

where $\pi_{*}$ is the pushforward of $\pi$.
This tensor field $\tilde{\mathcal{G}}$ is left invariant by virtue of $\mathcal{G}$ being $G$ invariant. It is also degenerate in the direction of $H$ and so yields a seminorm on $T G$.

Definition 10 (Seminorm on $T G$ ) Let $v \in T_{g} G$. Then, the metric tensor field $\mathcal{G}$ on $G / H$ induces the following seminorm:

$$
\begin{equation*}
\|\boldsymbol{v}\|_{\tilde{\mathcal{G}}}:=\sqrt{\left.\tilde{\mathcal{G}}\right|_{g}(\boldsymbol{v}, \boldsymbol{v})}:=\sqrt{\left.\mathcal{G}\right|_{g p_{0}}\left(\pi_{*} \boldsymbol{v}, \pi_{*} \boldsymbol{v}\right)} \tag{13}
\end{equation*}
$$

In the same fashion, we have an induced pseudometric on $G$ from the pseudometric tensor field on $G$.

Definition 11 (Pseudometric on $G$ ) Let $g_{1}, g_{2} \in G$. Then, we define:

$$
\begin{align*}
& d_{\tilde{\mathcal{G}}}\left(g_{1}, g_{2}\right):=d_{G, \tilde{\mathcal{G}}}\left(g_{1}, g_{2}\right):= \\
& \quad \inf _{\substack{\gamma \in \operatorname{Lip}([0,1], G) \\
\gamma(0)=g_{1}, \gamma(1)=g_{2}}} \int_{0}^{1} \sqrt{\left.\tilde{\mathcal{G}}\right|_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))} \mathrm{d} t . \tag{14}
\end{align*}
$$

This pseudometric has the property that $d_{\tilde{\mathcal{G}}}\left(h_{1}, h_{2}\right)=0$ for all $h_{1}, h_{1} \in H$; in fact for all $p \in G / H$ we have that $d_{\tilde{\mathcal{G}}}\left(g_{1}, g_{2}\right)=0$ for all $g_{1}, g_{2} \in p$.

By requiring $G$ and $H$ to be connected, we get the following strong correspondence between the metric structure on the homogeneous space and the pseudometric structure on the group.

Lemma 12 Let $g_{1}, g_{2} \in G$ so that $\pi\left(g_{2}\right)$ is away from the cut locus of $\pi\left(g_{1}\right)$, then:
$d_{\tilde{\mathcal{G}}}\left(g_{1}, g_{2}\right)=d_{\mathcal{G}}\left(\pi\left(g_{1}\right), \pi\left(g_{2}\right)\right)$.

Moreover, if $\gamma$ is a minimizing geodesic in the group $G$ connecting $g_{1}$ with $g_{2}$, then $\pi \circ \gamma$ is the unique minimizing geodesic in the homogeneous space $G / H$ that connects $\pi\left(g_{1}\right)$ with $\pi\left(g_{2}\right)$.

Proof Assuming it exists, let $\gamma \in \operatorname{Lip}([0,1], G)$ be a minimizing geodesic connecting $\gamma(0)=g_{1}$ with $\gamma(1)=g_{2}$ and let $\beta \in \operatorname{Lip}([0,1], G / H)$ be the unique minimizing geodesic connecting $\beta(0)=\pi\left(g_{1}\right)$ with $\beta(1)=\pi\left(g_{2}\right)$. Because of the pseudometric on $G$, minimizing geodesics are not unique, i.e. $\gamma$ is not unique. On $G / H$, we have a full metric and so staying away from the cut locus means $\beta$ is both unique and minimizing.

Denote the length functionals with:

$$
\begin{aligned}
& \operatorname{Len}_{G}(\gamma):=\int_{0}^{1} \sqrt{\left.\tilde{\mathcal{G}}\right|_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))} \mathrm{d} t \\
& \operatorname{Len}_{G / H}(\beta)
\end{aligned}=\int_{0}^{1} \sqrt{\left.\mathcal{G}\right|_{\beta(t)}(\dot{\beta}(t), \dot{\beta}(t))} \mathrm{d} t .
$$

Observe that by construction of the pseudometric tensor field $\tilde{\mathcal{G}}$ on $G$, we have: $\operatorname{Len}_{G}(\gamma)=\operatorname{Len}_{G / H}(\pi \circ \gamma)$.

Now we assume $\pi \circ \gamma \neq \beta$. Then, since $\beta$ is the unique geodesic, we have
$\operatorname{Len}_{G / H}(\beta)<\operatorname{Len}_{G / H}(\pi \circ \gamma)=\operatorname{Len}_{G}(\gamma)$.
But then we can find some $\gamma_{\text {lift }} \in \operatorname{Lip}([0,1], G)$ that is a preimage of $\beta$, i.e., $\pi \circ \gamma_{\text {lift }}=\beta$. The potential problem is that while $\gamma_{\text {lift }}(0) \in \pi\left(g_{1}\right)$ and $\gamma_{\text {lift }}(1) \in \pi\left(g_{2}\right)$, $\gamma_{\text {lift }}$ does not necessarily connect $g_{1}$ to $g_{2}$. But since the coset $\pi\left(g_{1}\right)$ is connected, we can find a curve wholly contained in it that connects $g_{1}$ with $\gamma_{\text {lift }}(0)$, call this curve $\gamma_{\text {head }} \in \operatorname{Lip}\left([0,1], \pi\left(g_{1}\right)\right)$. Similarly we can find a $\gamma_{\text {tail }} \in$ $\operatorname{Lip}\left([0,1], \pi\left(g_{2}\right)\right)$ that connects $\gamma_{\text {lift }}(1)$ to $g_{2}$. Both these curves have zero length since $\pi$ maps them to a single point on $G / H$, i.e., $\operatorname{Len}_{G}\left(\gamma_{\text {head }}\right)=\operatorname{Len}_{G}\left(\gamma_{\text {tail }}\right)=0$.

Now we can compose these three curves:
$\gamma_{\text {new }}(t):= \begin{cases}\gamma_{\text {head }}(3 t) & \text { if } t \in[0,1 / 3], \\ \gamma_{\text {lift }}(3 t-1) & \text { if } t \in[1 / 3,2 / 3], \\ \gamma_{\text {tail }}(3 t-2) & \text { if } t \in[2 / 3,1] .\end{cases}$
This new curve is again in $\operatorname{Lip}([0,1], G)$ and connects $g_{1}$ with $g_{2}$, but also:
$\operatorname{Len}_{G}\left(\gamma_{\text {new }}\right)=\operatorname{Len}_{G}(\gamma$ lift $)=\operatorname{Len}_{G / H}(\beta)<\operatorname{Len}_{G}(\gamma)$,
which is a contradiction since $\gamma$ is a minimizing geodesic between $g_{1}$ and $g_{2}$. We conclude $\pi \circ \gamma=\beta$ and thereby:
$d_{\tilde{\mathcal{G}}}\left(g_{1}, g_{2}\right)=\operatorname{Len}_{G}(\gamma)=\operatorname{Len}_{G / H}(\beta)=d_{\mathcal{G}}\left(\pi\left(g_{1}\right), \pi\left(g_{2}\right)\right)$.

This result allows us to more easily translate results from Lie groups to homogeneous spaces.

We end our theoretical preliminaries by introducing the space of positions and orientations $\mathbb{M}_{d}$.

### 3.3 Example: The Group $\operatorname{SE}(\boldsymbol{d})$ and the Homogeneous Space $\mathbb{M}_{\boldsymbol{d}}$

Our main example and Lie group of interest are the Special Euclidean group $S E(d)$ of the rotations and translations of $\mathbb{R}^{d}$, in particular for $d \in\{2,3\}$. When we take $H=\{0\} \times$ $S O(d-1)$, we obtain the space of positions and orientations
$\mathbb{M}_{d}=S E(d) /(\{0\} \times S O(d-1))$.
This homogeneous space and group will enable the construction of roto-translation equivariant networks. For the experiments in this paper, we restrict ourselves to $d=2$, but we include the case $d=3$ in some of our theoretical preliminaries.

As a set we identify $\mathbb{M}_{d}$ with $\mathbb{R}^{d} \times S^{d-1}$ and choose some reference direction $\boldsymbol{a} \in S^{d-1} \subset \mathbb{R}^{d}$ as the forward direction, we can set the reference point of the space as $p_{0}=(\mathbf{0}, \boldsymbol{a})$. We can then see that elements of $H$ are those rotations that map $\boldsymbol{a}$ to itself, i.e., rotations with the forward direction as their axis.

If we denote elements of $S E(d)$ as translation/rotation pairs $(\boldsymbol{y}, R) \in \mathbb{R}^{d} \times S O(d)$, then group multiplication is given by
$g_{1}=\left(\boldsymbol{y}_{1}, R_{1}\right), g_{2}=\left(\boldsymbol{y}_{2}, R_{2}\right) \in G:$
$g_{1} g_{2}=\left(\boldsymbol{y}_{1}, R_{1}\right)\left(\boldsymbol{y}_{2}, R_{2}\right)=\left(\boldsymbol{y}_{1}+R_{1} \boldsymbol{y}_{2}, R_{1} R_{2}\right)$,
and the group action on elements $p=(\boldsymbol{x}, \boldsymbol{n}) \in \mathbb{R}^{d} \times S^{d-1} \equiv$ $\mathbb{M}_{d}$ is given as
$g p=(\boldsymbol{y}, R)(\boldsymbol{x}, \boldsymbol{n})=(\boldsymbol{y}+R \boldsymbol{x}, R \boldsymbol{n})$.
What the G-invariant vector field and metric tensor fields look like on $\mathbb{M}_{d}$ is different for $d=2$ than for $d>2$. We first look at $d>2$.

Proposition 13 Let d $>2$ and let $\partial_{a} \in T_{p_{0}}\left(\mathbb{M}_{d}\right)$ be the tangent vector in the reference point in the main direction $\boldsymbol{a} \in S^{d-1}$, specifically:
$\partial_{\boldsymbol{a}} f:=\lim _{t \rightarrow 0} \frac{f((t \boldsymbol{a}, \boldsymbol{a}))-f((\mathbf{0}, \boldsymbol{a}))}{t}$,
where $f: \mathbb{M}_{d} \rightarrow \mathbb{R}$ is smooth in an open neighborhood of $p_{0}=(0, \boldsymbol{a})$; then, all $S E(d)$-invariant vector fields are spanned by the vector field:
$\left.p \mapsto \mathcal{A}_{1}\right|_{p}:=\left(L_{g_{p}}\right)_{*} \partial_{\boldsymbol{a}}$,
with $g_{p} \in p \in \mathbb{M}_{d}$.
Proof For $d>3$, we can see that (17) are the only left-invariant vector fields since for all $h \in H$ we have $\left(g_{p} h\right) p_{0}=p$, and so in order to be well defined we must require $\left(L_{h}\right)_{*} \boldsymbol{v}=\boldsymbol{v}$ on $T_{p_{0}}\left(\mathbb{M}_{d}\right)$, and this is true for $\partial_{\boldsymbol{a}}$ (and its scalar multiples) but not true for any other tangent vectors at $T_{p_{0}}\left(\mathbb{M}_{d}\right)$.

Proposition 14 For $d>2$, the only Riemannian metric tensor fields on $\mathbb{M}_{d}$ that are $S E(d)$-invariant are of the form:

$$
\begin{align*}
& \left.\mathcal{G}\right|_{(\boldsymbol{x}, \boldsymbol{n})}((\dot{\boldsymbol{x}}, \dot{\boldsymbol{n}}),(\dot{\boldsymbol{x}}, \dot{\boldsymbol{n}})) \\
& \quad=w_{M}|\dot{\boldsymbol{x}} \quad \boldsymbol{n}|^{2}+w_{L}\|\dot{\boldsymbol{x}} \wedge \boldsymbol{n}\|^{2}+w_{A}\|\dot{\boldsymbol{n}}\|^{2} \tag{18}
\end{align*}
$$

with $w_{M}, w_{L}, w_{A}>0$ weighing the main, lateral and angular motion, respectively, and where the inner product, outer product and norm are the standard Euclidean constructs.

Proof It follows that to satisfy the second condition of Corollary 5 at the tangent space $T_{(\boldsymbol{x}, \boldsymbol{n})}$ of a particular $(\boldsymbol{x}, \boldsymbol{n})$ the metric tensor needs to be symmetric with respect to rotations about $\boldsymbol{n}$ both spatially and angularly (i.e., we require isotropy in all angular and lateral directions) which leads to the three degrees of freedom contained in (18) irrespective of $d$.

For $d=2$, we represent the elements of $\mathbb{M}_{2}$ with $(x, y, \theta) \in \mathbb{R}^{3}$ where $x, y$ are the usual Cartesian coordinates and $\theta$ the angle with respect to the $x$-axis, so that $\boldsymbol{n}=(\cos \theta, \sin \theta)^{T}$. The reference element is then simply denoted by $(0,0,0)$.

It may be counter-intuitive but decreasing the number of dimensions to 2 gives more freedom to the $G$-invariant vector and metric tensor fields compared to $d>2$. This is a consequence of the subgroup $H$ being trivial, and so the symmetry conditions from Corollaries 3 and 5 also become trivial. The $S E(2)$-invariant vector fields are given as follows.

Proposition 15 On $\mathbb{M}_{2}$, the SE(2)-invariant vector fields are spanned by the following basis:

$$
\begin{cases}\left.\mathcal{A}_{1}\right|_{(x, y, \theta)} & =\left.\cos \theta \partial_{x}\right|_{(x, y, \theta)}+\left.\sin \theta \partial_{y}\right|_{(x, y, \theta)}  \tag{19}\\ \left.\mathcal{A}_{2}\right|_{(x, y, \theta)} & =-\left.\sin \theta \partial_{x}\right|_{(x, y, \theta)}+\left.\cos \theta \partial_{y}\right|_{(x, y, \theta)} \\ \left.\mathcal{A}_{3}\right|_{(x, y, \theta)} & =\left.\partial_{\theta}\right|_{(x, y, \theta)}\end{cases}
$$

Proof For $d=2$, we have $\mathbb{M}_{2} \equiv S E(d)$ and the group invariant vector fields on $\mathbb{M}_{2}$ are exactly the left-invariant vector fields on $S E(2)$ given by (19).

In a similar manner, $S E(2)$-invariant metric tensors are then given as follows.

Proposition 16 On $\mathbb{M}_{2}$, the $S E(2)$-invariant metric tensor fields are given by:

$$
\left.\mathcal{G}\right|_{(x, y, \theta)}(\boldsymbol{v}, \boldsymbol{w})=\left.\mathcal{G}\right|_{(0,0,0)}\left(\left(L_{(x, y, \theta)}^{-1}\right)_{*} \boldsymbol{v},\left(L_{(x, y, \theta)}^{-1}\right)_{*} \boldsymbol{w}\right)
$$

for any choice of inner product $\left.\mathcal{G}\right|_{(0,0,0)}$ at $e$.
Proof Since $S E(2) \equiv \mathbb{M}_{2}$, the $G$-invariant metric tensor fields are again exactly the left-invariant metric tensor fields.

This gives $S E$ (2)-invariant metric tensor fields 6 degrees of freedom and hence 6 trainable parameters on $\mathbb{M}_{2}$. Remarkably, the case $d=2$ allows for more degrees of freedom than the case $d=3$ where Proposition 14 applies. In our experiments so far, we have restricted ourselves to those metric tensors that are diagonal with respect to the frame from Proposition 15. A diagonal metric tensor would have just 3 degrees of freedom and have the same general form as (18), specifically:

$$
\begin{align*}
\left.\mathcal{G}\right|_{(x, y, \theta)} & ((\dot{x}, \dot{y}, \dot{\theta}),(\dot{x}, \dot{y}, \dot{\theta})) \\
= & w_{M}|\dot{x} \cos \theta+\dot{y} \sin \theta|^{2} \\
& +w_{L}|-\dot{x} \sin \theta+\dot{y} \cos \theta|^{2}  \tag{20}\\
& +w_{A}|\dot{\theta}|^{2}
\end{align*}
$$

We will expand into non-diagonal metric tensors in future work.

