



Total Variation and Mean Curvature PDEs on the Homogeneous Space of Positions and Orientations

Bart M. N. Smets¹ · Jim Portegies¹ · Etienne St-Onge² · Remco Duits¹

Received: 23 October 2019 / Accepted: 31 August 2020 / Published online: 18 September 2020
© The Author(s) 2020

Abstract

Two key ideas have greatly improved techniques for image enhancement and denoising: the lifting of image data to multi-orientation distributions and the application of nonlinear PDEs such as total variation flow (TVF) and mean curvature flow (MCF). These two ideas were recently combined by Chambolle and Pock (for TVF) and Citti et al. (for MCF) for two-dimensional images. In this work, we extend their approach to enhance and denoise images of arbitrary dimension, creating a unified geometric and algorithmic PDE framework, relying on (sub-)Riemannian geometry. In particular, we follow a different numerical approach, for which we prove convergence in the case of TVF by an application of Brezis–Komura gradient flow theory. Our framework also allows for additional data adaptation through the use of locally adaptive frames and coherence enhancement techniques. We apply TVF and MCF to the enhancement and denoising of elongated structures in 2D images via orientation scores and compare the results to Perona–Malik diffusion and BM3D. We also demonstrate our techniques in 3D in the denoising and enhancement of crossing fiber bundles in DW-MRI. In comparison with data-driven diffusions, we see a better preservation of bundle boundaries and angular sharpness in fiber orientation densities at crossings.

Keywords Total variation · Mean curvature · Sub-Riemannian geometry · Roto-translations · Denoising · Fiber enhancement

1 Introduction

In the last decade, many PDE-based image-analysis techniques for tracking and enhancement of curvilinear structures in images took advantage of lifting image data, typically defined on \mathbb{R}^d , to a multi-orientation distribution (e.g., an orientation score) defined on the homogeneous space \mathbb{M}_d of d -dimensional positions and orientations, see Fig. 1 and [5,8,11,14,20,53]. After lifting the image to a multi-orientation distribution, the distribution is taken as an initial condition of a PDE flow. After solving a limited number of iterations of the PDE model, one obtains a regularized version of the original distribution, and by integration over all orientations, one obtains a regularized version of the original image.

The key advantage of lifting the images from \mathbb{R}^d to the homogeneous space \mathbb{M}_d is that the PDE flow can act differently on substructures with different orientations [5,11,24]. For instance, if the image contains two crossing lines, the PDE can regularize the two lines independently, rather than regularizing the whole crossing. Similarly, if the image contains a corner, the corner is preserved in the regularized image.

This idea of lifting images has been successfully implemented for image enhancement [13,30], geodesic tracking [6,11,51], image denoising [31], contrast perception and optical illusions [7]. For instance, Perona–Malik diffusion has been lifted to the homogeneous space \mathbb{M}_d [17] and coherence enhancing diffusion (CED) [54] has been lifted to the method of *coherence enhancing diffusion on invertible orientation scores* (CED-OS) [30] and to its 3D generalization [33].

PDE flows on orientation lifts of 3D images are relevant for applications such as fiber enhancement [17,21,46,53] and fiber tracking [45] in diffusion-weighted magnetic resonance imaging (DW-MRI), and in enhancement [33] and tracking [15] of blood vessels in 3D images.

✉ Bart M. N. Smets
bsmetsjr@hotmail.com, b.m.n.smets@tue.nl

¹ CASA, Eindhoven University of Technology, P.O. Box 513,
5600 MB Eindhoven, The Netherlands

² Sherbrooke Connectivity Imaging Lab (SCIL), Sherbrooke,
Canada

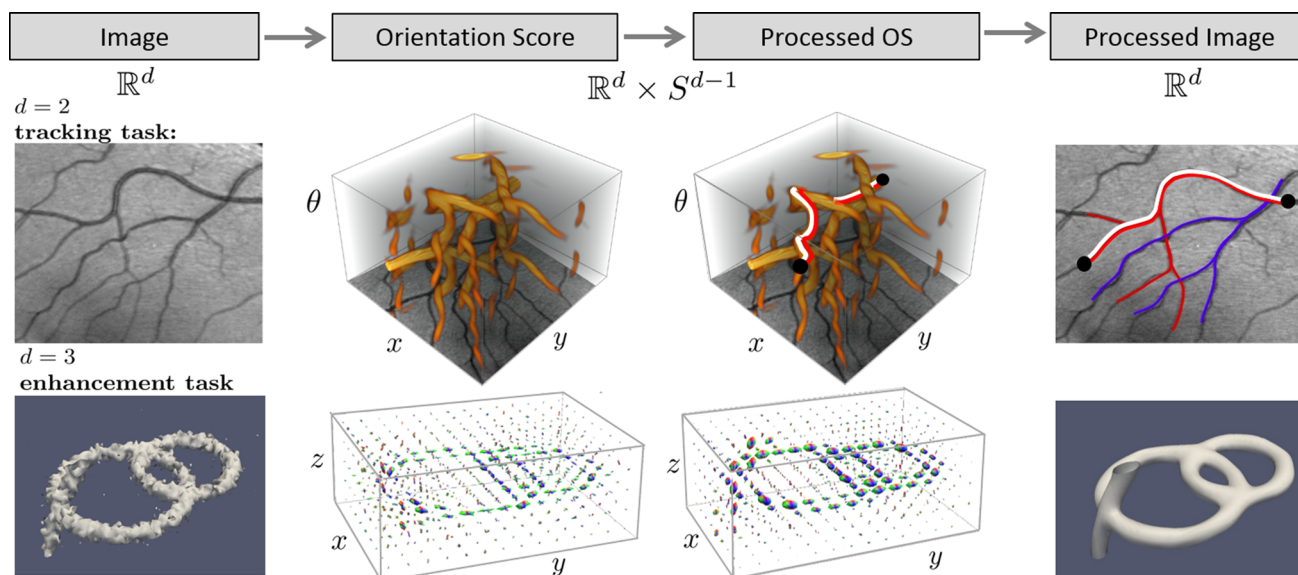


Fig. 1 Instead of direct PDE-based processing of an image, we apply PDE-based processing on a lifted image $U : \mathbb{R}^d \times S^{d-1} \rightarrow \mathbb{R}$ (e.g., an orientation score: OS). The OS is obtained by convolving the image with a set of rotated wavelets allowing for stable reconstruction [5,20,33]. 2nd row: vessel tracking in a 2D image via geodesic PDE flows in OS

