

PDE-based Group Equivariant Convolutional Neural Networks

Bart M.N. Smets · Jim Portegies · Erik J. Bekkers · Remco Duits

July 13, 2021

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 non-linear 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. Just like for linear convolution a morphological convolution is specified by a kernel that we train in our PDE-G-CNNs. 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.

Bart M.N. Smets · Jim Portegies · Remco Duits
 Department of Mathematics and Computer Science, Eindhoven University of Technology, The Netherlands
 E-mail: b.m.n.smets@tue.nl

Erik J. Bekkers
 Informatics Institute, University of Amsterdam, The Netherlands

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 Hamilton-Jacobi type PDE. This substitution is illustrated in Fig. 1 where we retain (channel-wise) affine combinations as the means of composing feature maps.

The PDE we propose to use in this setting comes from the geometric image analysis world [1–11]. It was chosen based on the fact that it exhibits similar behaviour 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 [12–24]. 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 non-linear terms of the PDE are a key aspect of our design. Normally, morphological convolutions are considered on \mathbb{R}^d [25, 26], but when extended to Lie groups such as $SE(d)$ they have many benefits in applications (e.g. crossing-preserving flow [27] or tracking [28, 29]). Using morphological convolutions allows our network to have trainable non-linearities in lieu 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 an 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.




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).

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 Section 2. Then we introduce the needed theoretical preliminaries from Lie group theory in Section 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 Section 4 we give the overall architecture of a PDE-G-CNN 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 Section 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 Section 7, and concluding remarks in Section 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 image-processing 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 [30–33] as regularization and our framework provides a natural 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.

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 [34]. 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 $G = SE(d)$, 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 , 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 [35–49].

Remark 1 (Generality of the architecture) Although not considered here, for other Lie groups applications (e.g. frequency scores [50], velocity scores, scale-orientation scores [51]) the same structure applies, therefore we keep our theory in the general setting of homogeneous spaces G/H .

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

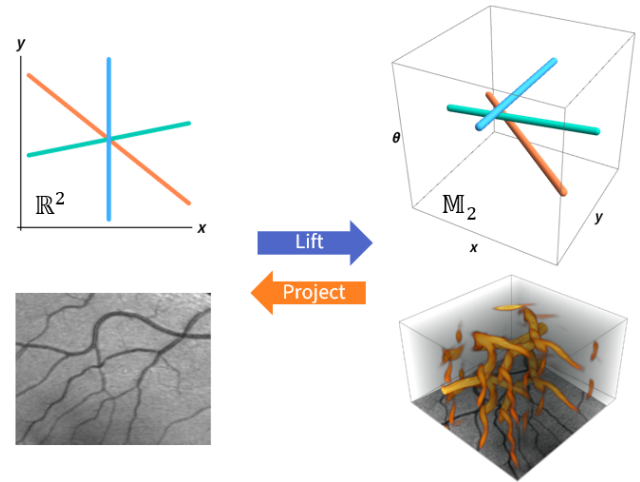


Fig. 3 Illustrating the process of lifting and projecting, in this case the advantage of lifting an image from \mathbb{R}^2 to the 2D space of positions and orientations \mathbb{M}_2 derives from the disentanglement of the lines at the crossings.

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 are reduced as the required symmetries are intrinsic to the network and need not be trained.




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 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 and so can confer the same symmetries.

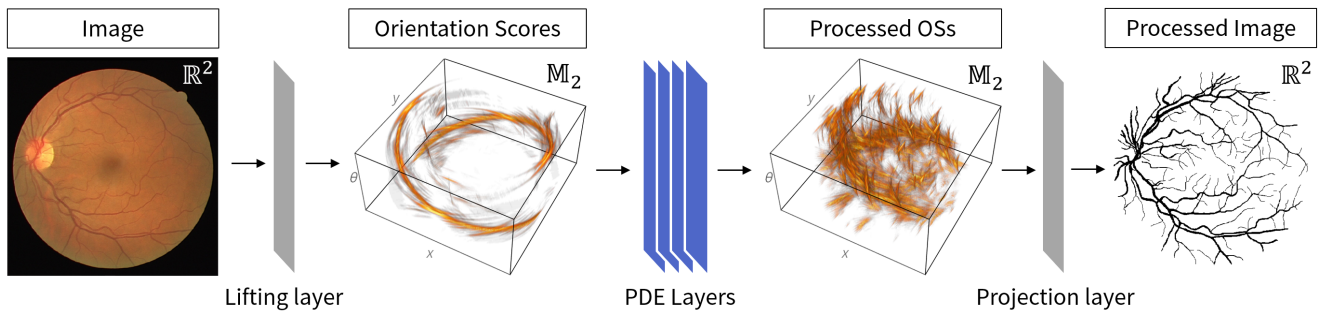


Fig. 5 Illustrating the overall architecture of a PDE-G-CNN (example: retinal vessel segmentation). An input image is lifted to a homogeneous space from which point on it can be fed through subsequent PDE layers (each PDE layer follow the structure of Fig.2) that replace the convolution layers in conventional CNNs. Finally the result is projected back to the desired output space.

2.4 Related Work

G-CNNs After the introduction of G-CNNs by Cohen & Welling [12] in the field of machine and deep learning, G-CNNs became popular. This resulted in many articles on showing the benefits of G-CNNs over classical spatial CNNs. These works can be roughly categorised as

- discrete G-CNNs [12–16],
- regular continuous G-CNNs [18–20, 28, 52, 53],
- and steerable continuous G-CNNs [21–24, 54] that rely on Fourier transforms on homogeneous spaces [48, 55].

Both regular and steerable G-CNNs naturally arise from linear mappings between functions on homogeneous spaces that are placed under equivariance constraints [21, 23, 52, 54]. 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 [56, 57], working with steerable representations puts constraints on non-linear activations within the networks which limits representation power of G-CNNs [54]. 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 [58]. 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 [59], orientation score theory [34], cortical perception models [37], G-CNNs [12, 19], and rigid-body motion scattering [60].

The proposed PDE-G-CNNs form a new, unique class of equivariant neural networks, and we show in section 6 how regular continuous G-CNNs arise as a special case of our PDE-G-CNNs.

Probabilistic-CNNs Our geometric PDEs relate to α -stable Lévy processes [48] and cost-processes akin to [25], but then on \mathbb{M}_d rather than \mathbb{R}^d . This relates to probabilistic equivariant numerical neural networks [61] that use anisotropic convection-diffusions on \mathbb{R}^d .

In contrast to 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 [25] 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. [62, 63] 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 E [64] explicitly explained the connection between neural networks and dynamical systems especially in the context of the ultradeep ResNet [65]. This point of view was further expanded by Lu et al. [66], 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. [67], 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. [68] who learn parameters in a reaction-diffusion 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 [69] with their PDE-Net.

More recent work in this direction includes integrating equivariant partial differential operators in steerable CNNs [70], 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 [71]. 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. [72].

The main contribution of our work in the field of PDE-related 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 & 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 functions defined on a space X to a space of 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. We further assume that the Lie group G acts smoothly on both 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 $g, h \in G$,

$$\rho_X(gh, x) = \rho_X(g, \rho_X(h, x))$$

and

$$\rho_Y(gh, x) = \rho_Y(g, \rho_Y(h, x)).$$

Additionally we will assume that the group G acts *transitively* on the spaces, meaning that for any two elements of the space 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* [73]. In particular, this means

that after selecting a reference element $x_0 \in X$ we can make the following isomorphism:

$$(X, G, x_0) \cong G/\text{Stab}_G(x_0) \quad (1)$$

using the mapping

$$x \mapsto \{g \in G \mid \rho_X(g, x_0) = x\}, \quad (2)$$

which is a bijection due to transitivity and the fact that $\text{Stab}_G(x_0)$ 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 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 $SE(d)$.

The elements of the quotient G/H consist of subsets of G which we will denote by the letter p , these subsets are known as left cosets of H since every one of them consists of the set $p = gH$ for some $g \in G$, the left cosets are a partition of G under the equivalence relation

$$g_1 \sim g_2 \iff g_1^{-1}g_2 \in H.$$

Under this notation the group G consists of the disjoint union

$$G = \bigsqcup_{p \in G/H} p. \quad (3)$$

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 π :

$$\pi(g) := gp_0 := gH. \quad (4)$$

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 action on a homogeneous space G/H induces an action on spaces of functions on G/H . 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 X and Y . Let Φ be an operator from functions (of some function class) on X to functions on Y , then we say that Φ is equivariant with respect to G if for all functions f (of that class) we have that:

$$\forall g \in G, y \in Y : (\Phi f)(\rho_Y(g, y)) = (\Phi [x \mapsto f(\rho_X(g, x))])(y). \quad (5)$$

Indeed, this definition expresses that one can either first apply the transformation g and then apply Φ , or first apply Φ and then apply the transformation g , and the result is the same.

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

$$L_g p := gp \quad \text{for all } p \in G/H. \quad (6)$$

In addition, we denote the left-regular representation of G on functions f defined on G/H by \mathcal{L}_g defined by

$$(\mathcal{L}_g f)(p) := f(g^{-1}p). \quad (7)$$

With this notation, condition (5) on our neural network operator Φ can be rewritten as: for all $g \in G$,

$$\boxed{\mathcal{L}_g \circ \Phi = \Phi \circ \mathcal{L}_g}, \quad (8)$$

or in words: the neural network commutes with transformations.

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 $p \in G/H$, the pushforward

$$(L_g)_* : T_p(G/H) \rightarrow T_{gp}(G/H)$$

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

$$((L_g)_* \mathbf{v}) f := \mathbf{v}(f \circ L_g). \quad (9)$$

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 (9). This algebraic viewpoint defines the action of the tangent vector $\dot{\gamma}(t)$ on a differentiable function f as

$$\dot{\gamma}(t)f := \frac{\partial}{\partial s} f(\gamma(s)) \Big|_{s=t}.$$

In the flat setting of $G = (\mathbb{R}^d, +)$, where the tangent spaces are isomorphic to the base manifold \mathbb{R}^d , when we have a tangent vector $\mathbf{c} \in \mathbb{R}^d$ its application to a function is the familiar directional derivative:

$$\mathbf{c}f = \mathbf{c} \cdot \nabla f = df(\mathbf{c}).$$

See [74, §2.1.1] for details on this double interpretation.

Vector fields that have the special property that the push forward $(L_g)_*$ maps them to themselves in the sense that

$$\forall g \in G, \forall \mathbf{v} \in T_p(G/H) : \mathbf{v}(p)f = \mathbf{v}(gp) [\mathcal{L}_g f], \quad (10)$$

for all differentiable function f and where $\mathbf{v} : p \mapsto T_p(G/H)$ is a vector field, are referred to as G -invariant.

Definition 2 (G -invariant vector field on a homogeneous space) A vector field \mathbf{v} on G/H is invariant with respect to G if it satisfies

$$\forall g \in G, \forall p \in G/H : \mathbf{v}(gp) = (L_g)_* \mathbf{v}(p). \quad (11)$$

It is straightforward to check that (10) and (11) are equivalent and that these imply the following.

Corollary 1 (Properties of left-invariant vector fields) On a homogeneous space G/H the left-invariant vector fields have the following properties:

1. they are fully determined by their value $\mathbf{v}(H)$ in H ,
2. $\forall h \in H, \forall \mathbf{v} \in T_H(G/H) : (L_h)_* \mathbf{v} = \mathbf{v}$.

We also introduce G -invariant metric tensor fields.

Definition 3 (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

$$\forall g \in G, \forall p \in G/H, \forall \mathbf{v}, \mathbf{w} \in T_p(G/H) :$$

$$\mathcal{G} \Big|_p(\mathbf{v}, \mathbf{w}) = \mathcal{G} \Big|_{gp}((L_g)_* \mathbf{v}, (L_g)_* \mathbf{w}). \quad (12)$$

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

Corollary 2 (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 $\mathcal{G} \Big|_{p_0}$ at $p_0 = H$,
2. $\forall h \in H, \forall \mathbf{v}, \mathbf{w} \in T_{p_0}(G/H) :$

$$\mathcal{G} \Big|_{p_0}(\mathbf{v}, \mathbf{w}) = \mathcal{G} \Big|_{p_0}((L_h)_* \mathbf{v}, (L_h)_* \mathbf{w}).$$

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

A metric tensor field yields a Riemannian metric in the usual manner.