## 4 Architecture

### 4.1 Lifting and Projecting

The key ingredient of what we call a PDE-G-CNN is the PDE layer that we detail in the next section; however, to make a complete network we need more. Specifically we need a layer that transforms the network's input into a format that is suitable for the PDE layers and a layer that takes the output of the PDE layers and transforms it to the desired output format. We call this input and output transformation lifting, respectively, projection; this yields the overall architecture of a PDE-G-CNN as illustrated in Fig. 5.

As our theoretical preliminaries suggest, we aim to do processing on homogeneous spaces, but the input and output of the network do not necessarily live on that homogeneous space. Indeed, in the case of images the data live on $\mathbb{R}^{2}$ and not on $\mathbb{M}_{2}$ where we propose to do processing.

This necessitates the addition of lifting and projection layers to first transform the input to the desired homogeneous space and end with transforming it back to the required output space. Of course, for the entire network to be equivariant we require these transformation layers to be equivariant as well. In this paper, we focus on the design of the PDE layers, details on appropriate equivariant lifting and projection layers in the case of $S E(2)$ is found in $[20,80]$.

Remark 4 (General equivariant linear transformations between homogeneous spaces) A general way to lift and project from one homogeneous space to another in a trainable fashion is the following. Consider two homogeneous spaces $G / H_{1}$ and $G / H_{2}$ of a Lie group $G$; let $f: G / H_{1} \rightarrow \mathbb{R}$ and $k: G / H_{2} \rightarrow \mathbb{R}$ with the following property:
$\forall h \in H_{1}, q \in G / H_{2}: k(h q)=k(q)$,
where $H_{1}$ is compact. Then, the operator $\mathcal{T}$ defined by
$\forall q \in G / H_{2}:(\mathcal{T} f)(q):=\int_{G} k\left(g^{-1} q\right) f\left(g H_{1}\right) \mathrm{d} \mu_{G}(g)$
transforms $f$ from a function on $G / H_{1}$ to a function on $G / H_{2}$ in an equivariant manner (assuming $f$ and $k$ are such that the


Fig. 5 Illustrating the overall architecture of a PDE-G-CNN (example: retinal vessel segmentation). An input image is lifted to a homogeneous space from which point on it can be fed through subsequent PDE layers
(each PDE layer follow the structure of Fig. 2) that replace the convolution layers in conventional CNNs. Finally, the result is projected back to the desired output space
integral exists). Here the kernel $k$ is the trainable part and $\mu_{G}$ is the left-invariant Haar measure on the group.

Moreover, it can be shown via the Dunford-Pettis [81] theorem that (under mild restrictions) all linear transforms between homogeneous spaces are of this form.

Remark 5 (Lifting and projecting on $\mathbb{M}_{2}$ ) Lifting an image (function) on $\mathbb{R}^{2}$ to $\mathbb{M}_{2}$ can either be performed by a nontrainable Invertible Orientation Score Transform [36] or a trainable lift [20] in the style of Remark 4.

Projecting from $\mathbb{M}_{2}$ back down to $\mathbb{R}^{2}$ can be performed by a simple maximum projection: let $f: \mathbb{M}_{2} \rightarrow \mathbb{R}$, then
$(x, y) \mapsto \max _{\theta \in[0,2 \pi)} f(x, y, \theta)$
is a roto-translation equivariant projection as used in [20]. A variation on the above projection is detailed in [80, Ch. 3.3.3].

### 4.2 PDE Layer

A PDE layer operates by taking its inputs as the initial conditions for a set of evolution equations; hence, there will be a PDE associated with each input feature. The idea is that we let each of these evolution equations work on the inputs up to a fixed time $T>0$. Afterward, we take these solutions at time $T$ and take affine combinations (really batch normalized linear combinations in practice) of them to produce the outputs of the layer and as such the initial conditions for the next set of PDEs.

If we index network layers (i.e., the depth of the network) with $\ell$ and denote the width (i.e., the number of features or channels) at layer $\ell$ with $M_{\ell}$, then we have $M_{\ell}$ PDEs and take $M_{\ell+1}$ linear combinations of their solutions. We divide a PDE layer into the PDE solvers that each apply the PDE evolution to their respective input channel and the affine combination unit. This design is illustrated in Fig. 1, but let us formalize it.

Let $\left(U_{\ell, c}\right)_{c=1}^{M_{\ell}}$ be the inputs of the $\ell$-th layer (i.e., some functions on $G / H)$; let $a_{\ell i j}$ and $b_{\ell i} \in \mathbb{R}$ be the coefficients of the affine transforms for $i=1 \ldots M_{\ell+1}$ and $j=1 \ldots M_{\ell}$. Let each PDE be parametrized by a set of parameters $\theta_{\ell j}$. Then, the action of a PDE layer is described as:
$U_{\ell+1, i}=\sum_{j=1}^{M_{\ell}} a_{\ell i j} \Phi_{T, \theta_{\ell j}}\left(U_{\ell j}\right)+b_{\ell i}$,
where $\Phi_{T, \theta}$ is the evolution operator of the PDE at time $T \geq 0$ and parameter set $\theta$. We define the operator $\Phi_{t, \theta}$ so that $(p, t) \mapsto\left(\Phi_{t, \theta} U\right)(p)$ satisfies the Hamilton-Jacobi type PDE that we introduce in just a moment. In this layer formula, the parameters $a_{\ell i j}, b_{\ell i}$ and $\theta_{\ell j}$ are the trainable weights, but the evolution time $T$ we keep fixed.

It is essential that we require the network layers and thereby all the PDE units to be equivariant. This has consequences for the class of PDEs that is allowed.

The PDE solver that we will consider in this article, illustrated in Fig. 2, computes the approximate solution to the PDE

$$
\left\{\begin{array}{rlr}
\frac{\partial W}{\partial t}(p, t)= & -\boldsymbol{c} W(p, t) & \underline{\text { (convection) }}  \tag{24}\\
& -\left(-\Delta_{\mathcal{G}_{1}}\right)^{\alpha} W(p, t) & \underline{\text { (diffusion) }} \\
& +\left\|\nabla_{\mathcal{G}_{2}^{+}} W(p, t)\right\|_{\mathcal{G}_{2}^{+}}^{2 \alpha} & \underline{\text { (dilation) }} \\
-\left\|\nabla_{\mathcal{G}_{2}^{-}} W(p, t)\right\|_{\mathcal{G}_{2}^{-}}^{2 \alpha} & \underline{\text { (erosion) }} \\
W(p, 0)=U(p) & \text { for } p \in G / H, t \geq 0 \\
& \text { for } p \in G / H
\end{array}\right.
$$

Here, $\boldsymbol{c}$ is a $G$-invariant vector field on $G / H$ (recall (17) and our use of tangent vectors as differential operators per Remark 3), $\alpha \in[1 / 2,1], \mathcal{G}_{1}$ and $\mathcal{G}_{2}^{ \pm}$are $G$-invariant metric tensor fields on $G / H, U$ is the initial condition and $\Delta_{\mathcal{G}}$,
and $\|\cdot\|_{\mathcal{G}}$ denote the Laplace-Beltrami operator and norm induced by the metric tensor field $\mathcal{G}$. As the labels indicate, the four terms have distinct effects:

- convection: moving data around,
- (fractional) diffusion: regularizing data (which relates to subsampling by destroying data),
- dilation: pooling of data,
- erosion: sharpening of data.

This is also why we refer to a layer using this PDE as a CDDE layer. Summarized parameters of this PDE are given by $\theta=\left(c, \mathcal{G}_{1}, \mathcal{G}_{2}^{+}, \mathcal{G}_{2}^{-}\right)$. The geometric interpretation of each of the terms in (24) is illustrated in Fig. 6 for $G=\mathbb{R}^{2}$ and in Fig. 7 for $G=\mathbb{M}_{2}$.

Since the convection vector field $\boldsymbol{c}$ and the metric tensor fields $\mathcal{G}_{1}$ and $\mathcal{G}_{2}^{ \pm}$are $G$-invariant, the PDE unit, and so the network layer, is automatically equivariant.

### 4.3 Training

Training the PDE layer comes down to adapting the parameters in the PDEs in order to minimize a given loss function (the choice of which depends on the application and we will not consider in this article). In this sense, the vector field and the metric tensors are analogous to the weights of this layer.

Since we required the convection vector field and the metric tensor fields to be $G$ invariant, the parameter space is finite-dimensional as a consequence of Corollaries 3 and 5 if we restrict ourselves to Riemannian metric tensor fields.

For our main application on $\mathbb{M}_{2}$, each PDE unit would have the following 12 trainable parameters:

- 3 parameters to specify the convection vector field as a linear combination of (19),
- 3 parameters to specify the fractional diffusion metric tensor field $\mathcal{G}_{1}$,
- and 3 parameters each to specify the dilation and erosion metric tensor fields $\mathcal{G}_{2}^{ \pm}$,
where the metric tensor fields are of the form (20) that are diagonal with respect to the frame from Proposition 15.

Surprisingly for higher dimensions, $\mathbb{M}_{d}$ has less trainable parameters than for $d=2$. This is caused by the $S E(d)$ invariant vector fields on $\mathbb{M}_{d}$ for $d \geq 3$ being spanned by a single basis element (per Proposition 13) instead of the three (19) basis elements available for $d=2$. Since the left-invariant metric tensor fields are determined by only 3 parameters irrespective of dimensions, we count a total of 7 parameters for each PDE unit for applications on $\mathbb{M}_{d}$ for $d \geq 3$.

In our own experiments, we always use some form of stochastic gradient descent (usually ADAM) with a small
amount of $L^{2}$ regularization applied uniformly over all the parameters. Similarly, we stick to a single learning rate for all the parameters. Given that in our setting different parameters have distinct effects, treating all of them the same is likely far from optimal, however, we leave that topic for future investigation.

## 5 PDE Solver

Our PDE solver will consist of an iteration of time step units, each of which is a composition of convection, diffusion, dilation and erosion substeps. These units all take their input as an initial condition of a PDE and produce as output the solution of a PDE at time $t=T$.

The convection, diffusion and dilation/erosion steps are implemented with, respectively, a shifted resample, linear convolution and two morphological convolutions, as illustrated in Fig. 8.

The composition of the substeps does not solve (24) exactly, but for small $\Delta t$, it approximates the solution by a principle called operator splitting.

We will now discuss each of these substeps separately.

### 5.1 Convection

The convection step has as input a function $U_{1}: G / H \rightarrow \mathbb{R}$ and takes it as initial condition of the PDE

$$
\begin{cases}\frac{\partial W_{1}}{\partial t}(p, t)=-\boldsymbol{c}(p) W_{1}(\cdot, t) & \text { for } p \in G / H, t \geq 0  \tag{25}\\ W_{1}(p, 0)=U_{1}(p) & \text { for } p \in G / H\end{cases}
$$

The output of the layer is the solution of the PDE evaluated at time $t=T$, i.e., the output is the function $p \mapsto W_{1}(p, T)$.

Proposition 17 (Convection solution) The solution of the convection PDE (25) is found by the method of characteristics and is given by

$$
\begin{align*}
W^{1}(p, t) & =\left(\mathcal{L}_{g_{p}^{-1}} U_{1}\right)\left(\gamma_{c}(t)^{-1} p_{0}\right) \\
& =U_{1}\left(g_{p} \gamma_{c}(t)^{-1} p_{0}\right)  \tag{26}\\
& =U_{1}\left(g_{p} \gamma_{-c}(t) p_{0}\right), \tag{27}
\end{align*}
$$

where $g_{p} \in p$ (i.e., $g_{p} p_{0}=p$ ) and $\gamma_{c}: \mathbb{R} \rightarrow G$ is the exponential curve that satisfies $\gamma_{c}(0)=e$ and
$\frac{\partial}{\partial t}\left(\gamma_{c}(t) p\right)(t)=\boldsymbol{c}\left(\gamma_{c}(t) p\right)$,
i.e., $\gamma_{c}$ is the exponential curve in the group $G$ that induces the integral curves of the $G$-invariant vector field $\boldsymbol{c}$ on $G / H$ when acting on elements of the homogeneous space.


Fig. 6 Geometric interpretation of the terms of the PDE (24) illustrated for $\mathbb{R}^{2}$. In this setting, the $G$-invariant vector field $\boldsymbol{c}$ is the constant vector field given by two translation parameters. For the other terms, we use Riemannian metric tensors parametrized by a positive definite $2 \times 2$
matrix in the standard basis. The kernels used in the diffusion, dilation and erosion terms are functions of the distance map induced by the metric tensors

Note that this exponential curve existing is a consequence of the vector field $\boldsymbol{c}$ being $G$-invariant, such exponential curves do not exist for general convection vector fields.

## Proof

$$
\begin{aligned}
& \frac{\partial W_{1}}{\partial t}(p, t) \\
& =\lim _{h \rightarrow 0} \frac{W_{1}(p, t+h)-W_{1}(p, t)}{h} \\
& \quad=\lim _{h \rightarrow 0} \frac{U_{1}\left(g_{p} \gamma_{c}(t+h)^{-1} p_{0}\right)-U_{1}\left(g_{p} \gamma_{c}(t)^{-1} p_{0}\right)}{h} \\
& \quad=\lim _{h \rightarrow 0} \frac{U_{1}\left(g_{p} \gamma_{c}(t)^{-1} \gamma_{c}(h)^{-1} p_{0}\right)-U_{1}\left(g_{p} \gamma_{c}(t)^{-1} p_{0}\right)}{h}
\end{aligned}
$$

due to the $G$-invariance of $\boldsymbol{c}$ this yields

$$
\begin{aligned}
& =-\boldsymbol{c}(p) \mathcal{L}_{g_{p}} \mathcal{L}_{\gamma_{\boldsymbol{c}}(t) g_{p}^{-1} U_{1}} \\
& =-\boldsymbol{c}(p)\left[p \mapsto U_{1}\left(g_{p} \gamma_{c}(t)^{-1} g_{p}^{-1} p\right)\right] \\
& =-\boldsymbol{c}(p)\left[p \mapsto U_{1}\left(g_{p} \gamma_{c}(t)^{-1} p_{0}\right)\right] \\
& =-\boldsymbol{c}(p) W_{1}(\cdot, t)
\end{aligned}
$$

In our experiments, Eq. (27) is numerically implemented as a resampling operation with trilinear interpolation to account for the off-grid coordinates.
now let $\bar{U}:=\mathcal{L}_{\gamma_{c}(t) g_{p}^{-1}} U_{1}$, then

$$
\begin{aligned}
& =\lim _{h \rightarrow 0} \frac{\bar{U}\left(\gamma_{c}(h)^{-1} p_{0}\right)-\bar{U}\left(p_{0}\right)}{h} \\
& =-\boldsymbol{c}\left(p_{0}\right) \bar{U} \\
& =-\left(L_{g_{p}}\right)_{*} \boldsymbol{c}\left(p_{0}\right) \mathcal{L}_{g_{p}} \bar{U}
\end{aligned}
$$

### 5.2 Fractional Diffusion

The (fractional) diffusion step solves the PDE

$$
\begin{cases}\frac{\partial W_{2}}{\partial t}=-\left(-\Delta_{\mathcal{G}_{1}}\right)^{\alpha} W_{2}(p, t) & \text { for } p \in G / H, t \geq 0  \tag{29}\\ W_{2}(p, 0)=U_{2}(p) & \text { for } p \in G / H\end{cases}
$$

| $\mathbb{M}_{2}$ | Transport | Regularization | Max pooling | Min pooling |
| :---: | :---: | :---: | :---: | :---: |
| PDE Term: | -cW | $-\left(-\Delta_{\mathcal{G}_{1}}\right)^{\alpha} W$ | $+\left\\|\nabla_{\mathcal{G}_{2}^{+}} W\right\\|_{\mathcal{G}_{2}^{+}}^{2 \alpha}$ | $-\left\\|\nabla_{\mathcal{G}_{2}^{-}} W\right\\|_{\mathcal{G}_{2}^{-}}^{2 \alpha}$ |
| Parameters: | G-invariant vector field | Riemannian metric tensor | Riemannian metric tensor | Riemannian metric tensor |
|  | $\left(\begin{array}{l} 1 \\ 0 \\ 2 \end{array}\right)$ | $\left(\begin{array}{lll} 4 & 0 & 0 \\ 0 & 6 & 0 \\ 0 & 0 & 1 \end{array}\right)$ | $\left(\begin{array}{ccc} 4 & 0 & 0 \\ 0 & 9 & 0 \\ 0 & 0 & 0.5 \end{array}\right)$ | $\left(\begin{array}{ccc} 6 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 1.5 \end{array}\right)$ |
| Operation: | resample <br> with offset | convolution with kernel | dilation with kernel | erosion with kernel |
|  |  |  |  |  |







Fig. 7 Geometric interpretation of the terms of the PDE (24) illustrated for $\mathbb{M}_{2}$. In this setting, the $G$-invariant vector field $\boldsymbol{c}$ is a left-invariant vector field given by two translation and one rotation parameter. For the other terms, we use Riemannian metric tensors parametrized by a positive definite $3 \times 3$ matrix in the left-invariant basis (the matrix does not
need to be diagonal, but we keep that for future work). The kernels used in the diffusion, dilation and erosion terms are functions of the distance map induced by the metric tensors and are visualized by partial plots of their level sets
Trainables:
c
$\mathcal{G}_{1}$
$\mathcal{G}_{2}^{+}$
$\mathcal{G}_{2}^{-}$

$$
\underset{\downarrow}{\downarrow}
$$

$\sigma_{2}^{-}$

Fig. 8 Evolving the PDE through operator splitting, each operation corresponds to a term of (24)

As with (fractional) diffusion on $\mathbb{R}^{n}$, there exists a smooth function
$K_{.}^{\alpha}:(0, \infty) \times(G / H) \rightarrow[0, \infty)$,
called the fundamental solution of the $\alpha$-diffusion equation, such that for every initial condition $U_{2}$; the solution to the PDE (29) is given by the convolution of the function $U_{2}$ with the fundamental solution $K_{t}^{\alpha}$ :

$$
\begin{equation*}
W^{2}(p, t)=\left(K_{t}^{\alpha} *_{G / H} U_{2}\right)(p) \tag{30}
\end{equation*}
$$

The convolution $*_{G / H}$ on a homogeneous space $G / H$ is specified by the following definition.