that underlay TVF: [5,11,24], with $\mathbf{n} = (\cos \theta, \sin \theta)^T \in S^1$. 3rd row: CED-OS diffusion of a 3D image [23,33] visualized as a field of angular profiles (see Remark 1). In this article we study image enhancement and denoising via TVF and MCF on $\mathbb{M}_d = \mathbb{R}^d \times S^{d-1}$ and compare to nonlinear diffusion methods

The general workflow is illustrated in Fig. 1. The original image is described by a function $f : \Omega_f \rightarrow \mathbb{R}^+$, where $\Omega_f \subset \mathbb{R}^d$ is its support. From $f \in \mathbb{L}_2(\Omega_f)$, one computes an orientation lift $U : \mathbb{M}_d \rightarrow \mathbb{C}$, compactly supported within

$$\Omega = \Omega_f \times S^{d-1} \subset \mathbb{M}_d. \quad (1)$$

There are various ways to construct such a lift: it can be (the real part of) an invertible orientation score [22] (cf. Fig. 1), a channel representation [28], a lift by Gabor wavelets [3], or a fiber orientation density [44]. In all of these approaches the absolute value $|U(\mathbf{x}, \mathbf{n})|$ can be regarded (after normalization) as a probability density of finding a fiber structure at position $\mathbf{x} \in \mathbb{R}^d$ with local orientation $\mathbf{n} \in S^{d-1}$. We set the orientation lift U as an initial condition of a PDE flow $U \mapsto \Phi_t(U)$ with evolution time $t > 0$. Finally, the processed multi-orientation representation $\Phi_t(U)$ is integrated over all orientations to obtain the enhanced image f_t . In this article, we will work with the orientation score, with the main motivation being that this operation is invertible [20], so that when taking $t \downarrow 0$, the output equals the input, i.e., $\lim_{t \downarrow 0} f_t = f$ in \mathbb{L}_2 -sense.

The enhanced image that one obtains after running a PDE flow, (the bottom-right picture in Fig. 1), naturally depends on the type of flow used. One flow may be more suitable than another, depending on the requirements imposed on the resulting image. In case it is important to preserve sharp tran-

sitions in the image, while maintaining plateaus, nonlinear flows such as total variation flows (TVF) and mean curvature flow (MCF) [49] are typically more suited than nonlinear diffusion flows (such as Perona and Malik diffusion [42] and coherence-enhancing diffusion [54]).

For $d = 2$, TVF and MCF were recently generalized to lifted images by Chambolle and Pock [11] and Citti et al. [13], respectively.

Their promising results have motivated us to generalize TVF and MCF to lifted images for general dimension d and provide a general geometric and algorithmic framework that can accommodate features such as locally adaptive frames and coherence enhancement.

The benefits of our approach are that we obtain a single unifying geometric and algorithmic framework for arbitrary d , with efficient algorithms (for $d = 2, 3$) that preserve crossing lines, corners, plateaus, edges and bundle boundaries and can improve curvature adaptation via the optional inclusion of locally adaptive frames. Such frames account for curvature of lines and allow us to remove bias toward sampled orientations in orientation scores.

Our PDE methods on \mathbb{M}_d are computationally more expensive than their counterparts acting only on \mathbb{R}^d , but they are still practical. Similar to crossing-preserving nonlinear diffusion on $SE(2) \equiv \mathbb{M}_2$, locally adaptive frames allow us to remove orientation sampling bias in orientation scores [29, Fig. 6.11] and to use only 4 orientation samples [30]. For our crossing-preserving MCF and TVF on \mathbb{M}_2 we sample our

Table 1 We formulate our PDEs independently of our choice of frame and metric and then apply it to both the left-invariant and gauge geometries

Choice of geometry	Left invariant	Gauge
Frame	$\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_{2d-1}$	$\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_{2d-1}$
Co-frame	$\omega^1, \omega^2, \dots, \omega^{2d-1}$	$\beta^1, \beta^2, \dots, \beta^{2d-1}$
Metric tensor	\mathcal{G}	\mathcal{J}
Gradient ∇	$\nabla_{\mathcal{G}}$	$\nabla_{\mathcal{J}}$
Divergence div (Riemannian)	$\sum_{i=1}^{2d-1} \mathcal{A}_i \circ \omega^i$	$\sum_{i=1}^{2d-1} \mathcal{B}_i \circ \beta^i$
Divergence div (sub-Riemannian)	$\sum_{i=d}^{2d-1} \mathcal{A}_i \circ \omega^i$	$\sum_{i=d}^{2d-1} \mathcal{B}_i \circ \beta^i$
Norm $\ \cdot\ $	$\ \cdot\ _{\mathcal{G}}$	$\ \cdot\ _{\mathcal{J}}$
Volume form $d\tilde{\mu}$	$\frac{\omega^1 \wedge \omega^2 \wedge \dots \wedge \omega^{2d-1}}{\sqrt{D_S} \sqrt{D_S D_A^{d-1}}}$	$\beta^1 \wedge \beta^2 \wedge \dots \wedge \beta^{2d-1}$

(processed) orientation scores only on 8 orientations. On \mathbb{M}_3 we compute regularized orientation lifts on a grid with 162 orientations, where we rely on efficient numerical schemes for PDEs on \mathbb{M}_3 relying on the low-order PDE discretization schemes explained in [17,37], instead of higher-order schemes via spherical harmonics [34, Ch. 3.4], in order to reduce computation time.