Definition 4 (Metric on G/H) Let $p_1, p_2 \in G/H$ then:

$$d_{\mathcal{G}}(p_1, p_2) := d_{G/H, \mathcal{G}}(p_1, p_2) := \inf_{\substack{\beta \in \text{Lip}([0,1], G/H) \\ \beta(0)=p_1, \beta(1)=p_2}} \int_0^1 \sqrt{\mathcal{G}|_{\beta(t)}(\dot{\beta}(t), \dot{\beta}(t))} dt.$$

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

Definition 5 The cut locus $\text{cut}(p) \subset G/H$ or $\text{cut}(g) \subset G$ is the set of points respectively group elements where the geodesics starting at p respectively g are no longer minimizing.

As long as we stay away from the cut locus the infimum from Def. 4 gives a unique geodesic and a well defined metric.

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

Proposition 1 (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:

$$\forall g \in G : d_{\mathcal{G}}(p_1, p_2) = d_{\mathcal{G}}(gp_1, gp_2).$$

Proof We observe that we can make a bijection from the set of Lipschitz curves between p_1 and p_2 and between gp_1 and gp_2 simply by left multiplication by g one way and g^{-1} the other way. Due to (12) 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 gp_1 to gp_2 , both having the same length. \square

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

Definition 6 (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 π^* is the pullback of the quotient map π from (4). This is equivalent to saying that for all $\mathbf{v}, \mathbf{w} \in T_g G$:

$$\tilde{\mathcal{G}}|_g(\mathbf{v}, \mathbf{w}) := \mathcal{G}|_{\pi(g)}(\pi_* \mathbf{v}, \pi_* \mathbf{w}),$$

where π_* is the pushforward of π .

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 TG .

Definition 7 (Seminorm on TG) Let $\mathbf{v} \in T_g G$ then the metric tensor field \mathcal{G} on G/H induces the following seminorm:

$$\|\mathbf{v}\|_{\tilde{\mathcal{G}}} := \sqrt{\tilde{\mathcal{G}}|_g(\mathbf{v}, \mathbf{v})} := \sqrt{\mathcal{G}|_{g p_0}(\pi_* \mathbf{v}, \pi_* \mathbf{v})}. \quad (13)$$

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

Definition 8 (Pseudometric on G) Let $g_1, g_2 \in G$ then we define:

$$d_{\tilde{\mathcal{G}}}(g_1, g_2) := d_{G, \tilde{\mathcal{G}}}(g_1, g_2) := \inf_{\substack{\gamma \in \text{Lip}([0,1], G) \\ \gamma(0)=g_1, \gamma(1)=g_2}} \int_0^1 \sqrt{\tilde{\mathcal{G}}|_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))} dt. \quad (14)$$

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 1 Let $g_1, g_2 \in G$ so that $\pi(g_2)$ is away from the cut locus of $\pi(g_1)$, then:

$$d_{\tilde{\mathcal{G}}}(g_1, g_2) = d_{\mathcal{G}}(\pi(g_1), \pi(g_2)).$$

Moreover if γ is a minimizing geodesic in the group connecting g_1 with g_2 then $\pi \circ \gamma$ is the unique minimizing geodesic in the homogeneous space that connects $\pi(g_1)$ with $\pi(g_2)$.

Proof Let $\gamma \in \text{Lip}([0, 1], G)$ be a minimizing geodesic connecting $\gamma(0) = g_1$ with $\gamma(1) = g_2$ and let $\beta \in \text{Lip}([0, 1], G/H)$ be the unique minimizing geodesic connecting $\beta(0) = \pi(g_1)$ with $\beta(1) = \pi(g_2)$. Because of the pseudometric on G , minimizing geodesics are not unique, i.e. γ is not unique. On G/H we have a full metric and so staying away from the cut locus means β is both unique and minimizing.

Denote the length functionals with:

$$\text{Len}_G(\gamma) := \int_0^1 \sqrt{\tilde{\mathcal{G}}|_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))} dt,$$

$$\text{Len}_{G/H}(\beta) := \int_0^1 \sqrt{\mathcal{G}|_{\beta(t)}(\dot{\beta}(t), \dot{\beta}(t))} dt.$$

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

Now we assume $\pi \circ \gamma \neq \beta$. Then since β is the unique geodesic we have

$$\text{Len}_{G/H}(\beta) < \text{Len}_{G/H}(\pi \circ \gamma) = \text{Len}_G(\gamma).$$

But then we can find some $\gamma_{\text{lift}} \in \text{Lip}([0, 1], G)$ that is a preimage of β , i.e. $\pi \circ \gamma_{\text{lift}} = \beta$. The potential problem is that while $\gamma_{\text{lift}}(0) \in \pi(g_1)$ and $\gamma_{\text{lift}}(1) \in \pi(g_2)$, γ_{lift} does not necessarily connect g_1 to g_2 . But since the coset $\pi(g_1)$ is connected we can find a curve wholly contained in it that connects g_1 with $\gamma_{\text{lift}}(0)$, call this curve $\gamma_{\text{head}} \in \text{Lip}([0, 1], \pi(g_1))$. Similarly we can find a $\gamma_{\text{tail}} \in \text{Lip}([0, 1], \pi(g_2))$ that connects $\gamma_{\text{lift}}(1)$ to g_2 . Both these curves have zero length since π maps them to a single point on G/H , i.e. $\text{Len}_G(\gamma_{\text{head}}) = \text{Len}_G(\gamma_{\text{tail}}) = 0$.

Now we can compose these three curves:

$$\gamma_{\text{new}}(t) := \begin{cases} \gamma_{\text{head}}(3t) & \text{if } t \in [0, 1/3], \\ \gamma_{\text{lift}}(3t - 1) & \text{if } t \in [1/3, 2/3], \\ \gamma_{\text{tail}}(3t - 2) & \text{if } t \in [2/3, 1]. \end{cases}$$

This new curve is again in $\text{Lip}([0, 1], G)$ and connects g_1 with g_2 , but also:

$$\text{Len}_G(\gamma_{\text{new}}) = \text{Len}_G(\gamma_{\text{lift}}) = \text{Len}_{G/H}(\beta) < \text{Len}_G(\gamma),$$

which is a contradiction since γ is a minimizing geodesic between g_1 and g_2 . We conclude $\pi \circ \gamma = \beta$ and thereby:

$$d_{\mathcal{G}}(g_1, g_2) = \text{Len}_G(\gamma) = \text{Len}_{G/H}(\beta) = d_{\mathcal{G}}(\pi(g_1), \pi(g_2)).$$

□

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 $SE(d)$ and the Homogeneous Space \mathbb{M}_d

Our main example and Lie group of interest is the *Special Euclidean* group $SE(d)$ of the rotations and translations of \mathbb{R}^d , in particular for $d \in \{2, 3\}$. When we take $H = \{0\} \times SO(d-1)$ we obtain the space of positions and orientations

$$\mathbb{M}_d = SE(d) / (\{0\} \times SO(d-1)). \quad (15)$$

This homogeneous space and group will enable the construction of roto-translation equivariant networks.

As a set we identify \mathbb{M}_d with $\mathbb{R}^d \times S^{d-1}$ and choose some reference direction $\mathbf{a} \in S^{d-1} \subset \mathbb{R}^d$ as the forward direction so that we can set the reference point of the space as $p_0 = (\mathbf{0}, \mathbf{a})$. We can then see that elements of H are those rotations that map \mathbf{a} to itself, i.e. rotations with the forward direction as their axis.

If we denote elements of $SE(d)$ as translation/rotation pairs $(\mathbf{y}, R) \in \mathbb{R}^d \times SO(d)$ then group multiplication is given by

$$g_1 = (\mathbf{y}_1, R_1), g_2 = (\mathbf{y}_2, R_2) \in G :$$

$$g_1 g_2 = (\mathbf{y}_1, R_1) (\mathbf{y}_2, R_2) = (\mathbf{y}_1 + R_1 \mathbf{y}_2, R_1 R_2),$$

and the group action on elements $p = (\mathbf{x}, \mathbf{n}) \in \mathbb{R}^d \times S^{d-1} \equiv \mathbb{M}_d$ is given as

$$gp = (\mathbf{y}, R) (\mathbf{x}, \mathbf{n}) = (\mathbf{y} + R\mathbf{x}, R\mathbf{n}). \quad (16)$$

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 2 *Let $d > 2$ and let $\partial_{\mathbf{a}} \in T_{p_0}(\mathbb{M}_d)$ be the tangent vector in the reference point in the main direction \mathbf{a} , specifically:*

$$\partial_{\mathbf{a}} f := \lim_{t \rightarrow 0} \frac{f((t\mathbf{a}, \mathbf{a})) - f((\mathbf{0}, \mathbf{a}))}{t},$$

then all $SE(d)$ -invariant vector fields are spanned by the vector field:

$$p \mapsto \mathcal{A}_1|_p := (L_{g_p})_* \partial_{\mathbf{a}}, \quad (17)$$

with $g_p \in p$.

Proof For $d > 3$ we can see that (17) are the only left-invariant vector fields since for all $h \in H$ we have $(g_p h) p_0 = p$ and so in order to be well-defined we must require $(L_h)_* \mathbf{v} = \mathbf{v}$ on $T_{p_0}(\mathbb{M}_d)$, and this is true for $\partial_{\mathbf{a}}$ (and its scalar multiples) but not true for any other tangent vectors at $T_{p_0}(\mathbb{M}_d)$. □

Proposition 3 *For $d > 2$ the only Riemannian metric tensor fields on \mathbb{M}_d that are $SE(d)$ -invariant are of the form:*

$$\mathcal{G}|_{(x, \mathbf{n})}((\dot{\mathbf{x}}, \dot{\mathbf{n}}), (\dot{\mathbf{x}}, \dot{\mathbf{n}})) = w_M |\dot{\mathbf{x}} \cdot \mathbf{n}|^2 + w_L \|\dot{\mathbf{x}} \wedge \mathbf{n}\|^2 + w_A \|\dot{\mathbf{n}}\|^2, \quad (18)$$

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 Cor. 2 at the tangent space $T_{(x, \mathbf{n})}$ of a particular (x, \mathbf{n}) the metric tensor needs to be symmetric with respect to rotations about \mathbf{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 elements of \mathbb{M}_2 with $(x, y, \theta) \in \mathbb{R}^3$ where x, y are the usual Cartesian coordinates and θ the angle with respect to the x -axis, 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 Cor. 1 and 2 also become trivial. The $SE(2)$ -invariant vector fields are given as follows.

Proposition 4 *On \mathbb{M}_2 the $SE(2)$ -invariant vector fields are spanned by the following basis:*

$$\begin{cases} \mathcal{A}_1|_{(x, y, \theta)} = \cos \theta \partial_x|_{(x, y, \theta)} + \sin \theta \partial_y|_{(x, y, \theta)}, \\ \mathcal{A}_2|_{(x, y, \theta)} = -\sin \theta \partial_x|_{(x, y, \theta)} + \cos \theta \partial_y|_{(x, y, \theta)}, \\ \mathcal{A}_3|_{(x, y, \theta)} = \partial_\theta|_{(x, y, \theta)}. \end{cases} \quad (19)$$

Proof For $d = 2$ we have $\mathbb{M}_2 \equiv SE(d)$ and the group invariant vector fields on \mathbb{M}_2 are exactly the left-invariant vector fields on $SE(2)$ given by (19). \square

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

Proposition 5 *On \mathbb{M}_2 the $SE(2)$ -invariant metric tensor fields are given by:*

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

for any choice of inner product $\mathcal{G}|_{(0,0,0)}$ at e .