Definition 18 (Linear group convolution) Let $p_{0}=H$ be compact, let $f \in L^{2}(G / H)$ and $k \in L^{1}(G / H)$ such that:
$\forall h \in H, p \in G / H: k(h p)=k(p)$,
(kernel compatibility)
then, we define:

$$
\begin{align*}
& \left(k *_{G / H} f\right)(p) \\
& \quad:=\frac{1}{\mu_{H}(H)} \int_{G} k\left(g^{-1} p\right) f\left(g p_{0}\right) \mathrm{d} \mu_{G}(g), \tag{31}
\end{align*}
$$

where $\mu_{H}$ and $\mu_{G}$ are the left-invariant Haar measures (determined up to scalar multiplication) on $H$, respectively, $G$.

Remark 6 In the remainder of this article, we refer the to the left-invariant Haar measure on $G$ as 'the Haar measure on $G^{\prime}$ as right-invariant Haar measures on $G$ do not play a role in our framework.

Remark 7 Compactness of $H$ is crucial as otherwise the integral in the right-hand side of (31) does not converge. To this end, we note that one can always decompose (by Weil's integral formula [82, Lem. 2.1]) the Haar measure $\mu_{G}$ on the group as a product of a measure on the quotient $G / H$ times the measure on the subgroup $H$. As Haar measures are determined up to a constant, we take the following convention: we normalize the Haar measure $\mu_{G}$ such that
$\mu_{G}\left(\pi^{-1}(A)\right)=\mu_{H}(H) \mu_{\mathcal{G}}(A), \quad \forall A \subset G / H$,
where $\mu_{\mathcal{G}}$ is the Riemannian measure induced by $\mathcal{G}$ and $\mu_{H}$ is a choice of Haar measure on $H$. Thereby, (32) boils down to Weil's integration formula:
$\mu_{H}(H) \int_{G / H} f(p) \mathrm{d} \mu_{\mathcal{G}}(p)=\int_{G} f(g H) \mathrm{d} \mu_{G}(g)$
whenever $f$ is measurable. Since $H$ is compact, we can indeed normalize the Haar measure $\mu_{H}$ so that $\mu_{H}(H)=1$.

In general, an exact analytic expression for the fundamental solution $K_{t}^{\alpha}$ requires complicated steerable filter operators [50, Thm. $1 \& 2$ ], and for that reason we contend ourselves with more easily computable approximations. For now, let us construct our approximations and address their quality and the involved asymptotics later.

Remark 8 In the approximations, we will make use of logarithmic map as the inverse of the Lie group exponential map $\exp _{G}$. Locally, such inversion can always be done by the inverse function theorem. Specifically, there is always a neighborhood $V \subset T_{e} G$ of the origin so that $\left.\exp _{G}\right|_{V}$ is a diffeomorphism between $V$ and $W=\exp _{G}(V) \subset G$, where $W$ is a neighborhood of $e$. Then we define the logarithmic map $\log _{G}: W \rightarrow V$ by $\exp _{G} \circ \log _{G}=\mathrm{id}_{W}$ and $\left.\log _{G} \circ \exp _{G}\right|_{V}=\mathrm{id}_{V}$. For the moment, for simplicity, we assume $V=T_{e}(G)$ in the general setting. ${ }^{1}$

The idea is that instead of basing our kernels on the metric $d_{\mathcal{G}}$ (which is hard to calculate [83]) we approximate it using the seminorm from Definition 10 (which is easy to calculate). We can use this seminorm on elements of the homogeneous space by using the group's logarithmic map $\log _{G}$. We can take

[^1]the group logarithm of all the group elements that constitute a particular equivalence class of $G / H$ and then pick the group element with the lowest seminorm:
$d_{\mathcal{G}}\left(p_{0}, p\right) \approx \inf _{g \in p}\left\|\log _{G} g\right\|_{\tilde{\mathcal{G}}}$.
Henceforth, we write this estimate as $d_{\mathcal{G}}\left(p_{0}, p\right) \approx \rho_{\mathcal{G}}(p)$ relying on the following definition.

Definition 19 (Logarithmic metric estimate) Let $\mathcal{G}$ be a $G$ invariant metric tensor field on the homogeneous space $G / H$, then we define

$$
\begin{align*}
\rho_{\mathcal{G}}(p) & :=\inf _{g \in p}\left\|\log _{G} g\right\|_{\tilde{\mathcal{G}}} \\
& :=\inf _{g \in p} \sqrt{\mathcal{G}\left(\pi_{*} \log _{G} g, \pi_{*} \log _{G} g\right)} \tag{35}
\end{align*}
$$

where $\pi_{*}$ is the push-forward of the projection map $\pi$ given by (4).

We can interpret this metric estimate as finding all exponential curves in $G$ whose actions on the homogeneous space connect $p_{0}$ (at $t=0$ ) to $p$ (at $t=1$ ), and then, from that set we choose the exponential curve that has the lowest constant velocity according to the seminorm in Definition 10 and use that velocity as the distance estimate.

Summarizing, Definition 19 and Eq. (34), can be intuitively reformulated as: 'instead of the length of the geodesic connecting two points of $G / H$, we take the length of the shortest exponential curve connecting those two points.'

The following lemma quantifies how well our estimate approximates the true metric.

Lemma 20 (Bounding the logarithmic metric estimate) For all $p \in G / H$ sufficiently close to $p_{0}$, we have
$d_{\mathcal{G}}\left(p_{0}, p\right)^{2} \leq \rho_{\mathcal{G}}(p)^{2} \leq d_{\mathcal{G}}\left(p_{0}, p\right)^{2}+O\left(d_{\mathcal{G}}\left(p_{0}, p\right)^{4}\right)$,
which has as a weaker corollary that for all compact neighborhoods of $p_{0}$ there exists a $C_{\text {metr }}>1$ so that
$d_{\mathcal{G}}\left(p_{0}, p\right) \leq \rho_{\mathcal{G}}(p) \leq C_{\mathrm{metr}} d_{\mathcal{G}}\left(p_{0}, p\right)$
for all $p$ in that neighborhood. Note that the constant $C_{\text {metr }}$ depends on both the choice of compact neighborhood and the metric tensor field.

The proof of this lemma can be found in Appendix A.1.
Remark 9 (Logarithmic metric estimate in principal homogeneous spaces) When we take a principal homogeneous space such as $\mathbb{M}_{2} \equiv S E(2)$ with a left-invariant metric tensor field, the metric estimate simplifies to
$\rho_{\mathcal{G}}(g)=\left\|\log _{G} g\right\|_{\mathcal{G}_{e}} ;$
hence, we see that this construction generalizes the logarithmic estimate, as used in [84,85], to homogeneous spaces other than the principal.

Remark 10 (Logarithmic metric estimate for $\mathbb{M}_{2}$ ) Using the $(x, y, \theta)$ coordinates for $\mathbb{M}_{2}$ and a left-invariant metric tensor field of the form (20), we formulate the metric estimate in terms of the following auxiliary functions called the exponential coordinates of the first kind:
$c^{1}(x, y, \theta):= \begin{cases}\frac{\theta}{2}\left(y+x \cot \frac{\theta}{2}\right) & \text { if } \theta \neq 0, \\ x & \text { if } \theta=0,\end{cases}$
$c^{2}(x, y, \theta):= \begin{cases}\frac{\theta}{2}\left(-x+y \cot \frac{\theta}{2}\right) & \text { if } \theta \neq 0, \\ y & \text { if } \theta=0,\end{cases}$
$c^{3}(x, y, \theta):=\theta$.
The logarithmic metric estimate for $S E(2)$ is then given by

$$
\begin{aligned}
& \rho_{\mathcal{G}}(x, y, \theta)= \\
& \quad \sqrt{w_{M} c^{1}(x, y, \theta)^{2}+w_{L} c^{2}(x, y, \theta)^{2}+w_{A} c^{3}(x, y, \theta)^{2}}
\end{aligned}
$$

this estimate is illustrated in Fig. 9 where it is contrasted against the exact metric.

We can see that the metric estimate $\rho_{\mathcal{G}}$ (and consequently any function of $\rho_{\mathcal{G}}$ ) has the necessary compatibility property to be a kernel used in convolutions per Definition 18.

Proposition 21 (Kernel compatibility of $\rho_{\mathcal{G}}$ ) Let $\mathcal{G}$ be a $G$ invariant metric tensor field on $G / H$, then we have
$\forall p \in G / H, \forall h \in H: \rho_{\mathcal{G}}(h p)=\rho_{\mathcal{G}}(p)$.
Note that, since we use left cosets, $p h=p$ but $h p \neq$ $p$ in general, this requirement is not trivial. Proof of this proposition is included in Appendix A.2.

Now that we have developed and analyzed the logarithmic metric estimate; we can use it to construct an approximation to the diffusion kernel for $\alpha=1$.

Definition 22 (Approximate $\alpha=1$ kernel)
$K_{t}^{1, \mathrm{appr}}(p):=\eta_{t} \exp \left(-\frac{\rho_{\mathcal{G}}(p)^{2}}{4 t}\right)$
where $\eta_{t}$ is a normalization constant for a given $t$; this can either be the $L^{1}$ normalization constant or in the case of groups of polynomial growth one typically sets $\eta_{t}=$ $\mu_{\mathcal{G}}\left(B\left(p_{0}, \sqrt{t}\right)\right)^{-1}$, see the definition of polynomial growth below.

On Lie groups of polynomial growth this approximate kernel be bounded from above and below by the exact kernels.


Fig. 9 Comparing the 'exact' Riemannian distance (top) obtained through numerically solving the Eikonal equation [29] versus the logarithmic metric estimate (bottom) on $S E$ (2) endowed with a left-invariant Riemannian metric tensor field (20) with $w_{M}=1, w_{L}=2, w_{A}=1 / \pi$. The relative $L^{1}$ error in the plotted volume is 0.20

Definition 23 (Polynomial growth) A Lie group $G$ with leftinvariant Haar measure $\mu_{G}$ is of polynomial growth when the volume of a sphere of radius $r$ around $g \in G$ :
$B(g, r)=\left\{g^{\prime} \in G \mid d_{\tilde{\mathcal{G}}}\left(g, g^{\prime}\right)<r\right\}$,
can be polynomialy bounded as follows: there exist constants $\delta>0$ and $C_{\text {grow }}>0$ so that
$\frac{1}{C_{\text {grow }}} r^{\delta} \leq \mu_{G}(B(g, r)) \leq C_{\text {grow }} r^{\delta}, \quad r \geq 1$,
take note that the exponent $\delta$ is the same on both the lower and upper bound. Since $\mu_{G}$ is left invariant, the choice of $g$ does not matter.

Lemma 24 Let $G$ be of polynomial growth, and let $K_{t}^{1}$ be the fundamental solution to the $\alpha=1$ diffusion equation on $G / H$, then there exist constants $C \geq 1, D_{1} \in(0,1)$ and $D_{2}>D_{1}$ so that for all $t>0$ the following holds:

$$
\begin{equation*}
\frac{1}{C} K_{D_{1} t}^{1}(p) \leq K_{t}^{1, \mathrm{appr}}(p) \leq C K_{D_{2} t}^{1}(p) \tag{38}
\end{equation*}
$$

for all $p \in G / H$.
Proof On a group of polynomial growth, we have $\eta_{t}=$ $\mu_{\mathcal{G}}\left(B\left(p_{0}, \sqrt{t}\right)\right)^{-1}$. If $G$ is of polynomial growth, we can apply [86, Thm. 2.12] to find that there exist constants $C_{1}, C_{2}>0$, and for all $\varepsilon>0$ there exist a constant $C_{\varepsilon}$ so that:

$$
\begin{aligned}
& C_{1} \eta_{t} \exp \left(-\frac{d_{\mathcal{G}}\left(p_{0}, p\right)^{2}}{4 C_{2} t}\right) \leq K_{t}^{1}(p) \\
& \quad \leq C_{\varepsilon} \eta_{t} \exp \left(-\frac{d_{\mathcal{G}}\left(p_{0}, p\right)^{2}}{4(1+\varepsilon) t}\right)
\end{aligned}
$$

Remark 11 (Left vs. right cosets) Note that Maheux [86] uses right cosets while we use left cosets. We can translate the results easily by inversion in view of $(g H)^{-1}=H^{-1} g^{-1}=$ $H g^{-1}$. We then apply the result of Maheux to the correct (invertible) $G$-invariant metric tensor field on $G / H$.

Also note the different (but equivalent) way Maheux relates distance on the group with distance on the homogeneous space. While we use a pseudometric on $G$ induced by a metric on $G / H$, Maheux uses a metric on $G / H$ induced by a metric on $G$ by:

$$
\begin{align*}
d_{G / H}^{\text {maheux }}\left(p_{1}, p_{2}\right) & =\inf _{g_{1} \in p_{1}} \inf _{g_{2} \in p_{2}} d_{G}^{\text {maheux }}\left(g_{1}, g_{2}\right) \\
& =\inf _{g_{2} \in p_{2}} d_{G}^{\text {maheux }}\left(q_{1}, g_{2}\right) \tag{39}
\end{align*}
$$

for any choice of $q_{1} \in p_{1}$. We avoid having to minimize inside the cosets as in (39) thanks to the inherent symmetries in our pseudometric.