The structure of this article is as follows. We start by recapitulating orientation scores and explaining the homogeneous space \mathbb{M}_d as a Lie group quotient in the rigid body motion group $\text{SE}(d)$ in Sect. 2 and explain the necessary geometric concepts. In Sect. 3, we introduce the PDEs for total variation and mean curvature flow on \mathbb{M}_d and explain our explicit discretization scheme. Our numerical scheme includes regularization for which we prove convergence in Sect. 4. In Sect. 5 we evaluate the potential of our methods with 2D and 3D experiments.

Remark 1 (Visualization of 3D orientation scores) In the 3rd row of Fig. 1, and henceforth, we visualize a lifted image $U : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}^+$ by a grid of angular profiles $\{\mu U(\mathbf{x}, \mathbf{n}) \mid \mathbf{x} \in \mathbb{Z}^3, \mathbf{n} \in S^2\}$, with fixed $\mu > 0$.

Remark 2 (Additional content in this version) This article is an extended version of the authors’ SSVM article by the same name [25]. The following content is new:

- A coordinate-free formulation of gauge-frame fitting in Sect. 2.4 that generalizes our previous coordinate-dependent approach of [23].
- An introduction of distinct geometric setups (with or without locally adaptive frames) for any dimension of \mathbb{M}_d that admits the formulation of TVF and MCF PDEs. A quick overview of the two distinct geometric approaches is provided in Table 1.
- A general formulation of our “coherence enhancement technique” for TVF and MCF on \mathbb{M}_d in Sect. 2.5.
- A proof for the theorem of the strong convergence, stability and accuracy of TV flows. This result was announced in [25] but not yet proven.

Table 2 The two geometries we developed on $\text{SE}(d)$ can be naturally applied to the homogeneous space \mathbb{M}_d using the natural extension of scalar functions (39) and vector fields (40)

\mathbb{M}_d tools	Left invariant	Gauge
$\nabla U(\mathbf{x}, \mathbf{n}) =$	$(\nabla_{\mathcal{G}} \tilde{U})(\mathbf{x}, \mathbf{R}_n)$	$(\nabla_{\mathcal{J}} \tilde{U})(\mathbf{x}, \mathbf{R}_n)$
$(\mathbf{v}, \mathbf{w})_{\mathbb{M}_d} \Big _{(\mathbf{x}, \mathbf{n})} =$	$\mathcal{G} \Big _{(\mathbf{x}, \mathbf{R}_n)}(\tilde{\mathbf{v}}, \tilde{\mathbf{w}})$	$\mathcal{J} \Big _{(\mathbf{x}, \mathbf{R}_n)}(\tilde{\mathbf{v}}, \tilde{\mathbf{w}})$
$\ \mathbf{v}\ _{\mathbb{M}_d} =$	$\ \tilde{\mathbf{v}}\ _{\mathcal{G}}$	$\ \tilde{\mathbf{v}}\ _{\mathcal{J}}$
$(\text{div} \mathbf{v})(\mathbf{x}, \mathbf{n}) =$	$(\text{div}_{\mathcal{G}} \tilde{\mathbf{v}})(\mathbf{x}, \mathbf{R}_n)$	$(\text{div}_{\mathcal{J}} \tilde{\mathbf{v}})(\mathbf{x}, \mathbf{R}_n)$
$d\mu(\mathbf{x}, \mathbf{n}) =$	$d\tilde{\mu}_{\mathcal{G}}(\mathbf{x}, \mathbf{R}_n)$	$d\tilde{\mu}_{\mathcal{J}}(\mathbf{x}, \mathbf{R}_n)$
$(E\mathbf{v})(\mathbf{x}, \mathbf{n}) =$	$(E_{\mathcal{G}} \tilde{\mathbf{v}})(\mathbf{x}, \mathbf{R}_n)$	$(E_{\mathcal{J}} \tilde{\mathbf{v}})(\mathbf{x}, \mathbf{R}_n)$

These objects are well defined despite the noninjective mapping used since the choice of particular $\mathbf{R}_n \in \text{SE}(d)$ that maps \mathbf{a} to \mathbf{n} does not affect any of them

- Extensions of our 2D denoising/enhancing experiments, Sect. 5, Figs. 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, and Table 4. These experiments now include a full comparison of isotropic and anisotropic processing, and the effect of the including coherence enhancement (via locally adaptive frames) in TVF and MCF. They also include additional comparisons to Perona–Malik [42] diffusion and to a well-established denoising method: BM3D [18,35].
- A more comprehensive treatment of the geometric tools used such as vector fields and metric tensors. We now clearly distinguish between the group $\text{SE}(d)$ and the homogeneous space \mathbb{M}_d . We also explain how to transfer geometric tools on these sets in Sect. 2.6 and Table 2.

2 Preliminary Theory

Before we can provide the generalized PDEs, which include TVF and MCF as special cases, we need to construct the necessary tools.