Proof Since $SE(2) \equiv \mathbb{M}_2$ the G -invariant metric tensor fields are again exactly the left-invariant metric tensor fields. \square

This gives $SE(2)$ -invariant metric tensor fields 6 degrees of freedom and hence 6 trainable parameters on \mathbb{M}_2 . In our experiments so far we have restricted ourselves to those metric tensors that are diagonal with respect to the frame from Prop. 4. A diagonal metric tensor would have just 3 degrees of freedom and have the same general form as (18), specifically:

$$\begin{aligned} \mathcal{G}|_{(x,y,\theta)}((\dot{x}, \dot{y}, \dot{\theta}), (\dot{x}, \dot{y}, \dot{\theta})) = & \\ & w_M \left(|\dot{x} \cos \theta|^2 + |\dot{y} \sin \theta|^2 \right) \\ & + w_L \left(|\dot{x} \sin \theta|^2 + |\dot{y} \cos \theta|^2 \right) \\ & + w_A |\dot{\theta}|^2. \end{aligned} \quad (20)$$

We will expand into non-diagonal metric tensors (to also include training of curvature) in future work.

4 Architecture

4.1 Lifting & 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 lives 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 $SE(2)$ can be found in [19, 75].

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 $f \in L^2(G/H_1)$ and $k \in L^1(G/H_2)$ with the following property:

$$\forall h \in H_1, q \in G/H_2 : k(hq) = k(q),$$

where H_1 is compact. then the operator \mathcal{T} defined by

$$\forall q \in G/H_2 : (\mathcal{T}f)(q) := \int_G k(g^{-1}q) f(gH_1) d\mu_G(g) \quad (21)$$

transforms f from a function on G/H_1 to a function on G/H_2 in an equivariant manner. Here the kernel k is the trainable part and μ_G is the left-invariant Haar measure on the group.

Moreover it can be shown via the Dunford-Pettis [76] 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 non-trainable *Orientation Score Transform* [34] or a trainable lift [19] 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

$$(x, y) \mapsto \max_{\theta \in [0, 2\pi)} f(x, y, \theta) \quad (22)$$

is a roto-translation equivariant projection as used in [19]. A variation on the above projection is detailed in [75, 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 . Afterwards, 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 ℓ and denote the width (i.e. the number of features or

channels) at layer ℓ with M_ℓ then we have M_ℓ 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 $(U_{\ell,c})_{c=1}^{M_\ell}$ be the inputs of the ℓ -th layer (i.e. some functions on G/H), let $a_{\ell ij}$ and $b_{\ell i} \in \mathbb{R}$ be the coefficients of the affine transforms for $i = 1 \dots M_{\ell+1}$ and $j = 1 \dots M_\ell$. Let each PDE be parametrized by a set of parameters $\theta_{\ell j}$. Then the action of a PDE layer is described as:

$$U_{\ell+1,i} = \sum_{j=1}^{M_\ell} a_{\ell ij} \Phi_{T, \theta_{\ell j}}(U_{\ell j}) + b_{\ell i}, \quad (23)$$

where $\Phi_{T, \theta}$ is the evolution operator of the PDE at time $T \geq 0$ and parameter set θ . We define the operator $\Phi_{T, \theta}$ so that $(p, t) \mapsto (\Phi_{T, \theta} U)(p)$ satisfies the Hamilton-Jacobi type PDE that we introduce in just a moment. In this layer formula the parameters $a_{\ell ij}$, $b_{\ell i}$ and $\theta_{\ell j}$ are the trainable weights, but the evolution time T we keep fixed.

It is essential that we require network layers, and thereby all 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

$$\left\{ \begin{array}{ll} \frac{\partial W}{\partial t}(p, t) = -\mathbf{c}W(p, t) & \text{(convection)} \\ -(-\Delta_{\mathcal{G}_1})^\alpha W(p, t) & \text{(diffusion)} \\ + \left\| \nabla_{\mathcal{G}_2^+} W(p, t) \right\|_{\mathcal{G}_2^+}^{2\alpha} & \text{(dilation)} \\ - \left\| \nabla_{\mathcal{G}_2^-} W(p, t) \right\|_{\mathcal{G}_2^-}^{2\alpha} & \text{(erosion)} \end{array} \right. \quad (24)$$

for $p \in G/H$, $t \geq 0$,

$$W(p, 0) = U(p) \quad \text{for } p \in G/H.$$

Here, \mathbf{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 sub sampling by destroying data),
- dilation: pooling of data,
- erosion: sharpening of data.

Which is also why we refer to a layer using this PDE as a CDDE layer. Summarized the parameters of this PDE are given by $\theta = (\mathbf{c}, \mathcal{G}_1, \mathcal{G}_2^+, \mathcal{G}_2^-)$. The geometric interpretation of each

of the terms in (24) is illustrated in Fig. 6 for $G = G/H = \mathbb{R}^2$ and in Fig. 7 for $G/H = \mathbb{M}_2$.

Since the convection vector field \mathbf{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 Cor. 1 and 2 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 Prop. 4.

Surprisingly for higher dimensions \mathbb{M}_d has less trainable parameters than for $d = 2$. This is caused by the $SE(d)$ -invariant vector fields on \mathbb{M}_d for $d \geq 3$ being spanned by a single basis element (17) 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 timestep-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 $t = T$.



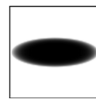


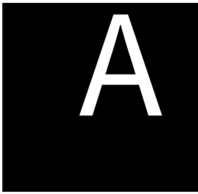



\mathbb{R}^2	Transport	Regularization	Max pooling	Min pooling
PDE Term:	$-cW$	$-(-\Delta_{G_1})^\alpha W$	$+\ \nabla_{G_2^+} W\ _{G_2^+}^{2\alpha}$	$-\ \nabla_{G_2^-} W\ _{G_2^-}^{2\alpha}$
Parameters:	transport vector $\begin{pmatrix} 100 \\ 120 \end{pmatrix}$	Riemannian metric tensor $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$	Riemannian metric tensor $\begin{pmatrix} 9 & 0 \\ 0 & 1 \end{pmatrix}$	Riemannian metric tensor $\begin{pmatrix} 5 & 11 \\ 11 & 40 \end{pmatrix}$
Operation:	resample with offset 	convolution with kernel 	dilation with kernel 	erosion with kernel 
				

Fig. 6 Geometric interpretation of the terms of the PDE (24) illustrated for \mathbb{R}^2 . In this setting the G -invariant vector field 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 2×2 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.

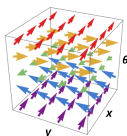
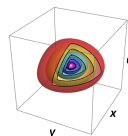
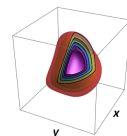
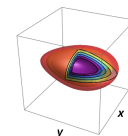
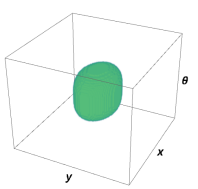
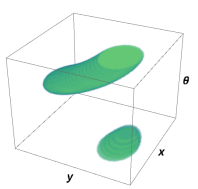
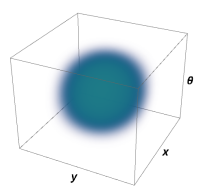
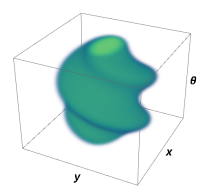
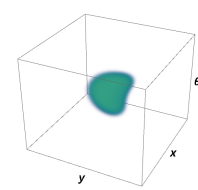
\mathbb{M}_2	Transport	Regularization	Max pooling	Min pooling
PDE Term:	$-cW$	$-(-\Delta_{G_1})^\alpha W$	$+\ \nabla_{G_2^+} W\ _{G_2^+}^{2\alpha}$	$-\ \nabla_{G_2^-} W\ _{G_2^-}^{2\alpha}$
Parameters:	G -invariant vector field $\begin{pmatrix} 1 \\ 0 \\ 2 \end{pmatrix}$	Riemannian metric tensor $\begin{pmatrix} 4 & 0 & 0 \\ 0 & 6 & 0 \\ 0 & 0 & 1 \end{pmatrix}$	Riemannian metric tensor $\begin{pmatrix} 4 & 0 & 0 \\ 0 & 9 & 0 \\ 0 & 0 & 0.5 \end{pmatrix}$	Riemannian metric tensor $\begin{pmatrix} 6 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 1.5 \end{pmatrix}$
Operation:	resample 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 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 3×3 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.

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 Δt , 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

$$\begin{cases} \frac{\partial W^1}{\partial t}(p, t) = -c(p)W^1(\cdot, t) & \text{for } p \in G/H, t \geq 0, \\ W^1(p, 0) = U^1(p) & \text{for } p \in G/H. \end{cases} \quad (25)$$

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

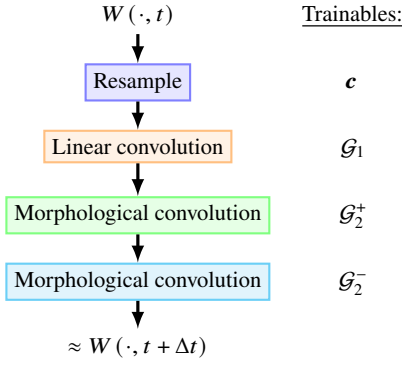


Fig. 8 Evolving the PDE through operator splitting, each operation corresponds to a term of (24).

Proposition 6 (Convection solution) *The solution of the convection PDE is found by the method of characteristics, and is given by*

$$\begin{aligned} W^1(p, t) &= \left(\mathcal{L}_{g_p^{-1}} U^1 \right) \left(\gamma_{\mathbf{c}}(t)^{-1} p_0 \right) \\ &= U^1 \left(g_p \gamma_{\mathbf{c}}(t)^{-1} p_0 \right), \end{aligned} \quad (26)$$

where $g_p \in p$ (i.e. $g_p p_0 = p$) and $\gamma_{\mathbf{c}} : \mathbb{R} \rightarrow G$ is the exponential curve that satisfies $\gamma_{\mathbf{c}}(0) = e$ and

$$\frac{\partial}{\partial t} (\gamma_{\mathbf{c}}(t)p) (t) = \mathbf{c} (\gamma_{\mathbf{c}}(t)p), \quad (27)$$

i.e. $\gamma_{\mathbf{c}}$ is the exponential curve in the group G that induces the integral curves of the G -invariant vector field \mathbf{c} on G/H when acting on elements of the homogeneous space.

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

Proof

$$\begin{aligned} \frac{\partial W^1}{\partial t}(p, t) &= \lim_{h \rightarrow 0} \frac{W^1(p, t+h) - W^1(p, t)}{h} \\ &= \lim_{h \rightarrow 0} \frac{U^1 \left(g_p \gamma_{\mathbf{c}}(t+h)^{-1} p_0 \right) - U^1 \left(g_p \gamma_{\mathbf{c}}(t)^{-1} p_0 \right)}{h} \\ &= \lim_{h \rightarrow 0} \frac{U^1 \left(g_p \gamma_{\mathbf{c}}(t)^{-1} \gamma_{\mathbf{c}}(h)^{-1} p_0 \right) - U^1 \left(g_p \gamma_{\mathbf{c}}(t)^{-1} p_0 \right)}{h}, \end{aligned}$$

now let $\bar{U} := \mathcal{L}_{\gamma_{\mathbf{c}}(t) g_p^{-1}} U^1$, then

$$\begin{aligned} &= \lim_{h \rightarrow 0} \frac{\bar{U} \left(\gamma_{\mathbf{c}}(h)^{-1} p_0 \right) - \bar{U} (p_0)}{h} \\ &= -\mathbf{c}(p_0) \bar{U} \\ &= - \left(\mathcal{L}_{g_p} \right)_* \mathbf{c}(p_0) \mathcal{L}_{g_p} \bar{U} \end{aligned}$$

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

$$\begin{aligned} &= -\mathbf{c}(p) \mathcal{L}_{g_p} \mathcal{L}_{\gamma_{\mathbf{c}}(t) g_p^{-1}} U^1 \\ &= -\mathbf{c}(p) \left[p \mapsto U^1 \left(g_p \gamma_{\mathbf{c}}(t)^{-1} g_p^{-1} p \right) \right] \\ &= -\mathbf{c}(p) \left[p \mapsto U^1 \left(g_p \gamma_{\mathbf{c}}(t)^{-1} p_0 \right) \right] \\ &= -\mathbf{c}(p) W^1(\cdot, t). \end{aligned}$$

□

Numerically equation (26) is implemented as a resampling operation with interpolation to account for the off-grid coordinates.