Now using the inequalities from Lemma 20, we obtain:

$$
\begin{aligned}
& C_{1} \eta_{t} \exp \left(-\frac{\rho_{\mathcal{G}}(p)^{2}}{4 C_{2} t}\right) \leq K_{t}^{1}(p) \\
& \quad \leq C_{\varepsilon} \eta_{t} \exp \left(-\frac{\rho_{\mathcal{G}}(p)^{2}}{4 C_{\mathrm{metr}}^{2}(1+\varepsilon) t}\right)
\end{aligned}
$$

which leads to:

$$
\begin{aligned}
& C_{1} \frac{\eta_{t}}{\eta_{c_{2} t}} K_{C_{2} t}^{1, \mathrm{appr}}(p) \leq K_{t}^{1}(p) \\
& \quad \leq C_{\varepsilon} \frac{\eta_{t}}{\eta_{C_{\operatorname{metr}}^{2}(1+\varepsilon) t}} K_{C_{\operatorname{metr}}^{2}(1+\varepsilon) t}^{1, \mathrm{appr}}(p)
\end{aligned}
$$

The group $G$ being of polynomial growth also implies $G / H$ is a doubling space [86, Thm. 2.17]. Using the volume doubling and reverse volume doubling property of doubling spaces [87, Prop. 3.2 and 3.3], we find that there exist constants $C_{3}, C_{4}, \beta, \beta^{\prime}>0$ so that:

$$
\begin{aligned}
\frac{\eta_{t}}{\eta_{c_{2} t}} \geq & C_{3}\left(\frac{\sqrt{t}}{\sqrt{C_{2} t}}\right)^{\beta}=C_{3} C_{2}^{-\beta / 2} \\
\frac{\eta_{t}}{\eta_{C_{\mathrm{metr}}^{2}(1+\varepsilon) t}} & \leq C_{4}\left(\frac{\sqrt{t}}{\sqrt{C_{\mathrm{metr}}^{2}(1+\varepsilon) t}}\right)^{\beta^{\prime}} \\
& =C_{4}\left(C_{\mathrm{metr}}^{2}(1+\varepsilon)\right)^{-\beta^{\prime} / 2}
\end{aligned}
$$

Applying these inequalities, we get:
$C_{1}^{\prime}:=C_{1} C_{3} C_{2}^{-\beta / 2}$
and
$C_{\varepsilon}^{\prime}:=C_{\varepsilon} C_{4}\left(C_{\mathrm{metr}}^{2}(1+\varepsilon)\right)^{-\beta^{\prime} / 2}$
; we obtain:
$C_{1}^{\prime} K_{C_{2} t}^{1, \mathrm{appr}}(p) \leq K_{t}^{1}(p) \leq C_{\varepsilon}^{\prime} K_{C_{\text {metr }}^{2}(1+\varepsilon) t}^{1, \text { appr }}(p)$.
Reparametrizing $t$ in both inequalities gives:

$$
\frac{1}{C_{\varepsilon}^{\prime}} K_{t /\left(C_{\operatorname{metr}}^{2}(1+\varepsilon)\right)}^{1}(p) \leq K_{t}^{1, \text { appr }}(p) \leq \frac{1}{C_{1}^{\prime}} K_{C_{2}^{-1} t}^{1}(p)
$$

Finally, we fix $\varepsilon>0$ and relabel constants:
$C:=\max \left\{C_{1}^{\prime-1}, C_{\varepsilon}^{\prime}, 1\right\}$,
$D_{1}:=\frac{1}{C_{\mathrm{metr}}^{2}(1+\varepsilon)}$,
$D_{2}:=\frac{1}{C_{2}} ;$


Fig. 10 Comparing the numerically computed heat kernel $K_{t}^{1}$ (left) with our approximation $K_{t}^{1, \text { appr }}$ based on the logarithmic norm estimate (right) for $G / H=S E$ (2). Shown here at $t=1$ with the same metric as in Fig. 9. Especially in deep learning applications where discretization is very coarse, our approximation is sufficiently accurate as long as the spatial anisotropies $w_{M} / w_{L}$ and $w_{L} / w_{M}$ do not become too high. In this case, with $w_{L} / w_{M}=2$ we have a relative $L^{2}$ error of 0.23 in the plotted volume
observe that since $\varepsilon>0$ and $C_{\text {metr }} \geq 1$, we have $0<D_{1}<$ 1.

Depending on the actually achievable constants, Lemma 24 provides a very strong or very weak bound on how much our approximation deviates from the fundamental solution. Fortunately in the $S E(2)$ case, our approximation is very close to the exact kernel in the vicinity of the origin, as shown in Fig. 10. In our experiments, we sample the kernel on a grid around the origin; hence, this approximation is good for reasonable values of the metric parameters, which we may expect from Lemma 20 providing a second-order relative error.

Now let us develop an approximation for values of $\alpha$ other than 1. From semi-group theory [88], it follows that semigroups generated by taking fractional powers of the generator (in our case $\Delta_{\mathcal{G}} \rightarrow-\left(-\Delta_{\mathcal{G}}\right)^{\alpha}$ ) amounts to the following key relation between the $\alpha$-kernel and the diffusion kernel:
$K_{t}^{\alpha}(p):=\int_{0}^{\infty} q_{t, \alpha}(\tau) K_{\tau}^{1}(p) \mathrm{d} \tau$,
for $\alpha \in(0,1)$ and $t>0$ where $q_{t, \alpha}$ is defined as follows.
Definition 25 Let $\mathcal{L}^{-1}$ be the inverse Laplace transform then
$q_{t, \alpha}(\tau):=\mathcal{L}^{-1}\left(r \mapsto e^{-t r^{\alpha}}\right)(\tau) \quad$ for $\tau \geq 0$.

For explicit formulas of this kernel, see [88, Ch. IX:11 eq. 17]. Since $e^{-t r^{\alpha}}$ is positive for all $r$, it follows that $q_{t, \alpha}$ is also positive everywhere.

Now instead of integrating $K_{t}^{1}$ to obtain the exact fundamental solution, we can replace it with our approximation $K_{t}^{1, \text { appr }}$ to obtain an approximate $\alpha$-kernel.

Definition 26 (Approximate $\alpha \in(0,1)$ kernel) Akin to (40), we set $\alpha \in(0,1), t>0$ and define:
$K_{t}^{\alpha, \operatorname{appr}}(p):=\int_{0}^{\infty} q_{t, \alpha}(\tau) K_{\tau}^{1, \text { appr }}(p) \mathrm{d} \tau \geq 0$,
for $p \in G / H$.
The bounding of $K_{t}^{1}$ we obtained in Lemma 24 transfers directly to our approximation for other $\alpha$.

Theorem 27 Let $G$ be of polynomial growth and let $K_{t}^{\alpha}$ be the fundamental solution to the $\alpha \in(0,1]$ diffusion equation on $G / H$, then there exists constants $C \geq 1, D_{1} \in(0,1)$ and $D_{2}>D_{1}$ so that for all $t>0$ and $p \in G / H$ the following holds:

$$
\begin{equation*}
\frac{1}{C} K_{D_{1}^{\alpha} t}^{\alpha}(p) \leq K_{t}^{\alpha, \text { appr }}(p) \leq C K_{D_{2}^{\alpha} t}^{\alpha}(p) \tag{42}
\end{equation*}
$$

Proof This is an consequence of Lemma 24 and the fact that $q_{t, \alpha}$ is positive, applying the integral from (40) yields:

$$
\begin{aligned}
K_{t}^{\alpha, \text { appr }}(p) & =\int_{0}^{\infty} q_{t, \alpha}(\tau) K_{\tau}^{1, \text { appr }}(p) \mathrm{d} \tau \\
(\text { Lemma } 24) & \leq C \int_{0}^{\infty} q_{t, \alpha}(\tau) K_{D_{2} \tau}^{1}(p) \mathrm{d} \tau \\
\left(\tau^{\prime}=D_{2} \tau\right) & =C \int_{0}^{\infty} \frac{1}{D_{2}} q_{t, \alpha}\left(\frac{\tau^{\prime}}{D_{2}}\right) K_{\tau^{\prime}}^{1}(p) \mathrm{d} \tau^{\prime} \\
\binom{\text { Bromwich }}{\text { integral }} & =C \int_{0}^{\infty} q_{D_{2}^{\alpha} t, \alpha}\left(\tau^{\prime}\right) K_{\tau^{\prime}}^{1}(p) \mathrm{d} \tau^{\prime} \\
& =C K_{D_{2}^{\alpha} t}^{\alpha}(p)
\end{aligned}
$$

The other inequality works the same way.
Although the approximation (41) is helpful in the proof above it contains some integration and is not an explicit expression. Our initial experiments with diffusion for $\alpha=1$ showed that (for the applications under consideration at least) adding diffusion did not improve performance. For that reason, we chose not to focus further on diffusion in this work. We leave developing a more explicit and computable approximation for diffusion kernels for $0<\alpha<1$ for future work.

### 5.3 Dilation and Erosion

The dilation/erosion step solves the PDE

$$
\begin{cases}\frac{\partial W_{3}}{\partial t}(p, t)= \pm\left\|\nabla_{\mathcal{G}_{2}^{ \pm}} W_{3}(p, t)\right\|_{\mathcal{G}_{2}^{ \pm}}^{2 \alpha} & \text { for } p \in G / H  \tag{43}\\ W_{3}(p, 0)=U_{3}(p) & t \geq 0 \\ & \text { for } p \in G / H\end{cases}
$$

By a generalization of the Hopf-Lax formula [89, Ch.10. 3], the solution is given by morphological convolution
$W_{3}(p, t)=-\left(k_{t}^{\alpha} \square_{G}-U_{3}\right)(p)$
for the $(+)$ (dilation) variant and
$W_{3}(p, t)=\left(k_{t}^{\alpha} \square_{G} U_{3}\right)(p)$
for the (-) (erosion) variant, where the kernel $k_{t}^{\alpha}$ (also called the structuring element in the context of morphology) is given as follows.

Definition 28 (Dilation/erosion kernels) The morphological convolution kernel $k_{t}^{\alpha}$ for small times $t$ and $\alpha \in(1 / 2,1]$ is given by
$k_{t}^{\alpha}(p):=v_{\alpha} t^{-\frac{1}{2 \alpha-1}} d_{\mathcal{G}_{2}}\left(p_{0}, p\right)^{\frac{2 \alpha}{2 \alpha-1}}$,
with $v_{\alpha}:=\left(\frac{2 \alpha-1}{(2 \alpha)^{2 \alpha /(2 \alpha-1)}}\right)$ and for $\alpha=1 / 2$ by
$k_{t}^{1 / 2}(p)= \begin{cases}0 & \text { if } d_{\mathcal{G}_{2}}\left(p_{0}, p\right) \leq t, \\ \infty & \text { if } d_{\mathcal{G}_{2}}\left(p_{0}, p\right)>t .\end{cases}$
In the above definition and for the rest of the section, we write $\mathcal{G}_{2}$ for either $\mathcal{G}_{2}^{+}$or $\mathcal{G}_{2}^{-}$depending on whether we are dealing with the dilation or erosion variant. The morphological convolution $\square_{G}$ (alternatively: the infimal convolution) is specified as follows.

Definition 29 (Morphological group convolution) Let $f \in$ $L^{\infty}(G / H)$, let $k: G / H \rightarrow \mathbb{R} \cup\{\infty\}$ be proper (not everywhere $\infty$ ), then we define:

$$
\begin{aligned}
\left(k \square_{G} f\right)(p) & :=\inf _{g \in G}\left\{k\left(g^{-1} p\right)+f\left(g p_{0}\right)\right\} \\
& =\inf _{g \in G}\left\{k\left(g^{-1} p\right)+f(g H)\right\} .
\end{aligned}
$$

Remark 12 (Grayscale morphology) Morphological convolution is related to the grayscale morphology operations $\oplus$ (dilation) and $\ominus$ (erosion) on $\mathbb{R}^{d}$ as follows:
$f_{1} \oplus f_{2}=-\left(-f_{1} \square_{\mathbb{R}^{d}}-f_{2}\right)$,
$f_{1} \ominus f_{2}=f_{1} \square_{\mathbb{R}^{d}}\left[x \mapsto-f_{2}(-x)\right]$,
where $f_{1}$ and $f_{2}$ are proper functions on $\mathbb{R}^{d}$. Hence, our use of the terms dilation and erosion, but mathematically we will only use $\square_{G}$ as the actual operation to be performed and avoid $\oplus$ and $\ominus$.

Combining morphological convolution with the structuring element $k_{t}^{\alpha}$ allows us to solve (43).

Theorem 30 Let $G$ be of polynomial growth, let $\alpha \in(1 / 2,1]$ and let $U_{3}: G / H \rightarrow \mathbb{R}$ be Lipschitz. Then, $W_{3}: G / H \times$ $(0, \infty) \rightarrow \mathbb{R}$ given by
$W_{3}(p, t):=\left(k_{t}^{\alpha} \square_{G} U_{3}\right)(p)$
is Lipschitz and solves the (-)-variant, the erosion variant, of the system (43) in the sense of Theorem 2.1 in [90], while

$$
W_{3}(p, t):=-\left(k_{t}^{\alpha} \square_{G}-U_{3}\right)(p)
$$

is Lipschitz and solves the (+)-variant, the dilation variant, of system (43) in the sense of Theorem 2.1 in [90]. The kernels satisfy the semigroup property
$k_{t}^{\alpha} \square_{G} k_{s}^{\alpha}=k_{t+s}^{\alpha}$
for all $s, t \geq 0$ and $\alpha \in(1 / 2,1]$.
Proof The Riemannian manifold $\left(G / H, \mathcal{G}_{2}\right)$ is a proper length space, and therefore, the theory of [90] applies. Moreover, since $G$ is of polynomial growth, we have that $G / H$ is a doubling space [86, Thm. 2.17] and also admits a Poincaré constant [86, Thm. 2.18]. So we meet the additional requirements of [90, Thm. 2.3 (vii) and (viii)].

The Hamiltonian $\mathcal{H}: \mathbb{R}_{+} \rightarrow \mathbb{R}_{+}$in [90] is given by $\mathcal{H}(x)=x^{2 \alpha}$. This Hamiltonian is indeed superlinear, convex, and satisfies $\mathcal{H}(0)=0$. The corresponding Lagrangian $\mathcal{L}$ : $\mathbb{R}_{+} \rightarrow \mathbb{R}_{+}$becomes
$\mathcal{L}(x)=v_{\alpha} x^{\frac{2 \alpha}{2 \alpha-1}}$.
According to [90], the solution (in the sense of their Theorem 2.1 ) to the $(-)$-variant of system (43) is given by

$$
\begin{aligned}
W_{3}(p, t) & =\inf _{x \in G / H}\left\{t \mathcal{L}\left(\frac{d_{\mathcal{G}_{2}}(p, x)}{t}\right)+U_{3}(x)\right\} \\
& =\inf _{g \in G}\left\{t \mathcal{L}\left(\frac{d_{\mathcal{G}_{2}}\left(p, g p_{0}\right)}{t}\right)+U_{3}\left(g p_{0}\right)\right\} \\
& =\inf _{g \in G}\left\{t \mathcal{L}\left(\frac{d_{\mathcal{G}_{2}}\left(g^{-1} p, p_{0}\right)}{t}\right)+U_{3}\left(g p_{0}\right)\right\} \\
& =\inf _{g \in G}\left\{v_{\alpha} \frac{d_{\mathcal{G}_{2}}\left(g^{-1} p, p_{0}\right)^{\frac{2 \alpha}{2 \alpha-1}}}{t^{\frac{2 \alpha}{2 \alpha-1}-1}}+U_{3}\left(g p_{0}\right)\right\} \\
& =\inf _{g \in G}\left\{v_{\alpha} t^{1-\frac{2 \alpha}{2 \alpha-1}} d_{\mathcal{G}_{2}}\left(g^{-1} p, p_{0}\right)^{\frac{2 \alpha}{2 \alpha-1}}+U_{3}\left(g p_{0}\right)\right\} \\
& =\inf _{g \in G}\left\{v_{\alpha} t^{\frac{-1}{2 \alpha-1}} d_{\mathcal{G}_{2}}\left(g^{-1} p, p_{0}\right)^{\frac{2 \alpha}{2 \alpha-1}}+U_{3}\left(g p_{0}\right)\right\} \\
& =\left(k_{t}^{\alpha} \square_{G} U_{3}\right)(p) .
\end{aligned}
$$

The ( + )-variant is proved analogously.
The semigroup property follows directly from [90, Thm 2.1(ii)].