In this section, we review orientation scores, the rigid-body motion group $\text{SE}(d)$, and the homogeneous space \mathbb{M}_d of

positions and orientations. For further reading on engineering applications and harmonic analysis on Lie group $SE(d)$ we refer to [12, Ch. 6]. For theory on homogeneous spaces we refer to [36, Ch. 21]. For image processing on $SE(2)$, see for example [5, 8, 11, 14], for image processing on $SE(3)$, see for example [39, 45, 47].

2.1 Orientation Scores: Lifting the Image Domain from \mathbb{R}^d to \mathbb{M}_d

In order to disentangle all local orientations in an image we lift the data from position space \mathbb{R}^d to the homogeneous space \mathbb{M}_d of positions and orientations. This means that we extend the domain of an image. See Fig. 1, where we lift the data from \mathbb{R}^d toward \mathbb{M}_d via invertible orientation scores.

Building an orientation score starts with selecting an orientation-sensitive filter (or wavelet) $\psi \in \mathbb{L}_1 \cap \mathbb{L}_2(\mathbb{R}^d)$. We can then (under appropriate conditions [20, 33]) filter out a particular direction from an image $f \in \mathbb{L}_2(\mathbb{R}^d)$ by convolving the image with this filter aligned to that direction. An orientation score $\mathcal{W}_\psi f$ can then be constructed by applying this filtering for all directions $\mathbf{n} \in S^{d-1}$:

$$\mathcal{W}_\psi f(\mathbf{x}, \mathbf{n}) = \int_{\mathbb{R}^d} \overline{\psi(\mathbf{R}_\mathbf{n}^{-1}(\mathbf{y} - \mathbf{x}))} f(\mathbf{y}) \, d\mathbf{y}, \tag{2}$$

for all $\mathbf{x} \in \mathbb{R}^d$ and rotations $\mathbf{R}_\mathbf{n}$ that map a reference axis $\mathbf{a} \in S^{d-1}$ to \mathbf{n} .

For this paper we will be using *cake wavelets* [20, 33] for our filter ψ , illustrated in Fig. 2 for $d = 2$. These wavelets are directional filters that have the property that we can accurately reconstruct the original image from the orientation score (again under appropriate conditions) by integration over S^{d-1} , i.e.,

$$f(\mathbf{x}) \approx \int_{S^{d-1}} \mathcal{W}_\psi f(\mathbf{x}, \mathbf{n}) \, d\sigma(\mathbf{n}), \tag{3}$$

where σ denotes the usual surface measure over S^{d-1} . We always use standard cake wavelet parameter settings from [37] in our experiments.

The explicit formulas for these cake wavelets that allow invertible orientation scores are available in [5, 20] and specifically for $d = 3$ in [33]. An intuitive illustration of an orientation score is given in Fig. 2.

2.2 $SE(d)$ and the Homogeneous Space of Positions and Orientations \mathbb{M}_d

Consider the rigid body motion group $SE(d) = \mathbb{R}^d \rtimes SO(d)$, the semidirect product of the translation group \mathbb{R}^d and the rotation group $SO(d)$ of orthogonal $d \times d$ matrices. We call




Fig. 2 Top: cake wavelets [20] for $d = 2$ used to directionally filter an image and construct an orientation score. Bottom: How orientation scores disentangle orientations

elements of $SE(d)$ *roto-translations*. The product of two roto-translations $g_i = (\mathbf{x}_i, \mathbf{R}_i)$ in $SE(d)$ is given by

$$g_1 g_2 = (\mathbf{x}_1, \mathbf{R}_1) (\mathbf{x}_2, \mathbf{R}_2) := (\mathbf{x}_1 + \mathbf{R}_1 \mathbf{x}_2, \mathbf{R}_1 \mathbf{R}_2). \tag{4}$$

These roto-translations act transitively on the space $\mathbb{R}^d \times S^{d-1}$ by

$$(\mathbf{x}, \mathbf{R}) \odot (\mathbf{y}, \mathbf{n}) := (\mathbf{x} + \mathbf{R}\mathbf{y}, \mathbf{R}\mathbf{n}), \tag{5}$$

for all $(\mathbf{y}, \mathbf{n}) \in \mathbb{R}^d \times S^{d-1}$ and all roto-translations $(\mathbf{x}, \mathbf{R}) \in SE(d)$.

We choose an *a priori reference vector* $\mathbf{a} \in S^{d-1}$, say $\mathbf{a} = (1, 0)^T$ if $d = 2$ or $\mathbf{a} = (0, 0, 1)^T$ if $d = 3$. Then the stabilizer of the element $(\mathbf{0}, \mathbf{a})$ is given by

$$H_d := \{g \in SE(d) \mid g \odot (\mathbf{0}, \mathbf{a}) = (\mathbf{0}, \mathbf{a})\}, \tag{6}$$

which is isomorphic to $SO(d - 1)$.

The homogeneous space of positions and orientations is the partition of left cosets

$$\mathbb{M}_d := \mathbb{R}^d \times S^{d-1} := SE(d)/H_d.$$

The left cosets are equivalence classes in $SE(d)$ with respect to the equivalence relation

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

For $d = 2$, the subgroup $H_2 = \{(\mathbf{0}, \mathbf{I})\}$ consists only of the unit element, and thereby the manifold \mathbb{M}_2 is diffeomorphic to $SE(2)$. However, for $d > 2$ the manifolds \mathbb{M}_d and $SE(d)$ are not diffeomorphic.