5.2 Fractional Diffusion

The (fractional) diffusion step solves the PDE

$$\begin{cases} \frac{\partial W^2}{\partial t} = -(-\Delta_{\mathcal{G}_1})^\alpha W^2(p, t) & \text{for } p \in G/H, t \geq 0, \\ W^2(p, 0) = U^2(p) & \text{for } p \in G/H. \end{cases} \quad (28)$$

As with (fractional) diffusion on \mathbb{R}^n , there exists a smooth function

$$K_t^\alpha : (0, \infty) \times (G/H) \rightarrow [0, \infty),$$

called the fundamental solution of the α -diffusion equation, such that for every initial condition U^2 , the solution to the PDE (28) is given by the convolution of the function U^2 with the fundamental solution K_t^α :

$$W^2(p, t) = \left(K_t^\alpha *_{G/H} U^2 \right) (p). \quad (29)$$

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

Definition 9 (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:

$$\forall h \in H, p \in G/H : k(hp) = k(p) \quad (\text{kernel compatibility})$$

then we define:

$$(k *_{G/H} f)(p) := \frac{1}{\mu_H(H)} \int_G k(g^{-1}p) f(gp_0) d\mu_G(g), \quad (30)$$

where μ_H and μ_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 to the left-invariant Haar measure on G as ‘the Haar measure on G ’ 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 righthand side of (30) does not converge. To this end we note that one can always decompose (by Weil’s integral formula [77, Lem.2.1]) the Haar measure μ_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 μ_G such that

$$\mu_G \left(\pi^{-1}(A) \right) = \mu_H(H) \mu_{\mathcal{G}}(A), \quad \forall A \subset G/H, \quad (31)$$

where $\mu_{\mathcal{G}}$ is the Riemannian measure induced by \mathcal{G} and μ_H is a choice of Haar measure on H . Thereby (31) boils down to Weil’s integration formula:

$$\mu_H(H) \int_{G/H} f(p) d\mu_{\mathcal{G}}(p) = \int_G f(gH) d\mu_G(g) \quad (32)$$

whenever f is measurable. Since H is compact we can indeed normalize the Haar measure μ_H so that $\mu_H(H) = 1$.

In general an exact analytic expression for the fundamental solution K_t^α requires complicated steerable filter operators [48, 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 exponential map in the Lie group. Locally, such inversion can always be done by the inverse function theorem. Globally, we need assumptions: we shall assume that the exponential map is surjective¹ and that the Lie group is connected with a surjective exponential map that is a diffeomorphism between some $V \subset T_e(G)$ and G . Then we define the logarithmic map $\log_G : G \rightarrow V$ by $\exp_G \circ \log_G = \text{id}_G$ and $\log_G \circ \exp_G|_V = \text{id}_V$. For the moment, for simplicity, we assume $V = T_e(G)$ in the general setting. We address the cases² where one has $V \subset T_e(G)$ with $V \neq T_e(G)$ later.

The idea is that instead of basing our kernels on the metric $d_{\mathcal{G}}$ (which is hard to calculate [78]) we approximate it using the seminorm from Def. 7 (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 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:

$$d_{\mathcal{G}}(p_0, p) \approx \inf_{g \in p} \|\log_G g\|_{\mathcal{G}}. \quad (33)$$

Henceforth, we write this estimate as $d_{\mathcal{G}}(p_0, p) \approx \rho_{\mathcal{G}}(p)$ relying on the following definition.

¹ For example the case if $G = \mathbb{R}^d$ and $G = SE(d)$ or if G is a connected compact Lie group such as $G = SO(d)$

² In our primary case of interest $G = SE(2)$ one has: we have $V = \{\sum_{k=1}^3 c^k A_k \mid c^3 \in [-\pi, \pi)\}$.

Definition 10 (Logarithmic metric estimate) Let \mathcal{G} be a G -invariant metric tensor field on the homogeneous space G/H then we define

$$\begin{aligned} \rho_{\mathcal{G}}(p) &:= \inf_{g \in p} \|\log_G g\|_{\mathcal{G}} \\ &:= \inf_{g \in p} \sqrt{\mathcal{G}(\pi_* \log_G g, \pi_* \log_G g)}, \end{aligned} \quad (34)$$

where π_* is the push-forward of the projection map π 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 Def. 7 and use that velocity as the distance estimate.

Summarizing, Def. 10 and Eq. (33), 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 2 (Bounding the logarithmic metric estimate)

There exists a $C_{\text{metr}} \geq 1$ so that

$$d_{\mathcal{G}}(p_0, p) \leq \rho_{\mathcal{G}}(p) \leq C_{\text{metr}} d_{\mathcal{G}}(p_0, p) \quad (35)$$

for all $p \in G/H$ in a compact neighborhood around p_0 away from the cut-locus.

The proof of this lemma can be found in Appendix A.

Remark 9 (Logarithmic metric estimate in principal homogeneous spaces) When we take a principal homogeneous space such as $\mathbb{M}_2 \equiv SE(2)$ with a left-invariant metric tensor field the metric estimate simplifies to

$$\rho_{\mathcal{G}}(g) = \|\log_G g\|_{\mathcal{G}|_e},$$

hence we see that this construction generalizes the logarithmic estimate, as used in [79, 80], to homogeneous spaces other than the principal.

Remark 10 (Logarithmic metric estimate for \mathbb{M}_2) Using the (x, y, θ) 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:

$$\begin{aligned} c^1(x, y, \theta) &:= \begin{cases} \frac{\theta}{2} (y + x \cot \frac{\theta}{2}) & \text{if } \theta \neq 0, \\ x & \text{if } \theta = 0, \end{cases} \\ c^2(x, y, \theta) &:= \begin{cases} \frac{\theta}{2} (-x + y \cot \frac{\theta}{2}) & \text{if } \theta \neq 0, \\ y & \text{if } \theta = 0, \end{cases} \\ c^3(x, y, \theta) &:= \theta. \end{aligned}$$

The logarithmic metric estimate for $SE(2)$ is then given by

$$\rho_{\mathcal{G}}(x, y, \theta) = \sqrt{w_M c^1(x, y, \theta)^2 + w_L c^2(x, y, \theta)^2 + w_A c^3(x, y, \theta)^2},$$

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




Fig. 9 Comparing the ‘exact’ Riemannian distance (left) obtained through numerically solving the Eikonal equation [28] versus the logarithmic metric estimate (right) on $SE(2)$ endowed with a left-invariant Riemannian metric tensor field (20) with $w_M = 1$, $w_L = 2$, $w_A = 1/\pi$.

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 Def. 9.

Proposition 7 (Kernel compatibility of $\rho_{\mathcal{G}}$) Let \mathcal{G} be a G -invariant metric tensor field on G/H then we have

$$\forall p \in G/H, \forall h \in H : \rho_{\mathcal{G}}(hp) = \rho_{\mathcal{G}}(p). \quad (36)$$

Proof We rewrite Def. 10 and find

$$\rho_{\mathcal{G}}(p) = \inf_{g \in \mathcal{P}} \left\| \frac{\partial}{\partial t} \exp_G(t \log_G g) p_0 \Big|_{t=0} \right\|_{\mathcal{G}|_{p_0}},$$

where we recognize that we are optimizing over left-invariant curves in G whose actions connect p_0 (at $t = 0$) to hp (at $t = 1$). Due to the G -invariance of \mathcal{G} and the fact that $hp_0 = p_0$ the following equality holds for all $h \in H$:

$$= \inf_{g \in \mathcal{P}} \left\| (L_h)_* \frac{\partial}{\partial t} \exp_G(t \log_G g) p_0 \Big|_{t=0} \right\|_{\mathcal{G}|_{p_0}}.$$

This can be rewritten as:

$$= \inf_{g \in \mathcal{P}} \left\| \frac{\partial}{\partial t} h \exp_G(t \log_G g) p_0 \Big|_{t=0} \right\|_{\mathcal{G}|_{p_0}}.$$

Now we see that we are optimizing over a set of left-invariant curves whose actions connect p_0 (at $t = 0$) to hp (at $t = 1$) i.e. we have:

$$\begin{aligned} &= \inf_{g \in hp} \left\| \frac{\partial}{\partial t} \exp_G(t \log_G g) p_0 \Big|_{t=0} \right\|_{\mathcal{G}|_{p_0}} \\ &= \rho_{\mathcal{G}}(hp). \end{aligned}$$

□

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 11 (Approximate $\alpha = 1$ kernel)

$$K_t^{1,\text{appr}}(p) := \eta_t \exp\left(-\frac{\rho_{\mathcal{G}}(p)^2}{4t}\right) \quad (37)$$

where η_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: $\eta_t = \mu_{\mathcal{G}}\left(B(p_0, \sqrt{t})\right)^{-1}$, 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.

Definition 12 (Polynomial growth) A Lie group G with left-invariant Haar measure μ_G is of polynomial growth when the volume of a sphere of radius r around $g \in G$:

$$B(g, r) = \{g' \in G \mid d_{\mathcal{G}}(g, g') < r\},$$

can be polynomially bounded, i.e. there exists constants $\delta > 0$ and $C_{\text{grow}} > 0$ so that

$$\frac{1}{C_{\text{grow}}} r^{\delta} \leq \mu_G(B(g, r)) \leq C_{\text{grow}} r^{\delta}, \quad r \geq 1.$$

Since μ_G is left-invariant the choice of g does not matter.

Lemma 3 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 exists constants $C \geq 1$, $D_1 \in (0, 1)$ and $D_2 > D_1$ so that for all $t > 0$ the following holds:

$$\frac{1}{C} K_{D_1 t}^1(p) \leq K_t^{1,\text{appr}}(p) \leq C K_{D_2 t}^1(p). \quad (38)$$

for all $p \in G/H$.