Remark 13 (Solution according to Balogh et al.) This theorem builds on the work by Balogh et al. [90] who provide a solution concept that is (potentially) different from the strong, weak or viscosity solution. The point of departure is to replace the norm of the gradient (i.e., the dual norm of the differential) with a metric subgradient, i.e., we replace $\left\|\nabla_{\mathcal{G}_{2}} W(p, t)\right\|_{\mathcal{G}_{2}}$ by:
$\underset{p^{\prime} \rightarrow p}{\lim \sup } \frac{\max \left(W(p, t)-W\left(p^{\prime}, t\right), 0\right)}{d_{\mathcal{G}_{2}}\left(p, p^{\prime}\right)}$,
and we get a solution concept in terms of this slightly different notion of a gradient.

Remark 14 Unique viscosity solutions For the case $\alpha=1 / 2$, we lose the superlinearity of the Hamiltonian and can no longer apply Balogh et al.'s approach [90]. The solution for $\alpha>1 / 2$ (46) converges pointwise to the solution for $\alpha=1 / 2$ (47) as $\alpha \downarrow^{1} / 2$. However, the solution concept changes from that of Balogh et al. to that of a viscosity solution [91,92]. In the general Riemannian homogeneous space setting, the result by Azagra [92, Thm 6.24] applies. It states that viscosity solutions of Eikonal PDEs on complete Riemannian manifolds are given by the distance map departing from the boundary of a given open and bounded set. As Eikonal equations directly relate to geodesically equidistant wavefront propagation on manifolds ( [93, ch. 3], [29, ch. 4,app. E], [89]) one expects that the solutions (44),(45) of (43) are indeed the viscosity solutions (for resp. the + and - -case) for $\alpha=1 / 2$.

In many matrix Lie group quotients, like the Heisenberg group $H(2 d+1)$ studied in [94], or in our case of interest: the homogeneous space $\mathbb{M}_{d}$ of positions and orientations) this is indeed the case. One can describe $G$-invariant vector fields via explicit coordinates and transfer HJB systems on $G / H$ directly toward HJB-systems on $\mathbb{R}^{n}$ or $\mathbb{R}^{d} \times S^{q}$, with $n=d+q=\operatorname{dim}(G / H)$. Then, one can directly apply results by Dragoni [91, Thm.4] and deduce that our solutions, the dilations in (44) resp. erosions in (45), are indeed the unique viscosity solutions of HJB-PDE system (43) for the + and -case, for all $\alpha \in[1 / 2,1]$. Details are left for future research.

To get an idea of how the kernel in (46) operates in conjunction with morphological convolution, we take $G=$ $G / H=\mathbb{R}$ and see how the operation evolves simple data, the kernels and results at $t=1$ are shown in Fig. 12. Observe that with $\alpha$ close to $1 / 2$ (kernel and result in red) we obtain what amounts to an equivariant version of max $/ \mathrm{min}$ pooling.

The level sets of the kernels $k_{t}^{\alpha}$ for $\alpha>1 / 2$ are of the same shape as for the approximate diffusion kernels, see Fig. 11; for $\alpha=1 / 2$ these are the stencils over which we would perform $\mathrm{min} / \mathrm{max}$ pooling.


Fig. 11 Shapes of the level sets of the kernels on $\mathbb{M}_{2}$ for solving fractional diffusion $\left(K_{t}^{\alpha}\right)$ and dilation/erosion $\left(k_{t}^{\alpha}\right)$ for various values of the trainable metric tensor field parameters $w_{M}, w_{L}$ and $w_{A}$. This shape is essentially what is being optimized during the training process of a metric tensor field on $\mathbb{M}_{2}$

Remark 15 The level sets in Fig. 11 are balls in $G / H=\mathbb{M}_{2}$ that do not depend on $\alpha$. It is only the growth of the kernel values when passing through these level sets that depends on $\alpha$. As such, the example $G / H=\mathbb{R}$ and Fig. 12 is very representative to the general $G / H$ case. In the general $G / H$ case, Fig. 11 still applies when one replaces the horizontal $\mathbb{R}$-axis with a signed distance along a minimizing geodesic in $G / H$ passing through the origin. In that sense, $\alpha \in[1 / 2,1]$ regulates soft-max pooling over Riemannian balls in $G / H$.

We can now define a more tangible approximate kernel by again replacing the exact metric $d_{\mathcal{G}_{2}}$ with the logarithmic approximation $\rho_{\mathcal{G}_{2}}$.

Definition 31 (Approximate dilation/erosion kernel) The approximate morphological convolution kernel $k_{t}^{\alpha, \text { appr }}$ for small times $t$ and $\alpha \in(1 / 2,1]$ is given by
$k_{t}^{\alpha, \text { appr }}(p):=v_{\alpha} t^{-\frac{1}{2 \alpha-1}} \rho_{\mathcal{G}_{2}}(p)^{\frac{2 \alpha}{2 \alpha-1}}$,
with $v_{\alpha}:=\left(\frac{2 \alpha-1}{(2 \alpha)^{2 \alpha /(2 \alpha-1)}}\right)$ and for $\alpha=1 / 2$ by
$k_{t}^{1 / 2, \text { appr }}(p)=\left\{\begin{array}{ll}0 & \text { if } \rho_{\mathcal{G}_{2}}(p) \leq t, \\ \infty & \text { if } \rho_{\mathcal{G}_{2}}(p)>t\end{array}\right.$.
We used this approximation in our parallel GPU-algorithms (for our PDE-G-CNNs experiments in Sect. 7). It is highly preferable over the 'exact' solution based on the true distance as this would require Eikonal PDE solvers ( $[29,95]$ which would not be practical for parallel GPU implementations of


Fig. 12 In the center, we have kernels of the type (46) in $\mathbb{R}$ (or the signed distance on a manifold of choice) for some $\alpha \in(1 / 2,1]$ and $t=1$, which solves dilation/erosion. For $\alpha \rightarrow 1 / 2$, this kernel converges to the type in (47), i.e., the solution is obtained by max/min pooling. On the left, we morphologically convolve a spike (in gray) with a few of these ker-
nels; we see that if $\alpha \rightarrow 1 / 2$ we get max pooling; conversely, we can call the case $\alpha>1 / 2$ soft max pooling. On the right, we similarly erode a plateau, which for $\alpha \rightarrow{ }^{1} / 2$ yields min pooling. The effects of these operations in the image processing context can also be seen in the last two columns of Fig. 6

## 6 Generalization of (Group-)CNNs

In this section, we point out the similarities between common (G-)CNN operations and our PDE-based approach. Our goal here is not so much claiming that our PDE approach serves as a useful model for analyzing (G-)CNNs, but that modern CNNs already bear some resemblance to a network of PDE solvers. Noticing that similarity, our approach is then just taking the next logical step by structuring a network to explicitly solve a set of PDEs.

### 6.1 Discrete Convolution as Convection and Diffusion

Now that we have seen how PDE-G-CNNs are designed we show how they generalize conventional G-CNNs. Starting with an initial condition $U$, we show how group convolution with a general kernel $k$ can be interpreted as a superposition of solutions (27) of convection PDEs:

$$
\begin{aligned}
& \left(k *_{G / H} U\right)(p) \\
& \quad=\frac{1}{\mu_{G}(H)} \int_{G} k\left(g^{-1} p\right) U\left(g p_{0}\right) \mathrm{d} \mu_{G}(g) \\
& \quad=\frac{1}{\mu_{G}(H)} \int_{G} k\left(g^{-1} g_{p} p_{0}\right) U\left(g p_{0}\right) \mathrm{d} \mu_{G}(g),
\end{aligned}
$$

for any $g_{p} \in p$, now change variables to $q=g_{p}^{-1} g$ and recall that $\mu_{G}$ is left invariant:
$=\frac{1}{\mu_{G}(H)} \int_{G} k\left(q^{-1} p_{0}\right) U\left(g_{p} q p_{0}\right) \mathrm{d} \mu_{G}(q)$.
In this last expression, we recognize (27) and see that we can interpret $p \mapsto U\left(g_{p} q p_{0}\right)$ as the solution of the convection PDE (25) at time $t=1$ for a convection vector field $\boldsymbol{c}$ that has flow lines given by $\gamma_{c}(t)=\exp _{G}\left(-t \log _{G} q\right) p_{0}$ so that $\left(\gamma_{\mathbf{c}}(1)\right)^{-1} p_{0}=q p_{0}$. As a result, the output $k *_{G / H} U$ can
then be seen as a weighted sum of solutions over all possible left invariant convection vector fields.

Using this result, we can consider what happens in the discrete case where we take the kernel $k$ to be a linear combination of displaced diffusion kernels $K_{t}^{\alpha}$ (for some choice of $\alpha$ ) as follows:
$k(p)=\sum_{i=1}^{n} k_{i} K_{t_{i}}^{\alpha}\left(g_{i}^{-1} p\right)$,
where for all $i$ we fix a weight $k_{i} \in \mathbb{R}$, diffusion time $t_{i} \geq 0$ and a displacement $g_{i} \in G$. Convolving with this kernel yields:

$$
\begin{aligned}
& \left(k *_{G / H} U\right)(p) \\
& \quad=\int_{G} \sum_{i=1}^{n} k_{i} K_{t_{i}}^{\alpha}\left(g_{i}^{-1} g^{-1} p\right) U\left(g p_{0}\right) \mathrm{d} \mu_{G}(g) \\
& \quad=\sum_{i=1}^{n} k_{i} \int_{G} K_{t_{i}}^{\alpha}\left(g_{i}^{-1} g^{-1} p\right) U\left(g p_{0}\right) \mathrm{d} \mu_{G}(g),
\end{aligned}
$$

we change variables to $h=g g_{i}$ :

$$
\begin{aligned}
& =\sum_{i=1}^{n} k_{i} \int_{G} K_{t_{i}}^{\alpha}\left(h^{-1} p\right) U\left(h g_{i}^{-1} p_{0}\right) \mathrm{d} \mu_{G}(h) \\
& =\sum_{i=1}^{n} k_{i}\left(K_{t_{i}}^{\alpha} *_{G / H}\left[q \mapsto U\left(g_{q} g_{i}^{-1} p_{0}\right)\right]\right)(p)
\end{aligned}
$$

Here again we recognize $q \mapsto U\left(g_{q} g_{i}^{-1} p_{0}\right)$ as the solution (27) of the convection PDE at $t=1$ with flow lines induced by $\gamma_{c}(t)=\exp _{G}\left(t \log _{G} g_{i}\right)$. Subsequently, we take these solutions and convolve them with a (fractional) diffusion kernel with scale $t_{i}$, i.e., after convection we apply the fractional diffusion PDE with evolution time $t_{i}$ and finally make a linear combination of the results.

We can conclude that G-CNNs fit in our PDE-based model by looking at a single discretized group convolution as a set of single-step PDE units working on an input, without the morphological convolution and with specific choices made for the convection vector fields and diffusion times.

### 6.2 Max Pooling as Morphological Convolution

The ordinary max pooling operation commonly found in convolutional neural networks can also be seen as a morphological convolution with a kernel for $\alpha=1 / 2$.

Proposition 34 (Max pooling) Let $f \in L^{\infty}(G / H)$; let $S \subset$ $G / H$ be nonempty and define $k_{S}: G / H \rightarrow \mathbb{R} \cup\{\infty\}$ as:
$k_{S}(p):= \begin{cases}0 & \text { if } p \in S, \\ \infty & \text { else } .\end{cases}$
Then,

$$
\begin{equation*}
-\left(k_{S} \square-f\right)(p)=\sup _{g \in G: g^{-1} p \in S} f\left(g p_{0}\right) \tag{52}
\end{equation*}
$$

We can recognize the morphological convolution as a generalized form of max pooling of the function $f$ with stencil $S$.

Proof Filling in (51) into Definition 29 yields:

$$
\begin{aligned}
- & \left(k_{S} \square-f\right)(p) \\
= & -\inf \left\{\inf _{g \in G: g^{-1}}-f\left(g p_{0}\right), \inf _{g \in G: g^{-1} p \notin S}\right. \\
& \left.\times-f\left(g p_{0}\right)+\infty\right\} \\
= & -\inf _{g \in G: g^{-1} p \in S}-f\left(g p_{0}\right) \\
= & \sup _{g \in G: g^{-1} p \in S} f\left(g p_{0}\right) .
\end{aligned}
$$

In particular cases, we recover a more familiar form of max pooling as the following corollary shows.

Corollary 35 (Euclidean Max Pooling) Let $G=G / H=\mathbb{R}^{n}$ and let $f \in C^{0}\left(\mathbb{R}^{n}\right)$ with $S \subset \mathbb{R}^{n}$ compact then:
$-\left(k_{S} \square \mathbb{R}^{n}-f\right)(x)=\max _{y \in S} f(x-y)$,
for all $x \in \mathbb{R}^{n}$.
The observation that max pooling is a particular limiting case of morphological convolution allows us to think of the case with $\alpha>1 / 2$ as a soft variant of max pooling, one that is better behaved under small perturbations in a discretized context.

### 6.3 ReLUs as Morphological Convolution

Max pooling is not the only common CNN operation that can be generalized by morphological convolution as the following proposition shows.

Proposition 36 Let $f$ be a compactly supported continuous function on $G / H$. Then, dilation with the kernel
$k_{\mathrm{ReLU}, f}(p):= \begin{cases}0 & \text { if } p=p_{0}, \\ \sup _{q \in G / H} f(q) & \text { else, }\end{cases}$
equates to applying a Rectified Linear Unit to the function $f$ :
$-\left(k_{R e L U, f} \square-f\right)(p)=\max \{0, f(p)\}$.
Proof Filling in $k$ into the definition of morphological convolution:

$$
\begin{aligned}
- & \left(k_{\operatorname{ReLU}, f} \square_{G}-f\right)(p) \\
= & -\inf _{g \in G} k_{\operatorname{ReLU}}\left(g^{-1} p\right)-f\left(g \cdot p_{0}\right) \\
= & -\inf _{g \in G}\left\{\inf _{g^{-1} p=p_{0}}-f\left(g p_{0}\right)\right. \\
& \left.\times \inf _{g^{-1} p \neq p_{0}}-f\left(g p_{0}\right)+\sup _{y \in G / H} f(y)\right\} \\
= & \sup \left\{f(p), \sup _{z \in G / H: z \neq p} f(z)-\sup _{y \in G / H} f(y)\right\}
\end{aligned}
$$

due to the continuity and compact support of $f$ its supremum exists, and moreover, we have $\sup _{z \in G / H: z \neq p_{0}} f(z)=$ $\sup _{y \in G / H} f(y)$, and thereby, we obtain the required result
$=\max \{f(p), 0\}$.

We conclude that morphological convolution allows us to:

- do pooling in an equivariant manner with transformations other then translation,
- do soft pooling that is continuous under domain transformations (illustrated in Fig. 12),
- learn the pooling region by considering the kernel $k$ as trainable,
- effectively fold the action of a ReLU into trainable nonlinearities.


### 6.4 Residual Networks

So-called residual networks [67] were introduced mainly as a means of dealing with the vanishing gradient problem in very deep networks, aiding trainability. These networks use so-called residual blocks, illustrated in Fig. 13, that feature a skip connection to group a few layers together to produce a delta-map that gets added to the input.

This identity + delta structure is very reminiscent of a forward Euler discretization scheme. If we had an evolution equation of the type
$\begin{cases}\frac{\partial U}{\partial t}(p, t)=\mathcal{F}(U(\cdot, t), p) & \text { for } p \in M, t \geq 0, \\ U(p, 0)=U_{0}(p) & \text { for } p \in M,\end{cases}$


Fig. 13 A residual block, like in [67], note the resemblance to a forward Euler discretization scheme
with some operator $\mathcal{F}: L^{\infty}(M) \times M \rightarrow \mathbb{R}$, we could solve it approximately by stepping forward with:
$U(p, t+\Delta t)=U(p, t)+\Delta t \mathcal{F}(U(\cdot, t), p)$,
for some time step $\Delta t>0$. We see that this is very similar to what is implemented in the residual block in Fig. 13 once we discretize it.