For $d = 3$, the stabilizer can be described by

$$H_3 = \left\{ \left| h_\alpha := (\mathbf{0}, \mathbf{R}_{\mathbf{a}, \alpha}) \mid \alpha \in [0, 2\pi) \right. \right\},$$

where $\mathbf{R}_{\mathbf{a}, \alpha}$ denotes a (counter-clockwise) rotation over angle α around the reference axis \mathbf{a} . This means that two roto-

translations $g_1 = (\mathbf{x}_1, \mathbf{R}_1)$ and $g_2 = (\mathbf{x}_2, \mathbf{R}_2)$ are equivalent if and only if

$$\mathbf{x}_1 = \mathbf{x}_2 \text{ and } \exists \alpha \in [0, 2\pi) : \mathbf{R}_1 = \mathbf{R}_2 \mathbf{R}_{\mathbf{a}, \alpha}.$$

The equivalence classes $[g] = \{g' \in \text{SE}(3) \mid g' \sim g\}$ are usually just denoted by $p = (\mathbf{x}, \mathbf{n})$ as they consist of all rigid body motions $g = (\mathbf{x}, \mathbf{R}_n)$ that map the reference point $(\mathbf{0}, \mathbf{a})$ onto $(\mathbf{x}, \mathbf{n}) \in \mathbb{R}^3 \times S^2$:

$$g \odot (\mathbf{0}, \mathbf{a}) = (\mathbf{x}, \mathbf{n}).$$

Remark 3 (Distinguishing $\text{SE}(d)$ from \mathbb{M}_d) As the distinction between the group $\text{SE}(d)$ and the homogeneous space \mathbb{M}_d (which is not a group for $d = 3$ and above) is important, we will use g, h for elements of $\text{SE}(d)$, and p, q for points in \mathbb{M}_d .

To understand why the situation changes from $d = 2$ to $d > 2$ observe that in 2 dimensions we need one angle to specify orientation and have one rotational degree of freedom, in 3 dimensions we need 2 angles to specify orientation, but we have 3 rotational degrees of freedom, i.e., \mathbb{M}_3 has one dimension less than $\text{SE}(3)$. See also Fig. 3 for an illustration of this difference.

For $d > 3$ this situation persists as we have more rotational degrees of freedom that do not change the orientation.

Remark 4 (Domain of an orientation score) The orientation score is well defined on the domain \mathbb{M}_d if we assume ψ is not affected under the action of subgroup H_d . For $d = 3$ this means we must impose axial symmetry on the wavelets, for details see [33].

2.3 Differential Structure on $\text{SE}(d), \mathbb{M}_d$

As a manifold, we view the group $\text{SE}(d)$ in a standard way as a submanifold of $\mathbb{R}^d \times \mathbb{R}^{d \times d}$. The Lie algebra is, as a vector space, the tangent space at the unit element (see [36, Ch. 7]). We view elements of tangent spaces (i.e., tangent vectors) as derivations acting on functions: If v is an ordinary vector in $\mathbb{R}^d \times \mathbb{R}^{d \times d}$ tangent to $\text{SE}(d)$, the corresponding derivation acting on a function $f \in C^1(\text{SE}(d))$ is just the derivative of f in the direction of v .

The Lie algebra has dimension $D = \frac{1}{2}d(d+1)$. We choose a basis $(A_i)_{i=1}^D$ for the Lie algebra of $\text{SE}(d)$ with the following properties. The basis is orthonormal with respect to the inner product belonging to the standard Euclidean metric on $\mathbb{R}^d \times \mathbb{R}^{d \times d}$; the vectors $\{A_1, \dots, A_d\}$ span the spatial part of the Lie algebra, which is isomorphic to \mathbb{R}^d with the vector A_d corresponding to the derivative in the direction of \mathbf{a} . Recall that for $d = 2$ the subgroup H_d is trivial. For $d \geq 3$ one has that the set $\{A_{2d}, \dots, A_D\}$ forms a basis for the stabilizer subgroup H_d . We take the convention that the Lie




Fig. 3 Illustration of the Lie algebra of left-invariant vector fields $(A_i)_{i=1}^D$ of $\text{SE}(d)$ for left: $d = 2$ and right: $d = 3$, note the orientation-preserving rotation generator A_6

algebra vector A_{i+d} generates the in-plane rotation in the plane spanned by A_i and \mathbf{a} for $i = 1, \dots, d - 1$.

For the case $d = 2$ this gives us two spatial generators, A_1 and A_2 , and one rotation generator A_3 . Moving to $d = 3$ we have 3 spatial degrees of freedom and 3 rotational degrees of freedom, but only 2 of those rotational degrees of freedom will change the orientation, we denote the generator corresponding to the rotation that preserves the reference axis by A_6 . This gives us the following basis:

$$\left\{ \underbrace{\left[A_1, A_2, A_3 \right]}_{\text{spatial generators}}, \underbrace{\left[A_4, A_5 \right]}_{\text{generators of rotations that move } \mathbf{a}}, \underbrace{\left[A_6 \right]}_{\text{generator of rotation that fixes } \mathbf{a}} \right\}.$$

As is illustrated in Fig. 3 for $d = 3$ we have a rotational degree of freedom that does not change the orientation reference axis. As a result \mathbb{M}_3 is not isomorphic to $\text{SE}(3)$. It is rather a 5-dimensional quotient of the 6-dimensional Lie group $\text{SE}(3)$, see also Remark 3.