Proof On a group of polynomial growth we have $\eta_t = \mu_{\mathcal{G}}\left(B(p_0, \sqrt{t})\right)^{-1}$. If G is of polynomial growth we can apply [81, Thm. 2.12] to find that there exists constants $C_1, C_2 > 0$ and for all $\varepsilon > 0$ there exists a constant C_{ε} so that:

$$\begin{aligned} C_1 \eta_t \exp\left(-\frac{d_{\mathcal{G}}(p_0, p)^2}{4C_2 t}\right) &\leq K_t^1(p) \\ &\leq C_{\varepsilon} \eta_t \exp\left(-\frac{d_{\mathcal{G}}(p_0, p)^2}{4(1+\varepsilon)t}\right). \end{aligned}$$

Note that Maheux [81] uses right cosets while we use left cosets. We can translate the results easily by inversion in view of $(gH)^{-1} = H^{-1}g^{-1} = Hg^{-1}$. 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:

$$\begin{aligned} d_{G/H}^{\text{maheux}}(p_1, p_2) &= \inf_{g_1 \in p_1} \inf_{g_2 \in p_2} d_G^{\text{maheux}}(g_1, g_2) \\ &= \inf_{g_2 \in p_2} d_G^{\text{maheux}}(q_1, g_2), \end{aligned} \quad (39)$$

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 2 we obtain:

$$\begin{aligned} C_1 \eta_t \exp\left(-\frac{\rho_{\mathcal{G}}(p)^2}{4C_2 t}\right) &\leq K_t^1(p) \\ &\leq C_\varepsilon \eta_t \exp\left(-\frac{\rho_{\mathcal{G}}(p)^2}{4C_{\text{metr}}^2(1+\varepsilon)t}\right), \end{aligned}$$

which leads to:

$$\begin{aligned} C_1 \frac{\eta_t}{\eta_{C_2 t}} K_{C_2 t}^{1,\text{appr}}(p) &\leq K_t^1(p) \\ &\leq C_\varepsilon \frac{\eta_t}{\eta_{C_{\text{metr}}^2(1+\varepsilon)t}} K_{C_{\text{metr}}^2(1+\varepsilon)t}^{1,\text{appr}}(p). \end{aligned}$$

The group G being of polynomial growth also implies G/H is a doubling space [81, Thm. 2.17]. Using the volume doubling and reverse volume doubling property of doubling spaces [82, Prop. 3.2 and 3.3] we find that there exist constants $C_3, C_4, \beta, \beta' > 0$ so that:

$$\begin{aligned} \frac{\eta_t}{\eta_{C_2 t}} &\geq C_3 \left(\frac{\sqrt{t}}{\sqrt{C_2 t}}\right)^\beta = C_3 C_2^{-\beta/2}, \\ \frac{\eta_t}{\eta_{C_{\text{metr}}^2(1+\varepsilon)t}} &\leq C_4 \left(\frac{\sqrt{t}}{\sqrt{C_{\text{metr}}^2(1+\varepsilon)t}}\right)^{\beta'} \\ &= C_4 \left(C_{\text{metr}}^2(1+\varepsilon)\right)^{-\beta'/2}. \end{aligned}$$

Applying these inequalities we get:

$$C'_1 := C_1 C_3 C_2^{-\beta/2}$$

and

$$C'_\varepsilon := C_\varepsilon C_4 \left(C_{\text{metr}}^2(1+\varepsilon)\right)^{-\beta'/2}$$

we obtain:

$$C'_1 K_{C_2 t}^{1,\text{appr}}(p) \leq K_t^1(p) \leq C'_\varepsilon K_{C_{\text{metr}}^2(1+\varepsilon)t}^{1,\text{appr}}(p).$$

Reparametrising t in both inequalities gives:

$$\frac{1}{C'_\varepsilon} K_{t/(C_{\text{metr}}^2(1+\varepsilon))}^1(p) \leq K_t^{1,\text{appr}}(p) \leq \frac{1}{C'_1} K_{C_2 t}^1(p).$$

Finally we fix $\varepsilon > 0$ and relabel constants:

$$C := \max\{C_1'^{-1}, C'_\varepsilon, 1\},$$

$$D_1 := \frac{1}{C_{\text{metr}}^2(1+\varepsilon)},$$

$$D_2 := \frac{1}{C_2},$$

observe that since $\varepsilon > 0$ and $C_{\text{metr}} \geq 1$ we have $0 < D_1 < 1$. \square

Depending on the actually achievable constants, Lem. 3 provides a very strong or very weak bound on how much our approximation deviates from the fundamental solution. Fortunately in the $SE(2)$ case our approximation is very close to the exact kernel in the vicinity of the origin, as can be seen 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.




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 $G/H = SE(2)$. 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.

Now let us develop an approximation for values of α other than 1. From semi-group theory [83] it follows that semi-groups generated by taking fractional powers of the generator (in our case $\Delta_{\mathcal{G}} \rightarrow -(-\Delta_{\mathcal{G}})^\alpha$) amounts to the following key relation between the α -kernel and the diffusion kernel:

$$K_t^\alpha(p) := \int_0^\infty q_{t,\alpha}(\tau) K_\tau^1(p) d\tau, \quad (40)$$

for $\alpha \in (0, 1)$ and $t > 0$ where $q_{t,\alpha}$ is defined as follows.

Definition 13 Let \mathcal{L}^{-1} be the inverse Laplace transform then

$$q_{t,\alpha}(\tau) := \mathcal{L}^{-1}\left(r \mapsto e^{-tr^\alpha}\right)(\tau) \quad \text{for } \tau \geq 0.$$

For explicit formulas of this kernel see [83, Ch. IX:11 eq. 17]. Since e^{-tr^α} 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 α -kernel.

Definition 14 (Approximate $\alpha \in (0, 1)$ kernel) Akin to (40) we set $\alpha \in (0, 1)$, $t > 0$ and define:

$$K_t^{\alpha, \text{appr}}(p) := \int_0^\infty q_{t, \alpha}(\tau) K_\tau^{1, \text{appr}}(p) d\tau \geq 0, \quad (41)$$

for $p \in G/H$.

The bounding of K_t^1 we obtained in Lem. 3 transfers directly to our approximation for other α .

Theorem 1 Let G be of polynomial growth and let K_t^α 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:

$$\frac{1}{C} K_{D_1 t}^\alpha(p) \leq K_t^{\alpha, \text{appr}}(p) \leq C K_{D_2 t}^\alpha(p). \quad (42)$$

Proof This is an consequence of Lem. 3 and the fact that $q_{t, \alpha}$ is positive, applying the integral from (40) yields:

$$\begin{aligned} K_t^{\alpha, \text{appr}}(p) &= \int_0^\infty q_{t, \alpha}(\tau) K_\tau^{1, \text{appr}}(p) d\tau \\ &\stackrel{\text{(Lem. 3)}}{\leq} C \int_0^\infty q_{t, \alpha}(\tau) K_{D_2 \tau}^1(p) d\tau \\ &\stackrel{(\tau' = D_2 \tau)}{=} C \int_0^\infty \frac{1}{D_2} q_{t, \alpha}\left(\frac{\tau'}{D_2}\right) K_{\tau'}^1(p) d\tau' \\ &\stackrel{\text{(Bromwich integral)}}{=} C \int_0^\infty q_{D_2 t, \alpha}(\tau') K_{\tau'}^1(p) d\tau' \\ &= C K_{D_2 t}^\alpha(p). \end{aligned}$$

The other inequality works the same way. \square

Although the approximation (41) is helpful in the proof above it contains some integration and is not an explicit expression. We leave developing a more explicit and computable approximation for $0 < \alpha < 1$ for future work.

5.3 Dilation and Erosion

The dilation/erosion step solves the PDE

$$\begin{cases} \frac{\partial W^3}{\partial t}(p, t) = \pm \left\| \nabla_{\mathcal{G}_2^\pm} W^3(p, t) \right\|_{\mathcal{G}_2^\pm}^{2\alpha} & \text{for } p \in G/H, \\ t \geq 0, & (43) \\ W^3(p, 0) = U^3(p) & \text{for } p \in G/H. \end{cases}$$

By a generalization of the Hopf-Lax formula [84, Ch.10.3], the solution is given by morphological convolution

$$W^3(p, t) = - \left(k_t^\alpha \square_G - U^3 \right) (p) \quad (44)$$

for the (+) (dilation) variant and

$$W^3(p, t) = \left(k_t^\alpha \square_G U^3 \right) (p) \quad (45)$$




Fig. 11 Shapes of the level sets of the kernels on \mathbb{M}_2 for solving fractional diffusion (K_t^α) and dilation/erosion (k_t^α) 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 .

for the (-) (erosion) variant, where the kernel k_t^α (also called the structuring element in the context of morphology) is given as follows.

Definition 15 (Dilation/erosion kernels) The morphological convolution kernel k_t^α for small times t and $\alpha \in (1/2, 1]$ is given by

$$k_t^\alpha(p) := \nu_\alpha t^{-\frac{1}{2\alpha-1}} d_{\mathcal{G}_2}(p_0, p)^{\frac{2\alpha}{2\alpha-1}}, \quad (46)$$

with $\nu_\alpha := \left(\frac{2\alpha-1}{(2\alpha)^{2\alpha/(2\alpha-1)}} \right)$ and for $\alpha = 1/2$ by

$$k_t^{1/2}(p) = \begin{cases} 0 & \text{if } d_{\mathcal{G}_2}(p_0, p)^{2\alpha} \leq t, \\ \infty & \text{if } d_{\mathcal{G}_2}(p_0, p)^{2\alpha} > t. \end{cases} \quad (47)$$

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

$$\begin{aligned} (k \square_G f)(p) &:= \inf_{g \in G} \{k(g^{-1}p) + f(gp_0)\} \\ &:= \inf_{g \in G} \{k(g^{-1}p) + f(gH)\}. \end{aligned}$$

Remark 11 (Grayscale morphology) Morphological convolution is related to the grayscale morphology operations \oplus (dilation) and \ominus (erosion) on \mathbb{R}^d as follows:

$$\begin{aligned} f_1 \oplus f_2 &= -(-f_1 \square_{\mathbb{R}^d} - f_2), \\ f_1 \ominus f_2 &= f_1 \square_{\mathbb{R}^d} [x \mapsto -f_2(-x)], \end{aligned}$$

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^α allows us to solve (43).

Theorem 2 *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*

$$W^3(p, t) := (k_t^\alpha \square_G U^3)(p)$$

is Lipschitz and solves the $(-)$ -variant, the erosion variant, of the system (43) in the sense of Theorem 2.1 in [85], while

$$W^3(p, t) := -(k_t^\alpha \square_G -U^3)(p)$$

is Lipschitz and solves the $(+)$ -variant, the dilation variant, of system (43) in the sense of Theorem 2.1 in [85]. 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 $(G/H, \mathcal{G}_2)$ is a proper length space, and therefore the theory of [85] applies. Moreover since G is of polynomial growth we have that G/H is a doubling space [81, Thm. 2.17] and also admits a Poincaré constant [81, Thm. 2.18]. So we meet the additional requirements of [85, Thm. 2.3 (vii) and (viii)].

The Hamiltonian $\mathcal{H} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ in [85] 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 [85] the solution (in the sense of their Theorem 2.1) to the $(-)$ -variant of system (43) is given by

$$\begin{aligned} W^3(p, t) &= \inf_{x \in G/H} \left\{ t \mathcal{L} \left(\frac{d_{\mathcal{G}_2}(p, x)}{t} \right) + U^3(x) \right\} \\ &= \inf_{g \in G} \left\{ t \mathcal{L} \left(\frac{d_{\mathcal{G}_2}(p, gp_0)}{t} \right) + U^3(gp_0) \right\} \\ &= \inf_{g \in G} \left\{ t \mathcal{L} \left(\frac{d_{\mathcal{G}_2}(g^{-1}p, p_0)}{t} \right) + U^3(gp_0) \right\} \\ &= \inf_{g \in G} \left\{ v_\alpha \frac{d_{\mathcal{G}_2}(g^{-1}p, p_0)^{\frac{2\alpha}{2\alpha-1}}}{t^{\frac{2\alpha}{2\alpha-1}-1}} + U^3(gp_0) \right\} \\ &= \inf_{g \in G} \left\{ v_\alpha t^{1-\frac{2\alpha}{2\alpha-1}} d_{\mathcal{G}_2}(g^{-1}p, p_0)^{\frac{2\alpha}{2\alpha-1}} + U^3(gp_0) \right\} \\ &= \inf_{g \in G} \left\{ v_\alpha t^{\frac{-1}{2\alpha-1}} d_{\mathcal{G}_2}(g^{-1}p, p_0)^{\frac{2\alpha}{2\alpha-1}} + U^3(gp_0) \right\} \\ &= (k_t^\alpha \square_G U^3)(p). \end{aligned}$$

The $(+)$ -variant is proven analogously.

The semigroup property follows directly from [85, Thm 2.1(ii)]. \square

Remark 12 (Solution according to Balogh et al.)

This theorem builds on the work by Balogh et al. [85] 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 $\|\nabla_{\mathcal{G}_2} W(p, t)\|_{\mathcal{G}_2}$ by:

$$\limsup_{p' \rightarrow p} \frac{\max(W(p, t) - W(p', t), 0)}{d_{\mathcal{G}_2}(p, p')},$$

and we get a solution concept in terms of this slightly different notion of a gradient.