The correspondence is far from exact given that multiple channels are being combined in residual blocks, so we can not easily describe a residual block with a PDE. Still, our takeaway is that residual networks and skip connections have moved CNNs from networks that change data to networks that evolve data.

For this reason, we speculate that deep PDE-G-CNNs will not need (or have the same benefit from) skip connections; we leave this subject for future investigation. More discussion on the relation between residual networks and PDEs can be found in [77].

## 7 Experiments

To demonstrate the viability of PDE-based CNNs, we perform two experiments where we compare the performance of PDE-G-CNNs against G-CNNs and classic CNNs. We will be doing a vessel segmentation and digit classification problem: two straightforward applications of CNNs. Examples of these two applications are illustrated in Fig. 14.

The goal of the experiments is to compare the basic building blocks of different types of networks in clearly defined feed-forward network architectures. So we test networks of

(a) Example from the DRIVE [96] dataset, showing a retinal image and its vessel segmentation.

(b) Examples from the RotNIST [97] dataset.

Fig. 14 We perform a segmentation experiment on retinal vessel images and a classification experiment on rotation augmented digits
modest size only and do not just aim for the performance that would be possible with large-scale networks.

### 7.1 Implementation

We implemented our PDE-based operators in an extension to the PyTorch deep learning framework [98]. Our package is called LieTorch and is open source. It is available at https:// gitlab.com/bsmetsjr/lietorch.

The operations we have proposed in the paper have been implemented in C++ for CPUs and CUDA for Nvidia GPUs but can be used from Python through PyTorch. Our package was also designed with modularity in mind: we provide a host of PyTorch modules that can be used together to implement the PDE-G-CNNs we proposed, but that can also be used separately to experiment with other architectures.

All the modules we provide are differentiable and so our PDE-G-CNNs are trainable through stochastic gradient descent (or its many variants) in the usual manner. In our experiments, we have had good results with using the ADAM [99] optimizer.

All the network models and training scripts used in the experiments are also available in the repository.

### 7.2 Design Choices

Several design choices are common to both experiments, we will go over these now.

First, we choose $G / H=\mathbb{M}_{2}$ for our G-CNNs and PDE-G-CNNs and so go for roto-translation equivariant networks. In all instances, we lift to 8 orientations.

Second, we use the convection, dilation and erosion version of (24); hence, we refer to these networks as PDE-CNNs of the CDE-type. Each PDE-layer is implemented as in Fig. 1 with the single-pass PDE solver from Fig. 2 without the convolution. So no explicit diffusion is used, and the layer consists of just resampling and two morphological convolutions. Since we do the resampling using trilinear interpolation, this does introduce a small amount of implicit diffusion.

Remark 16 (Role of diffusion) In these experiments, we found no benefit to adding diffusion to the networks. Diffusion likely would be of benefit when the input data is noisy but neither datasets we used are noisy and we have not yet performed experiments with adding noise. We leave this investigation for future work.

Third, we fix $\alpha=0.65$. We came to this value empirically; the networks performed best with $\alpha$-values in the range $0.6-$ 0.7. Looking at Fig. 12, we can conjecture that $\alpha=0.65$ is the "sweet spot" between sharpness and smoothness. When the kernel is too sharp ( $\alpha$ close to $1 / 2$ ), minor perturbations in the input can have large effects on the output; when the kernel is too smooth ( $\alpha$ close to 1 ) the output will be smoothed out too much as well.

Fourth, all our networks are simple feed-forward networks.

Finally, we use the ADAM optimizer [99] together with $L^{2}$ regularization uniformly over all parameters with a factor of 0.005 .

### 7.3 DRIVE Retinal Vessel Segmentation

The first experiment uses the DRIVE retinal vessel segmentation dataset [96]. The object is the generate a binary mask indicating the location of blood vessels from a color image of a retina as illustrated in Fig. 14a.

We test 6- and 12-layer variants of a CNN, a G-CNN and a CDE-PDE-CNN. The layout of the 6-layer networks is shown in Fig. 15; the 12-layer networks simply add more convolution, group convolution or CDE layers. All the networks were trained on the same training data and tested on the same testing data.

The output of the network is passed through a sigmoid function to produce a 2D map $a$ of values in the range [0,1] which we compare against the known segmentation map $b$


CDE-PDE-CNN 6


Fig. 15 Schematic of the 6-layer models used on our segmentation experiments. Kernel sizes and number of feature channels in each layer are indicated, depth indicates that the data lives on $\mathbb{M}_{2}$. Omitted are activation functions, batch normalization, padding and dropout modules. The 12-layer models are essentially the same but with double the number of layers but with reduced number of channels per layer (i.e., reduced width) for the CDE-PDE-CNN (hence the reduction in parameters going from 6 to 12 layers)
with values in $\{0,1\}$. We use the continuous DICE coefficient as the loss function:
$\operatorname{loss}(a, b)=1-\frac{2 \sum a b+\varepsilon}{\sum a+\sum b+\varepsilon}$,
where the sum $\sum$ is over all the values in the 2 D map. A relatively small $\varepsilon=1$ is used to avoid divide-by-zero issues and the $a \equiv b \equiv 0$ edge case.

The 6-layer networks were trained over 60 epochs, starting with a learning rate of 0.01 that we decay exponentially with a gamma of 0.95 . The 12-layer networks were trained over 80 epochs, starting from the same learning rate but with a learning rate gamma of 0.96 .

We measure the performance of the network by the DICE coefficient obtained on the 20 images of the testing dataset. We trained each model 10 times, the results of which are summarized in Table 1 and Fig. 16a.

We achieve similar or better performance than CNNs or GCNNs but with a vast reduction in parameters. Scaling from 6 to 12 layers even allows us to reduce the total number of parameters of the PDE-G-CNN while still increasing performance; this is achieved by reducing the number of channels (i.e., the width) of the network, see also Table 2.

### 7.4 RotNIST Digit Classification

The second experiment we performed is the classic digit classification experiment. Instead of using the plain MNIST dataset, we did the experiment on the RotNIST dataset [97].

Table 1 Average DICE coefficient achieved on the 20 images of the testing dataset and the number of trainable parameters of each model. The G-CNNs and CDE-PDE-CNNs are roto-translation equivariant by construction. Note the vast reduction in parameters allowed by using PDE-based networks

| Model | Parameters | DICE score $\pm$ std.dev. |
| :--- | :--- | :--- |
| CNN 6 | 47352 | $0.8058 \pm 0.0017$ |
| G-CNN 6 | 39258 | $0.8085 \pm 0.0022$ |
| CDE-PDE-CNN 6 | 4128 | $0.8115 \pm 0.0018$ |
| CNN 12 | 129432 | $0.8189 \pm 0.0005$ |
| G-CNN 12 | 114378 | $0.8192 \pm 0.0012$ |
| CDE-PDE-CNN 12 | 3678 |  |

(a) Performance on retinal vessel segmentation. We test 6- and 12-layer variants of conventional CNNs, G-CNNs and our PDECNNs, each network is trained 10 times, the chart shows the distribution of DICE performances on the test dataset.

(b) Performance of digit classification on the RotNIST dataset. We compare the classic 5-layer LeNet against a 4-layer G-CNN and PDE-CNN. LeNet was trained for 120 epochs, the other two for 60 epochs.

Fig. 16 Comparison of PDE-based networks against conventional CNNs and group CNNs on segmentation and classification tasks

Table 2 Allocation of parameters for the 6- and 12-layer CDE-PDE-CNNs used in the vessel segmentation experiment. The added depth of the networks allows us to shrink the width. With the network having less channels over all we can also shrink the number of channels in the lifting layer, which drastically reduces the total number of parameters

| Type of parameter | CDE-PDE-CNN 6 | CDE-PDE-CNN 12 |
| :--- | :--- | :--- |
| Lifting layer | 2352 | 1470 |
| Convection | 192 | 300 |
| Dilation | 192 | 300 |
| Erosion | 192 | 300 |
| Linear combinations | 1040 | 1076 |
| Batch normalization | 160 | 232 |



Fig. 17 Schematic of the three models tested with the RotNIST data. Kernel sizes and number of feature channels in each layer are indicated. Omitted are activation functions, batch normalization and dropout modules

RotNIST contains the same images as MNIST but rotated to various degrees. Even though classifying rotated digits is a fairly artificial problem, we include this experiment to show that PDE-G-CNNs also work in a context very different from the first segmentation experiment. While our choice of PDEs derives from more traditional image processing methods, this experiment shows their utility in a basic image classification context.

We tested three networks: the classic LeNet5 CNN [100] as a baseline, a 4-layer G-CNN and a 4-layer CDE-PDECNN. The architectures of these three networks are illustrated in Fig. 17.

All three networks were trained on the same training data and tested on the same testing data. We train with a learning rate of 0.05 and a learning rate gamma of 0.96 . We trained the LeNet5 model for 120 epochs and the G-CNN and CDE-PDE-CNN models for 60 epochs.

We measure the performance of the network by its accuracy on the testing dataset. We trained each model 10 times, the results of which are summarized in Table 3 and Fig. 16b.

We manage to get better performance than classic or group CNNs with far fewer parameters.

### 7.5 Computational Performance

Care was taken in optimizing the implementation to show that PDE-based networks can still achieve decent running times despite their higher computational complexity. In Table 4, we summarized the inferencing performance of each model we experimented with.

Our approach simultaneously gives us equivariance, a decrease in parameters and higher performance but at the cost of an increase in flops and memory footprint. While our implementation is reasonably optimized, it has had far less development time dedicated to it than the traditional CNN implementation provided by PyTorch/cuDNN, so we are confident more performance gains can be found.

In comparison with G-CNNs, our PDE-based networks are generally a little bit faster. Our G-CNN implementation is, however, less optimized compared to out PDE-G-CNN implementation. Were our G-CNN implementation equally optimized we expect G-CNNs to be slightly faster than the PDE-G-CNNs in our experiments.

Table 3 Accuracy of the digit classification models on the testing dataset and number of parameters for each model

| Model | Parameters | Error rate $\pm$ std.dev. |
| :--- | :--- | :--- |
| CNN (LeNet5) | 44426 | $2.59 \% \pm 0.66 \%$ |
| G-CNN Classifier 4 | 12700 | $1.14 \% \pm 0.21 \%$ |
| CDE-PDE-CNN Classifier 4 | 2542 | $1.10 \% \pm 0.10 \%$ |

Table 4 Time in seconds it took to run each model on the testing dataset of its respective experiment. The DRIVE testing dataset contains 20 images while the RotNIST testing dataset contains 10000 digits

|  | CNN | G-CNN | PDE-CNN |
| :--- | :--- | :--- | :--- |
| DRIVE 6-layer | 1.7 s | 6.5 s | 6.8 s |
| DRIVE 12-layer | 2.2 s | 14.1 s | 9.8 s |
| RotNIST | 0.1 s | 0.9 s | 0.7 s |

## 8 Conclusion

In this article, we presented the general mathematical framework of geometric PDEs on homogeneous spaces that underlies our PDE-G-CNNs. PDE-G-CNNs allow for a geometric and probabilistic interpretation of CNNs opening up new avenues for the study and development of these types of networks. We showed that additionally, PDE-G-CNNs have increased performance with a reduction in parameters.

PDE-G-CNNs ensure equivariance by design. The trainable parameters are geometrically relevant: they are leftinvariant vector and tensor fields.

PDE-G-CNNs have three types of layers: convection, diffusion and erosion/dilation layers. We have shown that these layers implicitly include standard nonlinear operations in CNNs such as max pooling and ReLU activation.

To efficiently evaluate PDE evolution in the layers, we provided tangible analytical approximations to the relevant kernel operators on homogeneous spaces. In this article, we have underpinned the quality of the approximations in Theorems 27 and 30.

With two experiments we have verified that PDE-G-CNNs can improve performance over G-CNNs in the context of automatic vessel segmentation and digit classification. Most importantly, the performance increase is achieved with a vast reduction in the amount of trainable parameters.

Acknowledgements We gratefully acknowledge the Dutch Foundation of Science NWO for financial support (Duits: Geometric learning for Image Analysis, VI.C 202-031, Bekkers: Context Aware AI, VI nr.17290).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecomm ons.org/licenses/by/4.0/.

## Appendix A

## A. 1 Proof of Lemma 20

The left inequality follows directly from the observation that $\rho_{\mathcal{G}}(p)$ is exactly the Riemannian length of the curve
$t \mapsto \exp _{G}\left(t \log _{G}\left(g_{p}\right)\right) p_{0}$
for $t \in[0,1]$ and $g_{p}=\arg \min _{g \in p}\left\|\log _{G} g\right\|_{\tilde{\mathcal{G}}}$. This continuous curve connects $p_{0}$ with $p$ and as such has a greater length than the minimal length curve between those two points.

For the right inequality, consider the function $F: T_{e} G \rightarrow$ $\mathbb{R}$ given by
$F(v):=d_{\mathcal{G}}\left(p_{0}, \pi \circ \exp _{G}(v)\right)^{2}$,
where we recall that $\pi: G \rightarrow G / H$ was given by (4). With the goal of making a Taylor expansion for this function, we note that:

- at the origin we have $F(0)=0$,
- due to the chain rule applied to the squaring we have $\left.d F\right|_{0}=0$.

Moreover, due to the to the $G$-invariance of $d_{\mathcal{G}}$, the function $F$ is even and consequently the 3rd order term of the Taylor expansion of $F$ is zero.

For the second-order term, we are looking for the Hessian $\mathcal{H}$ of $F$ at $v=0$. We split $F$ into $F_{1}:=\pi \circ \exp _{G}$ and $F_{2}(p):=d_{\mathcal{G}}\left(p_{0}, p\right)^{2}$ and find the Hessian of the composed function is

$$
\begin{aligned}
\mathcal{H} & \left.\left(F_{2} \circ F_{1}\right)\right|_{0}(v, w) \\
& =\left.\mathcal{H} F_{2}\right|_{p_{0}}\left(\left.d F_{1}\right|_{0}(v),\left.d F_{1}\right|_{0}(w)\right) \\
& =\left.2 \mathcal{G}\right|_{p_{0}}\left(\left.\left.d \pi\right|_{e} \circ d \exp _{G}\right|_{0}(v),\left.\left.d \pi\right|_{e} \circ d \exp _{G}\right|_{0}(w)\right) \\
& =\left.2 \tilde{\mathcal{G}}\right|_{e}\left(\left.d \exp _{G}\right|_{0}(v),\left.d \exp _{G}\right|_{0}(w)\right) \\
& =\left.2 \tilde{\mathcal{G}}\right|_{e}(v, w)
\end{aligned}
$$

Putting these facts together, we find:
$F(v)=\left.\tilde{\mathcal{G}}\right|_{e}(v, v)+O\left(\|v\|^{4}\right)$,
where $\|\cdot\|$ denotes some arbitrary norm on $T_{e} G$.
Now we take a linear subspace $V$ of $T_{e} G$ that is independent from $T_{e} H$ but so that the span of $T_{e} H$ and $V$ equals the entire $T_{e} G$, so that $T_{e} H \oplus V=T_{e} G$. Note that $\left.\tilde{\mathcal{G}}\right|_{e}$ is only degenerated along $T_{e} H$, and so is a full norm when restricted to $V$, i.e., for all $v \in V,\left.\tilde{\mathcal{G}}\right|_{e}(v, v)=0$ only if $v=0$. Therefore, there exists a $c>0$ such that for all $v \in V$,
$\left.\tilde{\mathcal{G}}\right|_{e}(v, v)>c\|v\|^{2}$,
and so by (53) we have:
$F(v)=d_{\mathcal{G}}\left(p_{0}, \pi \circ \exp _{G}(v)\right)^{2}>c\|v\|^{2}+O\left(\|v\|^{4}\right)$.
Hence, for all $v \in V$ close enough to 0 we have:
$d_{\mathcal{G}}\left(p_{0}, \pi \circ \exp _{G}(v)\right)^{2}>\frac{c}{2}\|v\|^{2}$.
In a neighborhood of the origin, the Lie group exponential $\operatorname{map} \exp _{G}: T_{e} G \rightarrow G$ is a diffeomorphism to a neighborhood of $e$; at the same time $\pi: G \rightarrow G / H$ is a smooth submersion by the Homogeneous Space Construction Theorem [101, Thm 21.17]. Consequently, $\left.d\left(\pi \circ \exp _{G}\right)\right|_{0}: V \rightarrow$ $T_{p_{0}}(G / H)$ has full rank since $\pi \circ \exp _{G}$ is a local diffeomorphism between two spaces ( $V$ and $G / H$ ) with the same dimension; it follows by the inverse function theorem that there exists a neighborhood $V_{0}$ of 0 in $V$ and a neighborhood $P_{0}$ of $p_{0}$ in $G / H$ such that $\pi \circ \exp _{G}$ is a diffeomorphism from $V_{0}$ to $P_{0}$. By possibly choosing $V_{0}$ smaller, we may assume by (54) that there exists a $C^{\prime}>0$ such that for all $v \in V_{0}$ :
$\left.\tilde{\mathcal{G}}\right|_{e}(v, v) \leq F(v)+C^{\prime}\|v\|^{4}$,
which, by using (55) yields
$\left.\tilde{\mathcal{G}}\right|_{e}(v, v) \leq d_{\mathcal{G}}\left(p_{0}, \pi \circ \exp _{G}(v)\right)^{2}+C d_{\mathcal{G}}\left(p_{0}, \pi \circ \exp _{G}(v)\right)^{4}$,
for all $v \in V_{0}$ and $C=C^{\prime} \frac{c}{2}>0$.
Now take a $p \in P_{0}$, then there exists a $w \in V_{0}$ so that
$\pi \circ \exp _{G} w=p$.