Remark 5 (Generalization for $d > 3$) Generalizing this scheme for $d > 3$ we would have the following basis for the Lie algebra:

$$\left\{ \underbrace{\left[A_1, \dots, A_d \right]}_{\text{spatial generators}}, \underbrace{\left[A_{d+1}, \dots, A_{2d-1} \right]}_{\text{generators of rotations that move } \mathbf{a}}, \underbrace{\left[A_{2d}, \dots, A_D \right]}_{\text{generators of rotations that fix } \mathbf{a}} \right\}.$$

We extend the vectors A_i to left-invariant vector fields \mathcal{A}_i as follows. The group acts on itself by left multiplication,

$$L_g h := gh, \quad \forall g, h \in \text{SE}(d), \tag{7}$$

and the derivation $(\mathcal{A}_i)_g$, evaluated in a point g , is given by the pushforward

$$\mathcal{A}_i|_g f := ((L_g)_* \mathcal{A}_i) f = A_i(f \circ L_g), \tag{8}$$

for all $f \in C^\infty(\text{SE}(d), \mathbb{R})$. We denote the corresponding covector fields by $\omega^i : g \mapsto \omega^i|_g$. For each $g \in \text{SE}(d)$, the covector $\omega^i|_g$ is an element of the dual to the tangent space of $\text{SE}(d)$ at g . The covector fields are characterized by

$$\langle \omega^i, \mathcal{A}_j \rangle = \delta_j^i, \tag{9}$$

where δ_j^i denotes the Kronecker delta.

Note that

$$A_d = \mathbf{a} \cdot \nabla_{\mathbb{R}^d} \Big|_{\mathbf{0}} = \sum_{i=1}^d a^i \frac{\partial}{\partial x^i} \Big|_{\mathbf{0}}, \tag{10}$$

and so for all $g = (\mathbf{x}, \mathbf{R}_n) \in \text{SE}(d)$:

$$\begin{aligned} \mathcal{A}_d|_g &= (L_g)_* A_d = \sum_{i=1}^d (\mathbf{R}_n \mathbf{a})^i \frac{\partial}{\partial x^i} \Big|_{\mathbf{x}} \\ &= \sum_{i=1}^d n^i \frac{\partial}{\partial x^i} \Big|_{\mathbf{x}} = \mathbf{n} \cdot \nabla_{\mathbb{R}^d} \Big|_{\mathbf{x}}, \end{aligned} \tag{11}$$

from which we infer that the left-invariant frame is aligned with the direction $\mathbf{n} \in S^{d-1}$.

Remark 6 (Left-Invariant Basis in 2D) We can represent an element of $\text{SE}(2)$ by its position and angle as $(\mathbf{x}, \theta) \in \mathbb{R}^2 \times [0, 2\pi)$ which allows us to write the left-invariant vector fields \mathcal{A}_i as:

$$\begin{aligned} \mathcal{A}_1|_{(\mathbf{x}, \theta)} &= -\sin \theta \partial_x|_{\mathbf{x}} + \cos \theta \partial_y|_{\mathbf{x}}, \\ \mathcal{A}_2|_{(\mathbf{x}, \theta)} &= \cos \theta \partial_x|_{\mathbf{x}} + \sin \theta \partial_y|_{\mathbf{x}} = \mathbf{n} \cdot \nabla_{\mathbb{R}^2} \Big|_{\mathbf{x}}, \\ \mathcal{A}_3|_{(\mathbf{x}, \theta)} &= \partial_\theta. \end{aligned}$$

For an explicit form of the left-invariant vector fields \mathcal{A}_i in case $d = 3$, see ‘‘Appendix A.’’

We introduce the following metric tensor field in terms of the left-invariant co-vector fields $(\omega^i)_{i=1}^{2d-1}$.

Definition 1 (Left-invariant metric tensor field) Given positive constants $D_S > 0$ and $D_A > 0$, and a nonnegative real number $\epsilon \geq 0$, we define the left-invariant metric tensor field \mathcal{G} by

$$\mathcal{G} = \begin{cases} \frac{\omega^d \otimes \omega^d}{D_S} + \frac{\sum_{i=1}^{d-1} \omega^i \otimes \omega^i}{\epsilon^2 D_S} + \frac{\sum_{i=d+1}^{2d-1} \omega^i \otimes \omega^i}{D_A} & \text{if } \epsilon > 0, \\ \frac{\omega^d \otimes \omega^d}{D_S} + \frac{\sum_{i=d+1}^{2d-1} \omega^i \otimes \omega^i}{D_A} & \text{if } \epsilon = 0. \end{cases} \tag{12}$$

Remark 7 (Sub-Riemannian case) Henceforth we refer to $\epsilon = 0$ as the sub-Riemannian case where tangent vectors are constrained to the span of $\mathcal{A}_d, \dots, \mathcal{A}_{2d-1}$. Intuitively when $\epsilon \downarrow 0$ the other tangent directions get infinite cost and

become prohibited. This means that we restrict ourselves to so-called horizontal tangent vectors:

$$(\dot{\mathbf{x}}, \dot{\mathbf{n}}) \in T_{(\mathbf{x}, \mathbf{n})}(\mathbb{M}_d) \text{ is horizontal} \iff \mathbf{n} \wedge \dot{\mathbf{x}} = 0. \tag{13}$$

Observe that this sub-Riemannian metric tensor is defined (and invertible) on a sub-bundle of the tangent bundle on the group as it does not contain any of the covectors dual to the sub-bundle induced by subgroup H_d . Furthermore it is spatially isotropic orthogonal to the primary spatial direction. Also spherically we impose isotropy in the metric as can be seen from the last term in the above definition.

This metric induces an associated norm: If $\dot{g} \in T_g(\text{SE}(d))$, then

$$\|\dot{g}\|_{\mathcal{G}} := \sqrt{\mathcal{G}|_g(\dot{g}, \dot{g})}, \tag{14}$$

where again in the sub-Riemannian case we only allow \dot{g} to be in the span of $\mathcal{A}_d, \dots, \mathcal{A}_{2d-1}$.