Remark 13 (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 [85]. 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 [86, 87]. In the general Riemannian homogeneous space setting the result by Azagra [87, 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 ([88, ch. 3], [28, ch. 4, app. E], [84]) 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 $H(2d+1)$ studied in [89], 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 towards HJB-systems on \mathbb{R}^n or $\mathbb{R}^d \times S^q$, with $n = d + q = \dim(G/H)$. Then one can directly apply results by Dragoni [86, 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 α close to $1/2$ (kernel and result in red) that we obtain what amounts to an equivariant version of max/min pooling.

The level sets of the kernels k_t^α 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 min/max pooling.




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 kernels, 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.

Remark 14 The level sets in Fig. 11 are balls in $G/H = \mathbb{M}_2$ that do not depend on α . It is only the growth of the kernel values when passing through these level sets that depends on α . 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 17 (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

$$k_t^{\alpha, \text{appr}}(p) := \nu_\alpha t^{-\frac{1}{2\alpha-1}} \rho_{\mathcal{G}_2}(p)^{\frac{2\alpha}{2\alpha-1}}, \quad (48)$$

with $\nu_\alpha := \left(\frac{2\alpha-1}{(2\alpha)^{2\alpha/(2\alpha-1)}} \right)$ and for $\alpha = 1/2$ by

$$k_t^{1/2, \text{appr}}(p) = \begin{cases} 0 & \text{if } \rho_{\mathcal{G}_2}(p) \leq t, \\ \infty & \text{if } \rho_{\mathcal{G}_2}(p) > t \end{cases}. \quad (49)$$

We used this approximation in our parallel GPU-algorithms (for our PDE-G-CNNs experiments in Section 7). It is highly preferable over the ‘exact’ solution based on the true distance as this would require Eikonal PDE solvers ([28, 90] which would not be practical for parallel GPU implementations of PDE-G-CNNs. Again the approximations are reasonable, likewise Fig. 10, even on coarse grids in practice [75], and again the approximations rely on the core idea of Lemma 2.

Next we formalize the theoretical underpinning of the approximations in the upcoming corollary.

An immediate consequence of Def. 17 and Lem. 2 (keeping in mind that the kernel expressions in Def. 17 are monotonic w.r.t. $\rho := \rho_{\mathcal{G}_2}(p)$) is that we can enclose our approximate morphological kernel with the exact morphological

kernels in the same way as we did for the (fractional) diffusion kernel in Theorem 1. This proves the following Corollaries:

Corollary 3 Let $\alpha \in (1/2, 1]$ then for all $t > 0$

$$k_t^\alpha(p) \leq k_t^{\alpha, \text{appr}}(p) \leq C_{\text{metr}}^{\frac{2\alpha}{2\alpha-1}} k_t^\alpha(p) \quad \text{for } p \in G/H.$$

For the case $\alpha = 1/2$ the approximation is exact in an inner and outer region:

$$\begin{aligned} k_t^{1/2, \text{appr}}(p) &= k_t^{1/2}(p) = 0 && \text{if } \rho_{\mathcal{G}_2}(p)^{2\alpha} \leq t, \\ k_t^{1/2, \text{appr}}(p) &= k_t^{1/2}(p) = \infty && \text{if } d_{\mathcal{G}_2}(p_0, p)^{2\alpha} > t, \end{aligned}$$

but in the intermediate region where $\rho_{\mathcal{G}_2}(p)^{2\alpha} > t$ and $d_{\mathcal{G}_2}(p_0, p)^{2\alpha} \leq t$ we have $k_t^{1/2, \text{appr}} = \infty$ while $k_t^{1/2} = 0$.

Alternatively, instead of bounding by value we can bound in time, in which case we do not need to distinguish different cases of α .

Corollary 4 Let $\alpha \in [1/2, 1], t > 0$ then for all $p \in G/H$ one has

$$k_t^\alpha(p) \leq k_t^{\alpha, \text{appr}}(p) \leq k_{C_{\text{metr}}^{-2\alpha} t}^\alpha(p)$$

With these two bounds on our approximate morphological kernels we end our theoretical results.

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 & 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 (26) of convection PDEs:

$$\begin{aligned} (k *_{G/H} U)(p) &= \frac{1}{\mu_G(H)} \int_G k(g^{-1}p) U(gp_0) d\mu_G(g) \\ &= \frac{1}{\mu_G(H)} \int_G k(g^{-1}g_p p_0) U(gp_0) d\mu_G(g), \end{aligned}$$

for any $g_p \in p$, now change variables to $q = g_p^{-1}g$ and recall that μ_G is left invariant:

$$= \frac{1}{\mu_G(H)} \int_G k(q^{-1}p_0) U(g_p q p_0) d\mu_G(q).$$

In this last expression we recognize (26) and see that we can interpret $p \mapsto U(g_p q p_0)$ as the solution of the convection PDE (25) at time $t = 1$ for a convection vector field \mathbf{c} that has flow lines given by $\gamma_{\mathbf{c}}(t) = \exp_G(-t \log_G q) p_0$ so that $(\gamma_{\mathbf{c}}(1))^{-1} p_0 = q p_0$. 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^α (for some choice of α) as follows:

$$k(p) = \sum_{i=1}^n k_i K_{t_i}^\alpha(g_i^{-1}p), \quad (50)$$

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:

$$\begin{aligned} (k *_{G/H} U)(p) &= \int_G \sum_{i=1}^n k_i K_{t_i}^\alpha(g_i^{-1}g^{-1}p) U(gp_0) d\mu_G(g) \\ &= \sum_{i=1}^n k_i \int_G K_{t_i}^\alpha(g_i^{-1}g^{-1}p) U(gp_0) d\mu_G(g), \end{aligned}$$

we change variables to $h = g g_i$:

$$\begin{aligned} &= \sum_{i=1}^n k_i \int_G K_{t_i}^\alpha(h^{-1}p) U(h g_i^{-1} p_0) d\mu_G(h) \\ &= \sum_{i=1}^n k_i \left(K_{t_i}^\alpha *_{G/H} \left[q \mapsto U(g_q g_i^{-1} p_0) \right] \right)(p). \end{aligned}$$

Here again we recognize $q \mapsto U(g_q g_i^{-1} p_0)$ as the solution (26) of the convection PDE at $t = 1$ with flow lines induced

by $\gamma_{\mathbf{c}}(t) = \exp_G(t \log_G g_i)$. 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, sans 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 8 (Max pooling) *Let $f \in L^\infty(G/H)$, let $S \subset G/H$ be non empty and define $k_S : G/H \rightarrow \mathbb{R} \cup \{\infty\}$ as:*

$$k_S(p) := \begin{cases} 0 & \text{if } p \in S, \\ \infty & \text{else.} \end{cases} \quad (51)$$

Then:

$$-(k_S \square -f)(p) = \sup_{g \in G: g^{-1}p \in S} f(gp_0). \quad (52)$$

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 Def. 16 yields:

$$\begin{aligned} &-(k_S \square -f)(p) \\ &= -\inf \left\{ \inf_{g \in G: g^{-1}p \in S} -f(gp_0), \inf_{g \in G: g^{-1}p \notin S} -f(gp_0) + \infty \right\} \\ &= -\inf_{g \in G: g^{-1}p \in S} -f(gp_0) \\ &= \sup_{g \in G: g^{-1}p \in S} f(gp_0). \\ &\square \end{aligned}$$

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

Corollary 5 (Euclidean Max Pooling) *Let $G = G/H = \mathbb{R}^n$ and let $f \in C^0(\mathbb{R}^n)$ with $S \subset \mathbb{R}^n$ compact then:*

$$-(k_S \square_{\mathbb{R}^n} -f)(x) = \max_{y \in S} f(x - y),$$

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 9 *Let f be a compactly supported continuous function on G/H . Then dilation with the kernel*

$$k_{\text{ReLU},f}(p) := \begin{cases} 0 & \text{if } p = p_0, \\ \sup_{p \in G/H} f(p) & \text{else,} \end{cases}$$

equates to applying a Rectified Linear Unit to the function f :

$$-(k_{\text{ReLU},f} \square -f)(p) = \max\{0, f(p)\}.$$

Proof Filling in k into the definition of morphological convolution:

$$\begin{aligned} & -(k_{\text{ReLU},f} \square -f)(p) \\ &= -\inf_{g \in G} k_{\text{ReLU}}(g^{-1}p) - f(g \cdot p_0) \\ &= -\inf_{g \in G} \left\{ \inf_{g^{-1}p=p_0} -f(gp_0), \inf_{g^{-1}p \neq p_0} -f(gp_0) + \sup_{y \in G/H} f(y) \right\} \\ &= \sup \left\{ f(p), \sup_{z \in G/H: z \neq p} f(z) - \sup_{y \in G/H} f(y) \right\}, \end{aligned}$$

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 than 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* [65] 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




Fig. 13 A residual block, like in [65], note the resemblance to a forward Euler discretization scheme.

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

$$\begin{cases} \frac{\partial U}{\partial t}(p, t) = \mathcal{F}(U(\cdot, t), p) & \text{for } p \in M, t \geq 0, \\ U(p, 0) = U_0(p) & \text{for } p \in M, \end{cases}$$

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 timestep $\Delta t > 0$. 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 *transformer networks* towards what could be described as *evolution networks*.

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 [72].

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 the different types of networks in clearly defined feed-forward network architectures. So we test networks of modest size only and do not just aim for the performance that would be possible with large-scale networks.



(a) Example from the DRIVE [91] dataset, showing a retinal image and its vessel segmentation.



(b) Examples from the RotNIST [92] dataset.

Fig. 14 We perform a segmentation experiment on retinal vessel images and a classification experiment on rotation augmented digits.

7.1 Implementation

We implemented our PDE-based operators in an extension to the PyTorch deep learning framework [93]. 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 [94] 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.

Third, we fix $\alpha = 0.65$. We came to this value empirically; the networks performed best with α -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 (α close to $1/2$) minor perturbations in the input can have large effects on the output, when the kernel is too smooth (α 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 [94] 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 [91]. 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. 14(a).

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 with values in $\{0, 1\}$. We use the continuous DICE coefficient as the loss function:

$$\text{loss}(a, b) = 1 - \frac{2 \sum ab + \epsilon}{\sum a + \sum b + \epsilon},$$

where the sum \sum is over all the values in the 2D map. A small $\varepsilon > 0$ 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 Tbl. 1 and Fig. 16(a).

We achieve similar or better performance than CNNs or G-CNNs 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.

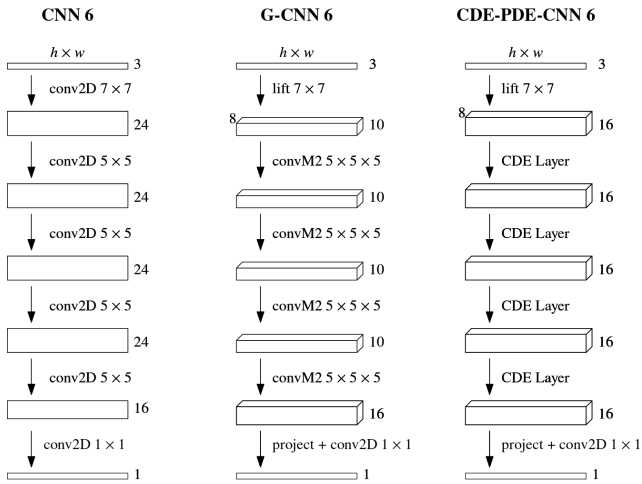


Fig. 15 Schematic of the 6-layer models used on our segmentation experiments. Kernel sizes and number of feature channels in each layer are indicated, depth indicates that the data lives on \mathbb{M}_2 . Omitted are activation functions, batch normalization, padding and dropout modules. The 12-layer models are essentially the same but with double the number of layers but with reduced width for the CDE-PDE-CNN (hence the reduction in parameters going from 6 to 12 layers).