Call $g_{p}=\exp _{G} w$, then the previous inequality gives
$\left\|\log _{G} g_{p}\right\|_{\tilde{\mathcal{G}}}^{2} \leq d_{\mathcal{G}}\left(p, p_{0}\right)^{2}+C d_{\mathcal{G}}\left(p, p_{0}\right)^{4}$.
Clearly $g_{p} \in p$. Since $\rho_{\mathcal{G}}(p)$ is the infinum of $\left\|\log _{G} g\right\|$ for all $g \in p$, it follows that $\rho_{\mathcal{G}}(p)$ must also satisfy:
$\rho_{\mathcal{G}}(p)^{2} \leq\left\|\log _{G} g_{p}\right\|_{\tilde{\mathcal{G}}}^{2} \leq d_{\mathcal{G}}\left(p, p_{0}\right)^{2}+C d_{\mathcal{G}}\left(p, p_{0}\right)^{4}$,
for all $p \in P_{0}$, i.e., all $p$ sufficiently close to $p_{0}$.
As a corollary, we get that for any compact neighborhood
$K \subset G / H$ of $p_{0}$
$\rho_{\mathcal{G}}(p) \leq C_{\text {metr }} d_{\mathcal{G}}\left(p_{0}, p\right)$
for all $p \in K$. We can see this by choosing $C_{1}^{2}=1+$ $C \sup _{p \in P)} d_{\mathcal{G}}\left(p_{0}, p\right)^{2}$, then for all by choosing $p \in P_{0}$ we have
$\rho_{\mathcal{G}}(p) \leq C_{1} d_{\mathcal{G}}\left(p_{0}, p\right)$.

Let $K \subset G / H$ be compact so that it contains $P_{0}$. Then, on $\overline{K \backslash P_{0}}$ we have that both $\rho_{\mathcal{G}}$ and $d_{\mathcal{G}}\left(p_{0}, \cdot\right)$ are strictly positive, continuous and so bounded functions. Consequently,
$\rho_{\mathcal{G}}(p) \leq \sup _{p \in K \backslash P_{0}} \rho_{\mathcal{G}}(p)=: M<\infty$,
and
$d_{\mathcal{G}}\left(p_{0}, p\right) \geq \sup _{p \in P_{0}} d_{\mathcal{G}}\left(p_{0}, p\right)=: m>0$,
for all $p \in K \backslash P_{0}$, which leads to
$\rho_{\mathcal{G}}(p) \leq M \leq \frac{M}{d_{\mathcal{G}}\left(p_{0}, p\right)} d_{\mathcal{G}}\left(p_{0}, p\right) \leq \frac{M}{m} d_{\mathcal{G}}\left(p_{0}, p\right)$.
for all $p \in K \backslash P_{0}$. Now choose $C_{\text {metr }}=\max \left\{C_{1}, C_{2}\right\}$, then we obtain the corollary. Remark that $C_{\text {metr }}$ depends on both the parameters of the metric tensor field and the choice of $K$ and so may become very large indeed.

## A. 2 Proof of Proposition 21

As a preliminary, we prove the following lemma.
Lemma For all $g \in G$, let $L_{g}: G \rightarrow G$ be the left group multiplication given by $L_{g} h=g h$ and let $R_{g}: G \rightarrow G$ be the right group multiplication given by $R_{g} h=h g$. Let $H$ be a closed subgroup of $G$ with the projection map $\pi: G \rightarrow$ $G / H$ given by $\pi(g)=g H$.

Then, for all $h \in H$ we have the following relations for the push forwards:

1. $\pi_{*} \circ\left(R_{h}\right)_{*}=\pi_{*}$,
2. $\left(L_{h}\right)_{*} \circ \pi_{*}=\pi_{*} \circ\left(L_{h}\right)_{*}$.

Proof 1. $\pi \circ R_{h}=\pi$ since $g h H=g H$,
2. $L_{h} \circ \pi=\pi \circ L_{h}$ since $h(g H)=(h g) H$.

Now for the proof of Proposition 21. Consider the set of all exponential curves in the group whose action connects $p_{0} \in G / H$ to $p \in G / H$ :

$$
\begin{aligned}
\Gamma_{p_{0}, p} & =\{\gamma \in \operatorname{Lip}([0,1], G) \mid \\
\gamma(0) & \left.=e, \gamma(1) p_{0}=p, \gamma(t+s)=\gamma(t) \gamma(s)\right\}
\end{aligned}
$$

We can then restate $\rho_{\mathcal{G}}$ equivalently in terms of these curves as

$$
\rho_{\mathcal{G}}(p):=\inf _{g \in P}\left\|\log _{G} g\right\|_{\tilde{\mathcal{G}}}=\inf _{\gamma \in \Gamma_{p_{0}, p}}\|\dot{\gamma}(0)\|_{\tilde{\mathcal{G}}}
$$

since for each $g \in p$ we have an exponential curve $t \mapsto$ $\exp _{G}\left(t \log _{G} g\right)$ in $\Gamma_{p_{0}, p}$, and for each exponential curve $\gamma$ in $\Gamma_{p_{0}, p}$ we have $\gamma(1) \in p$.

Let $\gamma \in \Gamma_{p_{0}, p}$ and let $h \in H$, then

1. $h \gamma(0) h^{-1}=h e h^{-1}=e$,
2. $h \gamma(1) h^{-1} p_{0}=h \gamma(1) p_{0}=h p$,
3. $h \gamma(a+b) h^{-1}=h \gamma(a) \gamma(b) h^{-1}=\left(h \gamma(a) h^{-1}\right)$ $\left(h \gamma(b) h^{-1}\right)$,
from which we conclude that $h \gamma(\cdot) h^{-1} \in \Gamma_{p_{0}, h p}$ and so there is a bijection between $\Gamma_{p_{0}, p}$ and $\Gamma_{p_{0}, h p}$ given by

$$
\Gamma_{p_{0}, h p}=h \Gamma_{p_{0}, p} h^{-1}
$$

Moreover, the bijection preserves the seminorm due to the G-invariance of $\mathcal{G}$ :

$$
\begin{aligned}
\left\|\left(h \gamma(\cdot) h^{-1}\right)(0)\right\|_{\tilde{\mathcal{G}}} & =\left\|\left(L_{h}\right)_{*}\left(R_{h^{-1}}\right)_{*} \dot{\gamma}(0)\right\|_{\tilde{\mathcal{G}}} \\
& =\left\|\pi_{*}\left(L_{h}\right)_{*}\left(R_{h^{-1}}\right)_{*} \dot{\gamma}(0)\right\|_{\mathcal{G}}
\end{aligned}
$$

(using the previous lemma)

$$
\begin{aligned}
& =\left\|\left(L_{h}\right)_{*} \pi_{*}\left(R_{h^{-1}}\right)_{*} \dot{\gamma}(0)\right\|_{\mathcal{G}} \\
& =\left\|\left(L_{h}\right)_{*} \pi_{*} \dot{\gamma}(0)\right\|_{\mathcal{G}}
\end{aligned}
$$

(using the G-invariance of $\mathcal{G}$ )

$$
\begin{aligned}
& =\left\|\pi_{*} \dot{\gamma}(0)\right\|_{\mathcal{G}} \\
& =\|\dot{\gamma}(0)\|_{\tilde{\mathcal{G}}} .
\end{aligned}
$$

It follows that

$$
\begin{aligned}
\rho_{\mathcal{G}}(p) & =\inf _{\gamma \in \Gamma_{p_{0}, p}}\|\dot{\gamma}(0)\|_{\tilde{\mathcal{G}}} \\
& =\inf _{\gamma \in \Gamma_{p_{0}, p}}\left\|h \dot{\gamma}(0) h^{-1}\right\|_{\tilde{\mathcal{G}}} \\
& =\inf _{\gamma \in \Gamma_{p_{0}, h p}}\|\dot{\gamma}(0)\|_{\tilde{\mathcal{G}}} \\
& =\rho_{\mathcal{G}}(h p) .
\end{aligned}
$$

## References

1. Welk, M., Weickert, J.: PDE evolutions for M-smoothers: from common myths to robust numerics. In: International Conference on Scale Space and Variational Methods in Computer Vision, pp. 236-248. Springer (2019)
2. Fadili, J., Kutyniok, G., Peyré, G., Plonka-Hoch, G., Steidl, G.: Guest editorial: mathematics and image analysis. J. Math. Imaging Vis. 52(3), 315-316 (2015)
3. Peyré, G., Péchaud, M., Keriven, R., Cohen, L.D.: Geodesic methods in computer vision and graphics. Found. Trends. Comput. Graph. Vis. 5(3), 197-397 (2010)
4. Dubrovina-Karni, A., Rosman, G., Kimmel, R.: Multi-region active contours with a single level set function. IEEE PAMI 37(8), 1585-1601 (2015)
5. Burger, M., Sawatzky, A., Steidl, G.: First Order Algorithms in Variational Image Processing. Springer, Cham (2016)
6. Duran, J., Moeller, M., Sbert, C., Cremers, D.: Collaborative total variation: a general framework for vectorial tv models. SIAM SIIMS 9(1), 116-151 (2016)
7. Weickert, J., Grewenig, S., Schroers, C., Bruhn, A.: Cyclic schemes for PDE-based image analysis. Int. J. Comput. Vis. 118(3), 275-299 (2016)
8. Sapiro, G.: Geometric Partial Differential Equations and Image Analysis. Cambridge University Press, Cambridge (2001). https:// doi.org/10.1017/CBO9780511626319
9. Sethian, J.: Level Set Methods and Fast Marching Methods. Cambridge University Press, Cambridge (1999)
10. Weickert, J.: Theoretical foundations of anisotropic diffusion in image processing. Comput. Suppl. 11, 221-236 (1996)
11. Morel, J., Solimini, S.: Variational Methods in Image Segmentation: With Seven Image Processing Experiments. Progress in Nonlinear Differential Equations and Their Applications. Birkhäuser, Basel (1995)
12. Duits, R., Burgeth, B.: Scale spaces on Lie groups. In: International Conference on Scale Space and Variational Methods in Computer Vision, pp. 300-312. Springer (2007)
13. Cohen, T.S., Welling, M.: Group equivariant convolutional networks. In: International of Conference on Machine Learning, pp. 2990-2999 (2016)
14. Dieleman, S., De Fauw, J., Kavukcuoglu, K.: Exploiting cyclic symmetry in convolutional neural networks (2016). arXiv preprint arXiv:1602.02660
15. Dieleman, S., Willett, K.W., Dambre, J.: Rotation-invariant convolutional neural networks for galaxy morphology prediction. Mon. Not. R. Astron. Soc. 450(2), 1441-1459 (2015)
16. Winkels, M., Cohen, T.S.: 3D G-CNNs for pulmonary nodule detection (2018). arXiv preprint arXiv: 1804.04656
17. Worrall, D., Brostow, G.: Cubenet: Equivariance to 3D rotation and translation. In: Proceedings of the European Conference on Computer Vision (ECCV), pp. 567-584 (2018)
18. Bekkers, E.J., Loog, M., ter Haar Romeny, B.M., Duits, R.: Template matching via densities on the roto-translation group. IEEE Trans. Pattern Anal. Mach. Intell. 40(2), 452-466 (2017)
19. Oyallon, E., Mallat, S.: Deep roto-translation scattering for object classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2865-2873 (2015)
20. Bekkers, E.J., Lafarge, M.W., Veta, M., Eppenhof, K.A., Pluim, J.P., Duits, R.: Roto-translation covariant convolutional networks for medical image analysis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 440-448. Springer (2018). arXiv:1804.03393
21. Weiler, M., Hamprecht, F.A., Storath, M.: Learning steerable filters for rotation equivariant cnns. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 849-858 (2018)
22. Cohen, T.S., Geiger, M., Weiler, M.: A general theory of equivariant cnns on homogeneous spaces. In: Advances in Neural Information Processing Systems, vol. 32 (2019)
23. Worrall, D.E., Garbin, S.J., Turmukhambetov, D., Brostow, G.J.: Harmonic networks: deep translation and rotation equivariance. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5028-5037 (2017)
24. Kondor, R., Trivedi, S.: On the generalization of equivariance and convolution in neural networks to the action of compact groups.