Now that we have $\text{SE}(d)$ equipped with a (sub)-Riemannian metric tensor, we can derive the basic tools that are required to formulate our geometric PDEs. These basic tools include the gradient, its norm, and the divergence of a vector field. Let us relabel our parameters as

$$D_i = \begin{cases} D_S & \text{if } i = d, \\ \epsilon^2 D_S & \text{if } 1 \leq i \leq d - 1, \\ D_A & \text{if } d + 1 \leq i \leq 2d - 1. \end{cases} \tag{15}$$

Let $\tilde{U} : \text{SE}(d) \rightarrow \mathbb{R}$ carry the axial symmetry:

$$\tilde{U}(\mathbf{x}, \mathbf{R}) = U(\mathbf{x}, \mathbf{R}\mathbf{a}), \tag{16}$$

for some $U : \mathbb{R}^d \times S^{d-1} \rightarrow \mathbb{R}$.

Then in the Riemannian setting the gradient of a differentiable function $\tilde{U} : \text{SE}(d) \rightarrow \mathbb{R}$ on the group induced by this metric tensor becomes

$$\nabla_{\mathcal{G}} \tilde{U} = \sum_{i=1}^{2d-1} D_i (\mathcal{A}_i \tilde{U}) \mathcal{A}_i, \tag{17}$$

where the sum only runs to $2d - 1$ and not to $\dim(\text{SE}(d)) = D = \frac{1}{2}d(d + 1)$ since (16) implies that

$$\mathcal{A}_i \tilde{U} = 0 \text{ for } 2d - 1 < i \leq D. \tag{18}$$

The gradient then has the following norm

$$\|\nabla_{\mathcal{G}} \tilde{U}\|_{\mathcal{G}}^2 = \sum_{i=1}^{2d-1} D_i (\mathcal{A}_i \tilde{U})^2. \tag{19}$$

The divergence of a vector field is given by

$$\operatorname{div}_{\mathcal{G}} \mathbf{v} = \sum_{i=1}^{2d-1} \mathcal{A}_i v^i \quad \text{for all } \mathbf{v} = \sum_{i=1}^{2d-1} v^i \mathcal{A}_i.$$

In the sub-Riemannian setting, where we restrict ourselves to vector fields spanned by $(\mathcal{A}_i)_{i=d}^{2d-1}$, we have

$$\begin{aligned} \nabla_{\mathcal{G}} \tilde{U} &= \sum_{i=d}^{2d-1} D_i (\mathcal{A}_i \tilde{U}) \mathcal{A}_i, \text{ and} \\ \operatorname{div}_{\mathcal{G}} \mathbf{v} &= \sum_{i=d}^{2d-1} \mathcal{A}_i v^i \quad \text{for all } \mathbf{v} = \sum_{i=d}^{2d-1} v^i \mathcal{A}_i. \end{aligned} \tag{20}$$

2.4 Locally Adaptive Frames on $SE(d)$ as SVD of the Hessian

As an alternative to the left-invariant frame we can choose a frame (and subsequently a metric tensor field) that is adapted to the data (which we also refer to as *gauge frames* in analogy with [23]). Specifically, instead of having the vector field $\mathcal{A}_d = \mathbf{n} \cdot \nabla$ as a static forward direction we want to choose a vector field \mathcal{B}_d that locally aligns with the data [23]. In particular, \mathcal{B}_d can take on an angular component, meaning the local “straight forward” will follow the curve of the data; consequently, flows can better follow curved structures, see Fig. 4 for an example.

Remark 8 (Fitting a frame) We can induce an entire frame in $SE(d)$ from a choice of main vector, see [23, Appendix A] for details. For an intuitive illustration see Fig. 4. In this article we will focus on the method by which the main gauge vector is obtained.

Next we will present a singular value decomposition of the Hessian; we will choose the eigenvector associated with the smallest eigenvalue as \mathcal{B}_d . Geometrically, this can be seen as the direction in which the gradient changes the least. Before we can formulate this procedure we explain the concept of exponential curves (see Fig. 5).

Definition 2 (Exponential curve) Let $\dot{g} \in T_g(SE(d))$ then the exponential curve parameterized by t through g with tangent vector \dot{g} is written as $e_g^{\dot{g}t}$ and is the curve for which $e_g^{\dot{g}0} = g$ and which has the property that for all $t \in \mathbb{R}$:

$$\left(\frac{d}{dt} e_g^{\dot{g}t} \right) (t) = \sum_{i=1}^{2d-1} \omega^i|_g(\dot{g}) \mathcal{A}_i|_{e_g^{\dot{g}t}}.$$




Fig. 4 Illustrating gauge frame fitting. Top: left-invariant frame where $\mathcal{A}_d = \mathbf{n} \cdot \nabla_{\mathbb{R}^2}$, recall (11), the red line indicates the exponential curve associated with \mathcal{A}_2 . Bottom: we choose a frame that takes into account the local curvature; here the green line indicates the exponential curve associated with \mathcal{B}_2

Or more explicitly in coordinates, if $\dot{g} = \sum_{i=1}^{2d-1} \dot{g}^i \mathcal{A}_i|_g$ we have that:

$$\left(\frac{d}{dt} e_g^{\dot{g}t} \right) (t) = \sum_{i=1}^{2d-1} \dot{g}^i \mathcal{A}_i|_{e_g^{\dot{g}t}}.$$

Hence the exponential curves are those curves whose tangent vector components with respect to the left-invariant frame do not change. For an illustration of such curves for the case $d = 2$ see Fig. 5.