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 [92]. RotNIST contains the same images as MNIST but rotated to various degrees. We made this choice to accentuate the roto-translational equivariance properties of the G-CNN and PDE-G-CNNs.

We tested three networks: the classic LeNet5 CNN [95] as a baseline, a 4-layer G-CNN and a 4-layer CDE-PDE-CNN.

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	0.8220 \pm 0.0007

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.

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 Tbl. 2 and Fig. 16(b).

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

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 2 Accuracy of the digit classification models on the testing dataset and number of parameters for each model.

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 Tbl. 3 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



(a) Performance on retinal vessel segmentation. We test 6- and 12-layer variants of conventional CNNs, G-CNNs and our PDE-CNNs, 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.

is however less optimized compared to our 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.

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

Table 3 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.

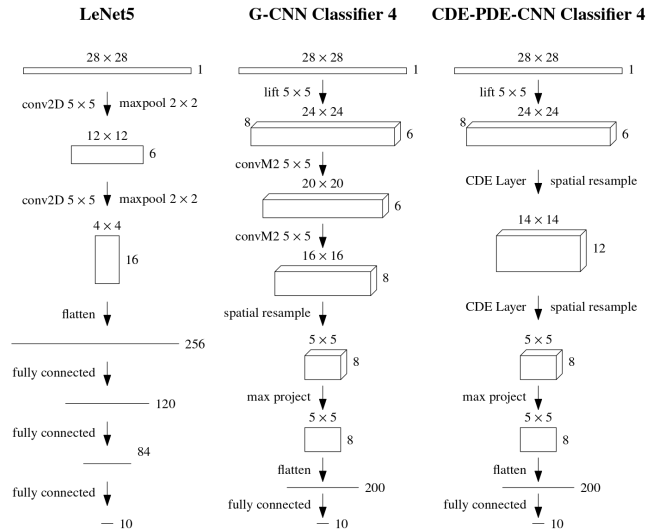


Fig. 17 Schematic of the three models tested with the RotNIST data. Kernel sizes and number of feature channels in each layer are indicated. Omitted are activation functions, batch normalization and dropout modules.

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 of parameters.

PDE-G-CNNs ensure equivariance by design. The trainable parameters are geometrically relevant: they are left-invariant 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 Theorem 1 and Theorem 2.

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.

Appendix A Proof of Lemma 2

The left inequality follows directly from the observation that $\rho_{\mathcal{G}}(p)$ is exactly the Riemannian length of the curve

$$t \mapsto \exp_G(t \log_G(g_p))p_0$$

for $t \in [0, 1]$ and $g_p = \arg \min_{g \in p} \|\log_G g\|_{\tilde{\mathcal{G}}}$. This continuous curve connects p_0 with p and as such has a greater length than the minimal-length curve between those two points.

Proving the right inequality requires some preliminaries. Let $G_0 \subset G$ be a compact neighborhood of e so that

- the map π restricted to G_0 is a submersion,
- it is away from the cut-locus,
- the group logarithm is a diffeomorphism from G_0 to $\log_G(G_0) \subset T_e G$,

For each of these three conditions there exists a compact neighborhood of e that satisfied it so we can always obtain G_0 by taking their intersection.

Now we construct three coordinate charts:

- $x : G_0 \rightarrow X = \mathbb{R}^{\dim(G/H)} \oplus \mathbb{R}^{\dim(H)}$,
- $y : \pi(G_0) \rightarrow Y = \mathbb{R}^{\dim(G/H)}$,
- $z : \log_G(G_0) \rightarrow Z = \mathbb{R}^{\dim(G/H)}$.

First, since $\pi : G_0 \rightarrow G/H$ is a submersion it follows from the *Submersion Theorem* that there exists a smooth chart x from G_0 to $X = \mathbb{R}^{\dim(G/H)} \oplus \mathbb{R}^{\dim(H)}$ and a smooth chart x' from $\pi(G_0)$ to $\mathbb{R}^{\dim(G/H)}$ so that

$$\begin{aligned} x' \circ \pi \circ x^{-1} \left(x^1, \dots, x^{\dim(G/H)}, x^{\dim(G/H)+1}, \dots, x^{\dim(G)} \right) \\ = \left(x^1, \dots, x^{\dim(G/H)} \right), \end{aligned}$$

i.e. in X the coordinates of the equivalence classes of G/H are orthogonal to the first $\dim(G/H)$ coordinates.

Second, since $\pi(G_0)$ is away from the cut-locus it admits a normal coordinates chart $y : \pi(G_0) \rightarrow Y = \mathbb{R}^{\dim(G/H)}$ with respect to \mathcal{G} . Since the coordinates given by y are normal we have that:

$$\forall p \in \pi(G_0) : |y(p)| = d_{\mathcal{G}}(p_0, p),$$

i.e. the Euclidean norm on Y agrees with the metric on G/H .

The third mapping $z : T_e G \rightarrow Z = \mathbb{R}^{\dim(G/H)}$ is a concatenation of chart and projection, we are only interested in the coordinates of the vector component that is orthogonal to $T_e H$, i.e. we disregard the degenerate part of $\tilde{\mathcal{G}}|_e$. In the non-degenerate subspace we want the coordinates to be normal with respect to $\tilde{\mathcal{G}}|_e$, or:

$$\forall g \in G_0 : |z(\log_G(g))| = \|\log_G(g)\|_{\tilde{\mathcal{G}}},$$

i.e. the Euclidean norm on Z agrees with the seminorm on $T_e G$.

The relation between all these mappings and spaces is illustrated in Fig. 18.

Let us now define the composite maps $\Omega := y \circ \pi \circ x^{-1}$ and $\Gamma := z \circ \log_G \circ x^{-1}$. By construction we have the following properties:

$$\begin{aligned} |\Omega(\mathbf{x})|^2 &= d_{\mathcal{G}} \left(p_0, \pi \left(x^{-1}(\mathbf{x}) \right) \right)^2 \quad \text{and} \\ |\Gamma(\mathbf{x})|^2 &= \left\| \log_G \left(x^{-1}(\mathbf{x}) \right) \right\|_{\tilde{\mathcal{G}}}^2, \end{aligned} \tag{53}$$

for all $\mathbf{x} \in x(G_0)$. Our inclusion of the squared Euclidean norm here is for a reason. Suppose you have a smooth function $F_1 : \mathbb{R}^m \rightarrow \mathbb{R}^n$ with $F_1(0) = 0$ and a smooth function $F_2 : \mathbb{R}^n \rightarrow \mathbb{R}$ also with $F_2(0) = 0$ and $D|_0 F_2 = 0$. Then $D|_0 (F_2 \circ F_1) = 0$ and also

$$\begin{aligned} D^2|_0 (F_2 \circ F_1) (\mathbf{v}, \mathbf{w}) \\ = D^2|_0 F_2 (D|_0 F_1(\mathbf{v}), D|_0 F_1(\mathbf{w})) + D|_0 F_2 \left(D^2|_0 F_2(\mathbf{v}, \mathbf{w}) \right) \\ = D^2|_0 F_2 (D|_0 F_1(\mathbf{v}), D|_0 F_1(\mathbf{w})). \end{aligned}$$

Now we substitute $F_2 = \Omega$ or $F_2 = \Gamma$ and $F_1 = \|\cdot\|^2$. As a result we have:

$$|\Omega(0)|^2 = |\Gamma(0)|^2 = 0, \quad D|_0 \left(|\Omega|^2 \right) = D|_0 \left(|\Gamma|^2 \right) = 0,$$

and

$$D^2|_0 \left(|\Omega|^2 \right) = 2 |D|_0 \Omega|^2, \quad D^2|_0 \left(|\Gamma|^2 \right) = 2 |D|_0 \Gamma|^2.$$

Taking these properties we can write the Taylor series of $|\Omega|^2$ and $|\Gamma|^2$ around 0 as follows:

$$\begin{aligned} |\Omega(\mathbf{x})|^2 &= |\Omega(0)|^2 + \left[D|_0 \left(|\Omega|^2 \right) \right] (\mathbf{x}) \\ &\quad + \frac{1}{2!} \left[D^2|_0 \left(|\Omega|^2 \right) \right] (\mathbf{x}, \mathbf{x}) + O(|\mathbf{x}|^3) \\ &= |D|_0 \Omega(\mathbf{x})|^2 + O(|\mathbf{x}|^3), \end{aligned}$$

$$\begin{aligned} |\Gamma(\mathbf{x})|^2 &= |\Gamma(0)|^2 + \left[D|_0 \left(|\Gamma|^2 \right) \right] (\mathbf{x}) \\ &\quad + \frac{1}{2!} \left[D^2|_0 \left(|\Gamma|^2 \right) \right] (\mathbf{x}, \mathbf{x}) + O(|\mathbf{x}|^3) \\ &= |D|_0 \Gamma(\mathbf{x})|^2 + O(|\mathbf{x}|^3). \end{aligned}$$

Now the linear maps $D|_0 \Omega$ and $D|_0 \Gamma$ on the space $X = \mathbb{R}^{\tilde{D}} \oplus \mathbb{R}^k$ both vanish in the \mathbb{R}^k subspace and are non-singular in the $\mathbb{R}^{\tilde{D}}$ subspace (i.e. invertible), hence we can conclude that:

$$\exists C > 0, \forall \mathbf{x} \in x(G_0) : |\Gamma(\mathbf{x})| \leq C |\Omega(\mathbf{x})|.$$

Or stated differently using (53):

$$\exists C > 0, \forall g \in G_0 : \|\log_G(g)\|_{\tilde{\mathcal{G}}} \leq C d_{\mathcal{G}}(p_0, gp_0).$$

Now we are almost there. Given any $p \in G/H$ so that $p \cap G_0 \neq \emptyset$ we know that for all $g \in p \cap G_0$ the above inequality holds. Since taking the infimum of the seminorm




Fig. 18 Relation of mappings used for the proof of Lemma 2. Through the local submersion theorem we obtain a chart x from a compact neighborhood G_0 of e to $X = \mathbb{R}^{\dim(G)}$. In X moving inside the equivalence classes is orthogonal to moving between equivalence classes. From the coordinate space X we can construct smooth maps to coordinate spaces Y and Z of G/H respectively $T_e G$. The spaces Y and Z are constructed so that in them the Euclidean norm matches the metric d_G respectively the seminorm $\|\cdot\|_{\tilde{\mathcal{G}}}$ of the corresponding elements. Mapping the metric and seminorm to the shared linear space X allows their comparison.

over the whole equivalence class p will yield a value lower or equal than that of the seminorms in $p \cap G_0$ we have:

$$\forall p \in G_0/H : \inf_{g \in p} \|\log_G(g)\|_{\tilde{\mathcal{G}}} \leq C d_G(p_0, p).$$

□

References

1. Welk, M., Weickert, J.: PDE evolutions for M-smoothers: From common myths to robust numerics. In: J. Lellmann, M. Burger, J. Modersitzki (eds.) *Scale Space and Variational Methods in Computer Vision*, pp. 236–248. Springer International Publishing, Cham (2019)
2. Fadili, J., Kutyniok, G., Peyré, G., Plonka-Hoch, G., Steidl, G.: Guest editorial: Mathematics and image analysis. *Journal of Mathematical Imaging and Vision* **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 International Publishing, Cham (2016)
6. J. Duran M. Moeller, C.S., 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. *International Journal of Computer Vision* **118**(3), 275–299 (2016)
8. Sapiro, G.: *Geometric Partial Differential Equations and Image Analysis*. Cambridge University Press (2001). DOI 10.1017/CBO9780511626319
9. Sethian, J.: *Level Set Methods and Fast Marching Methods*. Cambridge University Press (1999)
10. Weickert, J.: Theoretical foundations of anisotropic diffusion in image processing. *Computing, 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 (1995)
12. Cohen, T.S., Welling, M.: Group equivariant convolutional networks. In: *Int. Conf. on Machine Learning*, pp. 2990–2999 (2016)
13. Dieleman, S., De Fauw, J., Kavukcuoglu, K.: Exploiting cyclic symmetry in convolutional neural networks. *arXiv preprint arXiv:1602.02660* (2016)
14. Dieleman, S., Willett, K.W., Dambre, J.: Rotation-invariant convolutional neural networks for galaxy morphology prediction. *Monthly notices of the royal astronomical society* **450**(2), 1441–1459 (2015)
15. Winkels, M., Cohen, T.S.: 3d g-cnns for pulmonary nodule detection. *arXiv preprint arXiv:1804.04656* (2018)
16. 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)
17. Bekkers, E.J., Loog, M., ter Haar Romeny, B.M., Duits, R.: Template matching via densities on the roto-translation group. *IEEE transactions on pattern analysis and machine intelligence* **40**(2), 452–466 (2017)