In: Dy, J., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 80, pp. 2747-2755. PMLR, Stockholmsmässan, Stockholm Sweden (2018). http://proceedings.mlr. press/v80/kondor18a.html
25. Esteves, C., Allen-Blanchette, C., Makadia, A., Daniilidis, K.: Learning SO (3) equivariant representations with spherical CNNs. In: Proceedings of the European Conference on Computer Vision (ECCV), pp. 52-68 (2018)
26. Akian, M., Quadrat, J.P., Viot, M.: Bellman processes. In: 11th International Conference on Analysis and Optimization of Systems Discrete Event Systems, pp. 302-311. Springer (1994)
27. Burgeth, B., Welk, M., Feddern, C., Weickert, J.: Morphological operations on matrix-valued images. In: European Conference on Computer Vision, pp. 155-167. Springer (2004)
28. Duits, R., Haije, T.D., Creusen, E., Ghosh, A.: Morphological and linear scale spaces for fiber enhancement in DW-MRI. J. Math. Imaging Vis. 46(3), 326-368 (2013)
29. Bekkers, E.J., Duits, R., Mashtakov, A., Sanguinetti, G.R.: A PDE approach to data-driven sub-Riemannian geodesics in $\mathrm{SE}(2)$. SIAM J. Imaging Sci. 8(4), 2740-2770 (2015)
30. Haije, T.C.D., Duits, R., Tax, C.M.: Sharpening fibers in diffusion weighted MRI via erosion. In: Visualization and Processing of Tensors and Higher Order Descriptors for Multi-valued Data, pp. 97-126. Springer (2014)
31. Duits, R., Smets, B., Bekkers, E.J., Portegies, J.M.: Equivariant deep learning via morphological and linear scale space PDEs on the space of positions and orientations. LNCS 12679, 27-39 (2021)
32. Citti, G., Franceschiello, B., Sanguinetti, G., Sarti, A.: SubRiemannian mean curvature flow for image processing. SIAM J. Imaging Sci. 9(1), 212-237 (2016)
33. Chambolle, A., Pock, T.: A first-order primal-dual algorithm for convex problems with applications to imaging. J. Math. Imaging Vis. 40(1), 120-145 (2011)
34. Chambolle, A., Pock, T.: Total roto-translational variation. Numer. Math. 142(3), 611-666 (2019). https://doi.org/10.1007/ s00211-019-01026-w
35. Smets, B.M., Portegies, J., St-Onge, E., Duits, R.: Total variation and mean curvature PDEs on the homogeneous space of positions and orientations. J. Math. Imaging Vis. 63(2), 237-262 (2021)
36. Duits, R., Duits, M., van Almsick, M., ter Haar Romeny, B.: Invertible orientation scores as an application of generalized wavelet theory. Pattern Recognit. Image Anal. 17(1), 42-75 (2007)
37. Janssen, M.H.J., Janssen, A.J.E.M., Bekkers, E.J., Bescós, J.O., Duits, R.: Design and processing of invertible orientation scores of 3D images. J. Math. Imaging Vis. $\mathbf{6 0}(9), 1427-1458$ (2018). https://doi.org/10.1007/s 10851-018-0806-0
38. Franceschiello, B., Mashtakov, A., Citti, G., Sarti, A.: Geometrical optical illusion via sub-Riemannian geodesics in the roto-translation group. Differ. Geom. Appl. 65, 55-77 (2019)
39. Citti, G., Sarti, A.: A cortical based model of perceptional completion in the roto-translation space. J. Math. Imaging Vis. 24(3), 307-326 (2006)
40. Duits, R., Franken, E.M.: Left invariant parabolic evolution equations on $\mathrm{SE}(2)$ and contour enhancement via invertible orientation scores, part I: linear left-invariant diffusion equations on $\mathrm{SE}(2)$. Q. Appl. Math. 68, 255-292 (2010)
41. Duits, R., Franken, E.M.: Left invariant parabolic evolution equations on $\mathrm{SE}(2)$ and contour enhancement via invertible orientation scores, part II: nonlinear left-invariant diffusion equations on invertible orientation scores. Q. Appl. Math. 68, 293-331 (2010)
42. Zhang, J., Duits, R., ter Haar Romeny, B., Sanguinetti, G.: Numerical approaches for linear left-invariant diffusions on $\mathrm{SE}(2)$, their comparisons to exact solutions, and their applications in retinal imaging. Numer. Math. Theory Methods Appl. 9(1), 1-50 (2016)
43. Boscain, U., Chertovskih, R.A., Gauthier, J.P., Remizov, A.O.: Hypoelliptic diffusion and human vision: a semidiscrete new twist. SIAM J. Imaging Sci. 7(2), 669-695 (2014)
44. Bertalmío, M., Calatroni, L., Franceschi, V., Franceschiello, B., Prandi, D.: A cortical-inspired model for orientation-dependent contrast perception: a link with Wilson-Cowan equations. In: Lellmann, J., Burger, M., Modersitzki, J. (eds.) Scale Space and Variational Methods in Computer Vision, pp. 472-484. Springer, Cham (2019)
45. Duits, R., Fuehr, H., Janssen, B., Florack, L., van Assen, H.: Evolution equations on Gabor transforms and their applications. ACHA 35(3), 483-526 (2013)
46. Barbieri, D., Citti, G., Cocci, G., Sarti, A.: A cortical-inspired geometry for contour perception and motion integration. J. Math. Imaging Vis. 49(3), 511-529 (2014)
47. Petitot, J.: The neurogeometry of pinwheels as a sub-Riemannian contact structure. J. Physiol. Paris 97, 265-309 (2003)
48. Felsberg, M., Forssen, P.E., Scharr, H.: Channel smoothing: efficient robust smoothing of low-level signal features. IEEE Trans. Pattern Anal. Mach. Intell. 28, 209-222 (2006)
49. Savadjiev, P., Strijkers, G., Bakermans, A., Piuze, E., Zucker, S., Siddiqi, K.: Heart wall myofibers are arranged in minimal surfaces to optimize organ function. PNAS 109(24), 9248-9253 (2012)
50. Duits, R., Bekkers, E., Mashtakov, A.: Fourier transform on the homogeneous space of 3D positions and orientations for exact solutions to linear PDEs. Entropy 21(1), 38 (2019)
51. Momayyez-Siahkal, P., Siddiqi, K.: 3D stochastic completion fields for fiber tractography. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (2009). https://doi.org/10.1109/CVPRW.2009.5204044
52. Duits, R., Führ, H., Janssen, B., Bruurmijn, M., Florack, L., van Assen, H.: Evolution equations on Gabor transforms and their applications. Appl. Comput. Harmon. Anal. 35(3), 483-526 (2013)
53. Boscain, U., Prandi, D., Sacchelli, L., Turco, G.: A bio-inspired geometric model for sound reconstruction. J. Math. Neurosci. 11(1), 1-18 (2021)
54. Baspinar, E., Citti, G., Sarti, A.: A geometric model of multiscale orientation preference maps via Gabor functions. J. Math. Imaging Vis. 60(6), 900-912 (2018)
55. Bekkers, E.J.: B-Spline CNNs on Lie groups. In: International Conference on Learning Representations (2019)
56. Finzi, M., Stanton, S., Izmailov, P., Wilson, A.G.: Generalizing convolutional neural networks for equivariance to Lie groups on arbitrary continuous data. In: III, H.D., Singh, A. (eds.) Proceedings of the 37th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 119, pp. 31653176. PMLR (2020). http://proceedings.mlr.press/v119/finzi20a. html
57. Weiler, M., Cesa, G.: General E(2)-equivariant steerable CNNs. In: Advances in Neural Information Processing Systems, pp. 14334-14345 (2019)
58. Chirikjian, G.S., Kyatkin, A.B.: An operational calculus for the Euclidean motion group with applications in robotics and polymer science. J. Fourier Anal. Appl. 6(6), 583-606 (2000)
59. Franken, E., van Almsick, M., Rongen, P., Florack, L., ter Haar Romeny, B.: An efficient method for tensor voting using steerable filters. In: European Conference on Computer Vision, pp. 228-240. Springer (2006)
60. Reisert, M.: Group integration techniques in pattern analysis: a kernel view. Ph.D. thesis, Albert-Ludwigs-Universität Freiburg (2008)
61. Ali, S.T., Antoine, J.P., Gazeau, J.P., et al.: Coherent States, Wavelets and Their Generalizations, vol. 1. Springer, Berlin (2000)
62. Sifre, L., Mallat, S.: Rigid-motion scattering for texture classification (2014). arXiv preprint arXiv:1403.1687
63. Finzi, M., Bondesan, R., Welling, M.: Probabilistic numeric convolutional neural networks (2020). arXiv preprint arXiv:2010.10876
64. Montobbio, N.: A metric model of the visual cortex. Ph.D. thesis, Università di Bologna - Sorbonne Université (2019)
65. Montobbio, N., Bonnasse-Gahot, L., Citti, G., Sarti, A.: KerCNNs: biologically inspired lateral connections for classification of corrupted images (2019). arXiv preprint arXiv: 1910.08336
66. Weinan, E.: A proposal on machine learning via dynamical systems. Commun. Math. Stat. 5(1), 1-11 (2017)
67. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770-778 (2016)
68. Lu, Y., Zhong, A., Li, Q., Dong, B.: Beyond finite layer neural networks: bridging deep architectures and numerical differential equations (2017). arXiv preprint arXiv:1710.10121
69. Chen, T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. In: Advances in Neural Information Processing Systems, pp. 6571-6583 (2018)
70. Chen, Y., Yu, W., Pock, T.: On learning optimized reaction diffusion processes for effective image restoration. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5261-5269 (2015)
71. Long, Z., Lu, Y., Ma, X., Dong, B.: PDE-net: Learning PDEs from data (2017). arXiv preprint arXiv:1710.09668
72. Ruthotto, L., Haber, E.: Deep neural networks motivated by partial differential equations. J. Math. Imaging Vis. 62(3), 352-364 (2020)
73. Shen, Z., He, L., Lin, Z., Ma, J.: Pdo-econvs: partial differential operator based equivariant convolutions. In: International Conference on Machine Learning, pp. 8697-8706. PMLR (2020)
74. Weiler, M., Forré, P., Verlinde, E., Welling, M.: Coordinate independent convolutional networks-isometry and gauge equivariant convolutions on Riemannian manifolds (2021). arXiv preprint arXiv:2106.06020
75. Jenner, E., Weiler, M.: Steerable partial differential operators for equivariant neural networks (2021). arXiv preprint arXiv:2106.10163
76. Haber, E., Ruthotto, L.: Stable architectures for deep neural networks. Inverse Prob. 34(1), 014004 (2017)
77. Alt, T., Peter, P., Weickert, J., Schrader, K.: Translating numerical concepts for PDEs into neural architectures (2021). arXiv preprint arXiv:2103.15419
78. Koda, T.: An introduction to the geometry of homogeneous spaces. In: Proceedings of the Thirteenth International Workshop on Differential Geometry, vol. 13, pp. 121-144 (2009)
79. Lee, J.M., Chow, B., Chu, S.C., Glickenstein, D., Guenther, C., Isenberg, J., Ivey, T., Knopf, D., Lu, P., Luo, F., et al.: Manifolds and differential geometry. Topology 643, 658 (2009)
80. Smets, B.: Geometric image denoising and machine learning. Master thesis, TU Eindhoven (2019). https://bmnsmets.com/ publication/smets $2019 \mathrm{msc} /$
81. Arendt, W., Bukhvalov, A.V.: Integral representations of resolvents and semigroups. Forum Math. 6, 111-136 (1994)
82. Führ, H.: Abstract Harmonic Analysis of Continuous Wavelet Transforms. 1863. Springer, Berlin (2005)
83. Duits, R., Meesters, S.P., Mirebeau, J.M., Portegies, J.M.: Optimal paths for variants of the 2D and 3D Reeds-Shepp car with applications in image analysis. J. Math. Imaging Vis. 60, 1-33 (2018)
84. Portegies, J., Sanguinetti, G., Meesters, S., Duits, R.: 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, pp. 40-52. Springer (2015)
85. Ter Elst, A., Robinson, D.W.: Weighted subcoercive operators on Lie groups. J. Funct. Anal. 157(1), 88-163 (1998)
86. Maheux, P.: Estimations du noyau de la chaleur sur les espaces homogenes. ESAIM Control Optim. Calc. Var. 8(1), 65-96 (1998)
87. Grigor'yan, A., Hu, J., Lau, K.S.: Heat kernels on metric spaces with doubling measure. In: Fractal Geometry and Stochastics IV, pp. 3-44. Springer (2009)
88. Yosida, K.: Functional Analysis. Springer, Berlin (1968)
89. Evans, L.C.: Partial Differential Equations, vol. 19. American Mathematical Society, Providence (2010)
90. Balogh, Z.M., Engulatov, A., Hunziker, L., Maasalo, O.E.: Functional inequalities and Hamilton-Jacobi equations in geodesic spaces. Potential Anal. 36(2), 317-337 (2012)
91. Dragoni, F.: Metric Hopf-Lax formula with semicontinuous data. Discrete Contin. Dyn. Syst. A 17(4), 713 (2007)
92. Azagra, D., Ferrera, J., López-Mesas, F.: Nonsmooth analysis and Hamilton-Jacobi equations on Riemannian manifolds. J. Funct. Anal. 220(2), 304-361 (2005)
93. Rund, H.: The Hamilton-Jacobi Theory in the Calculus of Variations: Its Role in Mathematics and Physics. Krieger Publishing Company, Malabar (1966)
94. Manfredi, J.J., Stroffolini, B.: A version of the Hopf-Lax formula in the Heisenberg group. Commun. Partial Differ. Equ. 27, 11391159 (2002)
95. Mirebeau, J.M., Portegies, J.: Hamiltonian fast marching: a numerical solver for anisotropic and non-holonomic eikonal PDEs. Image Process. On Line 9, 47-93 (2019)
96. Staal, J., Abràmoff, M.D., Niemeijer, M., Viergever, M.A., Van Ginneken, B.: Ridge-based vessel segmentation in color images of the retina. IEEE Trans. Med. Imaging 23(4), 501-509 (2004)
97. Baweja, C.: RotNIST (2018). https://github.com/ ChaitanyaBaweja/RotNIST
98. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: an imperative style, high-performance deep learning library. In: Advances in Neural Information Processing Systems, pp. 80248035 (2019). www.ytorch.org
99. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization (2014). arXiv preprint arXiv:1412.6980
100. LeCun, Y., Boser, B., Denker, J.S., Henderson, D., Howard, R.E., Hubbard, W., Jackel, L.D.: Backpropagation applied to handwritten zip code recognition. Neural Comput. 1(4), 541-551 (1989)
101. Lee, J.: Introduction to Smooth Manifolds. Graduate Texts in Mathematics, 2nd edn. Springer, Berlin (2013). https://doi.org/ 10.1007/978-1-4419-9982-5_1

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.


Bart M. N. Smets received his M.Sc. in applied mathematics at the TU/e in 2019. Now he is pursuing his $\mathrm{Ph} . \mathrm{D}$. under the supervision of Remco Duits and Jim Portegies, also at the TU/e. He created and teaches a course on the mathematics of neural networks. As part of his Ph.D. reasearch, he is investigating how PDEbased image analysis techniques can be utilized to improve and gain a better understanding of how neural networks work. He is interested in all the ways that mathematics can be used to understand and enchance deep learning.


Jim Portegies obtained his Ph.D. from the Courant Institute of Mathematical Sciences in 2014. Afterward, he was a postdoc at the Max Planck Institute for Mathematics in the Sciences in Leipzig. Since 2016, he is an assistant professor in mathematics in the Applied Analysis group at the Eindhoven University of Technology. Jim Portegies is a member of the Dutch Young Academy.


Erik J. Bekkers is an assistant professor in Geometric Deep Learning in the Machine Learning Lab of the University of Amsterdam. He did a post-doc in applied differential geometry at the Applied Mathematics Dept. of Technical University Eindhoven. In his Ph.D. thesis (cum laude, Biomedical Engineering, TU/e), he developed medical image analysis algorithms based on sub-Riemannian geometry in the Lie group SE(2) using mathematical principles that underlie models of human visual perception. His interest is in generalizing group convolutional networks and improving computational and representational efficiency through sparse (graphs) and adaptive learning mechanisms. Erik is a recipient of a MICCAI Young Scientist Award 2018, Philips Impact Award (MIDL 2018) and a personal VENI research grant awarded by the Dutch Research Council.


Remco Duits received his M.Sc. degree (cum laude) in Mathematics in 2001 at the TU/e, where he also received his Ph.D. (cum laude) at the Department of Biomedical Engineering. Now he is associate professor at the Department of Applied Mathematics \& Computer Science. He leads a research group 'Geometric Deep Learning and Differential Geometry' within CASA (Center for Analysis, Scientific Computing and Applications). He has organized various large conferences (6) and colloquia, coordinated the TU/e Mathematics department's research assessment (2015-2021) and serves in program committees of international conferences (e.g., SSVM, GSI). He is editor for JMIV, and he has led a special issue on 'Differential Geometry and Orientation Analysis in Image Processing.' He has been scientific reviewer for many (inter)national research proposals, conference proceedings, books and a wide range of journals. His research interests subtend analysis. differential geometry, Lie group theory, PDEs, mathematical and medical image analysis. He has received two prestigious personal grants: an ERC-StG grant (2014-2019) "Lie Group Analysis for Medical Image Processing" and an NWO-VICI grant (2021-2026) "Geometric Learning for Image Analysis" and acted as supervisor/applicant in other EU-projects (EU-MC, FP7) and TU/e (High Potential, EAISI) projects.


[^0]:    $\boxtimes$ Bart M. N. Smets
    b.m.n.smets@tue.nl

    1 Department of Mathematics and Computer Science, Cluster: CASA (Center for Analysis, Scientific Computing and Applications), Research Groups: Geometric Learning and Differential Geometry, Eindhoven University of Technology, Eindhoven, The Netherlands
    2 Applied Analysis, Eindhoven University of Technology, Eindhoven, The Netherlands
    3 Machine Learning Lab, Informatics Institute, University of Amsterdam, Amsterdam, The Netherlands

[^1]:    ${ }^{1}$ In our primary case of interest $G=S E(2)$, we have $V=$ $\left\{\sum_{k=1}^{3} c^{k} A_{k} \mid c^{3} \in[-\pi, \pi)\right\}$.