18. 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)
19. 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). URL <https://arxiv.org/abs/1804.03393>
20. 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)
21. Cohen, T., Geiger, M., Weiler, M.: A General Theory of Equivariant CNNs on Homogeneous Spaces. arXiv preprint arXiv:1811.02017 (2018)
22. 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)
23. Kondor, R., Trivedi, S.: On the Generalization of Equivariance and Convolution in Neural Networks to the Action of Compact Groups. In: J. Dy, A. Krause (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). URL <http://proceedings.mlr.press/v80/kondor18a.html>
24. 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)
25. 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)
26. 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)
27. Duits, R., Haije, T.D., Creusen, E., Ghosh, A.: Morphological and linear scale spaces for fiber enhancement in dw-mri. *Journal of Mathematical Imaging and Vision* **46**(3), 326–368 (2013)
28. Bekkers, E.J., Duits, R., Mashtakov, A., Sanguinetti, G.R.: A pde approach to data-driven sub-Riemannian geodesics in se (2). *SIAM Journal on Imaging Sciences* **8**(4), 2740–2770 (2015)
29. 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)
30. Citti, G., Franceschiello, B., Sanguinetti, G., Sarti, A.: Sub-Riemannian mean curvature flow for image processing. *SIAM Journal on Imaging Sciences* **9**(1), 212–237 (2016)
31. Chambolle, A., Pock, T.: A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision* **40**(1), 120–145 (2011)
32. Chambolle, A., Pock, T.: Total roto-translational variation. *Numerische Mathematik* **142**(3), 611–666 (2019). DOI 10.1007/s00211-019-01026-w. URL <https://doi.org/10.1007/s00211-019-01026-w>
33. Smets, B.M., Portegies, J., St-Onge, E., Duits, R.: Total variation and mean curvature pdes on the homogeneous space of positions and orientations. *Journal of Mathematical Imaging and Vision* **63**(2), 237–262 (2021)
34. Duits, R., Duits, M., van Almsick, M., ter Haar Romeny, B.: Invertible orientation scores as an application of generalized wavelet theory. *Pattern Recognition and Image Analysis* **17**(1), 42–75 (2007)
35. 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. *Journal of Mathematical Imaging and Vision* **60**(9), 1427–1458 (2018). DOI 10.1007/s10851-018-0806-0. URL <https://doi.org/10.1007/s10851-018-0806-0>
36. Franceschiello, B., Mashtakov, A., Citti, G., Sarti, A.: Geometrical optical illusion via sub-Riemannian geodesics in the roto-translation group. *Differential Geometry and its Applications* **65**, 55–77 (2019)
37. Citti, G., Sarti, A.: A cortical based model of perceptual completion in the roto-translation space. *Journal of Mathematical Imaging and Vision* **24**(3), 307–326 (2006)
38. Duits, R., Franken, E.M.: Left invariant parabolic evolution equations on $SE(2)$ and contour enhancement via invertible orientation scores, part I: Linear left-invariant diffusion equations on $SE(2)$. *Quarterly of Applied mathematics, AMS* **68**, 255–292 (2010)
39. Duits, R., Franken, E.M.: Left invariant parabolic evolution equations on $SE(2)$ and contour enhancement via invertible orientation scores, part II: Nonlinear left-invariant diffusion equations on invertible orientation scores. *Quarterly of Applied mathematics, AMS* **68**, 293–331 (2010)
40. Zhang, J., Duits, R., ter Haar Romeny, B., Sanguinetti, G.: Numerical approaches for linear left-invariant diffusions on $SE(2)$, their comparisons to exact solutions, and their applications in retinal imaging. *Numerical Mathematics: Theory Methods and Applications* **9**(1), 1–50 (2016)
41. Boscain, U., Chertovskih, R.A., Gauthier, J.P., Remizov, A.O.: Hypoelliptic diffusion and human vision: A semidiscrete new twist. *SIAM Journal on Imaging Sciences* **7**(2), 669–695 (2014)
42. 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: J. Lellmann, M. Burger, J. Modersitzki (eds.) *Scale Space and Variational Methods in Computer Vision*, pp. 472–484. Springer International Publishing, Cham (2019)
43. 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)
44. Barbieri, D., Citti, G., Cocci, G., Sarti, A.: A cortical-inspired geometry for contour perception and motion integration. *Journal of Mathematical Imaging and Vision* **49**(3), 511–529 (2014)
45. Petitot, J.: The neurogeometry of pinwheels as a sub-Riemannian contact structure. *Journal of Physiology - Paris* **97**, 265–309 (2003)
46. Felsberg, M., Forssen, P.E., Scharr, H.: Channel smoothing: Efficient robust smoothing of low-level signal features. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 209–222 (2006)
47. 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)
48. 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)
49. Momayyez-Siahkal, P., Siddiqi, K.: 3d Stochastic Completion Fields for Fiber Tractography. In: Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit, pp. 178–185 (2009). DOI 10.1109/CVPRW.2009.5204044
50. Boscain, U., Prandi, D., Sacchelli, L., Turco, G.: A bio-inspired geometric model for sound reconstruction. *The Journal of Mathematical Neuroscience* **11**(1), 1–18 (2021)
51. Baspinar, E., Citti, G., Sarti, A.: A geometric model of multi-scale orientation preference maps via gabor functions. *Journal of Mathematical Imaging and Vision* **60**(6), 900–912 (2018)
52. Bekkers, E.J.: B-spline cnns on lie groups (2019)
53. Finzi, M., Stanton, S., Izmailov, P., Wilson, A.G.: Generalizing convolutional neural networks for equivariance to lie groups on arbitrary continuous data. In: H.D. III, A. Singh (eds.) Proceedings of the 37th International Conference on Machine Learning, *Proceedings of Machine Learning Research*, vol. 119, pp. 3165–3176. PMLR (2020). URL <http://proceedings.mlr.press/v119/finzi20a.html>

54. Weiler, M., Cesa, G.: General $e(2)$ -equivariant steerable cnns. In: *Advances in Neural Information Processing Systems*, pp. 14334–14345 (2019)
55. Chirikjian, G.S., Kyatkin, A.B.: *An operational calculus for the euclidean motion group with applications in robotics and polymer science*. *Journal of Fourier Analysis and Applications* **6**(6), 583–606 (2000)
56. 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)
57. Reisert, M.: *Group integration techniques in pattern analysis—a kernel view*. Ph.D. thesis, Albert-Ludwigs-Universität Freiburg (2008)
58. 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)
59. Ali, S.T., Antoine, J.P., Gazeau, J.P., et al.: *Coherent states, wavelets and their generalizations*, vol. 1. Springer (2000)
60. Sifre, L., Mallat, S.: Rigid-motion scattering for texture classification. *arXiv preprint arXiv:1403.1687* (2014)
61. Finzi, M., Bondesan, R., Welling, M.: Probabilistic numeric convolutional neural networks. *arXiv preprint arXiv:2010.10876* (2020)
62. Montobbio, N.: *A metric model of the visual cortex*. Ph.D. thesis, Università di Bologna - Sorbonne Université (2019)
63. Montobbio, N., Bonnasse-Gahot, L., Citti, G., Sarti, A.: Kercnns: biologically inspired lateral connections for classification of corrupted images. *arXiv preprint arXiv:1910.08336* (2019)
64. Weinan, E.: A proposal on machine learning via dynamical systems. *Communications in Mathematics and Statistics* **5**(1), 1–11 (2017)
65. 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)
66. Lu, Y., Zhong, A., Li, Q., Dong, B.: Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. *arXiv preprint arXiv:1710.10121* (2017)
67. 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)
68. 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)
69. Long, Z., Lu, Y., Ma, X., Dong, B.: Pde-net: Learning pdes from data. *arXiv preprint arXiv:1710.09668* (2017)
70. Jenner, E., Weiler, M.: Steerable partial differential operators for equivariant neural networks. *arXiv preprint arXiv:2106.10163* (2021)
71. Haber, E., Ruthotto, L.: Stable architectures for deep neural networks. *Inverse Problems* **34**(1), 014004 (2017)
72. Alt, T., Peter, P., Weickert, J., Schrader, K.: Translating numerical concepts for pdes into neural architectures. *arXiv preprint arXiv:2103.15419* (2021)
73. Koda, T.: An introduction to the geometry of homogeneous spaces. In: *Proceedings of The Thirteenth International Workshop on Diff. Geom*, vol. 13, pp. 121–144 (2009)
74. 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)
75. Smets, B.: *Geometric image denoising and machine learning*. Master thesis, TU Eindhoven (2019). URL <https://bmsmets.com/publication/smets2019msc/>
76. Arendt, W., Bukhvalov, A.V.: Integral representations of resolvents and semigroups. In: *Forum Mathematicum*, vol. 6, pp. 111–136. Walter de Gruyter, Berlin/New York (1994)
77. Führ, H., Fuehr, H.: *Abstract harmonic analysis of continuous wavelet transforms*. 1863. Springer Science & Business Media (2005)
78. 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. *Journal of Mathematical Imaging and Vision* pp. 1–33 (2018)
79. 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)
80. Ter Elst, A., Robinson, D.W.: Weighted subcoercive operators on Lie groups. *Journal of Functional Analysis* **157**(1), 88–163 (1998)
81. P.Maheux: Estimations du noyau de la chaleur sur les espaces homogenes. *ESAIM: Control, Optimization and Calculus of Variations* **8**(1), 65–96 (1998)
82. 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)
83. Yosida, K.: *Functional analysis*. springer (1968)
84. Evans, L.C.: *Partial differential equations*. Providence, RI (1998)
85. Balogh, Z.M., Engulatov, A., Hunziker, L., Maasalo, O.E.: Functional inequalities and hamilton–jacobi equations in geodesic spaces. *Potential analysis* **36**(2), 317–337 (2012)
86. Dragoni, F.: Metric hopf-lax formula with semicontinuous data. *Discrete & Continuous Dynamical Systems-A* **17**(4), 713 (2007)
87. Azagra, D., Ferrera, J., López-Mesas, F.: Nonsmooth analysis and Hamilton–Jacobi equations on Riemannian manifolds. *Journal of Functional Analysis* **220**(2), 304–361 (2005)
88. Rund, H.: *The Hamilton-Jacobi theory in the calculus of variations: its role in mathematics and physics*. Krieger Publishing Company (1966)
89. Manfredi, J.J., Stroffolini, B.: A version of the Hopf-Lax formula in the heisenberg group. *Communications in Partial Differential Equations* (2002)
90. Mirebeau, J.M., Portegies, J.: Hamiltonian fast marching: a numerical solver for anisotropic and non-holonomic eikonal pdes. *Image Processing On Line* **9**, 47–93 (2019)
91. Staal, J., Abràmoff, M.D., Niemeijer, M., Viergever, M.A., Van Ginneken, B.: Ridge-based vessel segmentation in color images of the retina. *IEEE transactions on medical imaging* **23**(4), 501–509 (2004). URL <https://www.isi.uu.nl/Research/Databases/DRIVE/>
92. Baweja, C.: RotNIST. <https://github.com/ChaitanyaBaweja/RotNIST> (2018)
93. 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. 8024–8035 (2019). URL pytorch.org
94. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
95. 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 computation* **1**(4), 541–551 (1989)