In view of (16) and (18) we define

$$T(g) := \operatorname{span} \{ \mathcal{A}_1, \dots, \mathcal{A}_{2d-1} \} \subset T_g(SE(d)). \tag{21}$$

Now we want to select $\mathcal{B}_d|_g$ (normalized with respect to the existing metric tensor $\mathcal{G}|_g$) so that the gradient of the data $\tilde{U} \in C^1(SE(d), \mathbb{R})$ changes as little as possible (recall Fig. 4) in the following manner.

Definition 3 (Main gauge vector) We define the main gauge vector as

$$\mathcal{B}_d|_g := \operatorname{argmin}_{\substack{\dot{g} \in T(g), \\ \|\dot{g}\|_{\mathcal{G}} = 1}} \left\| \operatorname{Hess}|_g(\tilde{U})(\dot{g}, \cdot) \right\|_{\mathcal{G}^*}^2$$




Fig. 5 A set of exponential curves through a common point for $d = 2$. Exponential curves are those curves whose tangent vectors are part of a left-invariant vector field. Taken with permission from [5]

$$= \operatorname{argmin}_{\substack{\dot{g} \in T(g) \\ \|\dot{g}\|_{\mathcal{G}}=1}} \left\| \frac{d}{dt} \left[(L_{g e_g^{-t} \dot{g}})_* \nabla_{\mathcal{G}} \tilde{U} (e^{t \dot{g}}) \right] \Big|_{t=0} \right\|_{\mathcal{G}}^2,$$

where we assume that \tilde{U} is such that we have a unique minimizer. The Hessian in the previous equation is induced by a Cartan connection as outlined in [23, Appendix 4, (133)].

Writing the tangent vector in terms of the local left-invariant frame as $\dot{g} = \sum_{i=1}^{2d-1} \dot{g}^i \mathcal{A}_i|_g \in T(g)$ let us write out the Hessian as follows:

$$\begin{aligned} \mathcal{B}_d|_g &= \operatorname{argmin}_{\substack{\dot{g} \in T(g) \\ \|\dot{g}\|_{\mathcal{G}}=1}} \sum_{i=1}^{2d-1} D_i \left(\dot{g}^i (\mathcal{A}_i \tilde{U}) \right)^2 \\ &= \operatorname{argmin}_{\substack{\dot{g} \in T(g) \\ \|\dot{g}\|_{\mathcal{G}}=1}} \sum_{i=1}^{2d-1} D_i \left(\sum_{j=1}^{2d-1} \dot{g}^j \mathcal{A}_j (\mathcal{A}_i \tilde{U}) \right)^2. \end{aligned} \tag{22}$$

We can write this problem in terms of matrices by defining the following:

$$\begin{aligned} \dot{g} &= \left(\dot{g}^1, \dot{g}^2, \dots, \dot{g}^{2d-1} \right)^T, \\ M &= \operatorname{diag} \left(\sqrt{D_1}, \sqrt{D_2}, \dots, \sqrt{D_{2d-1}} \right), \\ K &= K(\tilde{U})|_g := \left(\mathcal{A}_j|_g (\mathcal{A}_i|_g \tilde{U}) \right)_{i,j}. \end{aligned}$$

with i as row index and j as column index.

Using these the objective function in (22) becomes

$$(MK\dot{g})^T (MK\dot{g}) = \dot{g}^T K^T M^2 K \dot{g}, \tag{23}$$

which we want to minimize under the constraint

$$\sum_{i=1}^{2d-1} \frac{(\dot{g}^i)^2}{D_i} = \|M^{-2} \dot{g}\|^2 = 1. \tag{24}$$

Taking the derivative of the Lagrangian of this convex optimization problem gives us optimality under the following condition ($\lambda \in \mathbb{R}$):

$$K^T M^2 K \dot{g} = \lambda M^{-2} \dot{g} \iff M^2 K^T M^2 K \dot{g} = \lambda \dot{g}, \tag{25}$$

i.e., \dot{g} needs to be an eigenvector of the matrix $M^2 K^T M^2 K$ with eigenvalue λ (serving as the Lagrangian multiplier). If for a moment we rewrite (25) as

$$(MKM)^T (MKM) (M^{-1} \dot{g}) = \lambda (M^{-1} \dot{g}), \tag{26}$$

we see that λ is indeed real since $(MKM)^T (MKM)$ is symmetric. With this eigenvalue and vector objective function (23) evaluates to

$$\begin{aligned} \dot{g}^T K^T M^2 K \dot{g} &= \dot{g}^T M^{-2} M^2 K^T M^2 K \dot{g} \\ &= \dot{g}^T M^{-2} \lambda \dot{g} \\ &= \lambda \dot{g}^T M^{-2} \dot{g} \\ &= \lambda \|\dot{g}\|_{\mathcal{G}}^2 = \lambda. \end{aligned}$$

This last equation incidentally proves that $M^2 K^T M^2 K$ is positive semidefinite and, more importantly, that to minimize the change in gradient we need to look at the eigenvector belonging to the smallest eigenvalue.

In practice, we do not immediately calculate the eigenvectors and eigenvalues from the scheme we have just proposed, but for the purpose of stability we first apply a component-wise Gaussian smoothing on the matrix K as follows:

$$\begin{aligned} \tilde{K}|_g &:= (G * K)(x, R_n) \\ &= \int_{\mathbb{R}^d} \int_{S^{d-1}} G(x - y, R_m^T n) K|_{(y, R_m)} d\sigma(m) dy, \end{aligned} \tag{27}$$

ergon with the usual surface measure σ on S^{d-1} and with the smoothing kernel

$$G(y, m) := G_{\rho_s}^{\mathbb{R}^d}(\mathbf{0}, y) \cdot G_{\rho_a}^{S^{d-1}}(a, m), \tag{28}$$

where G_{ρ}^M is the heat kernel on the Riemannian manifold M with timescale $\rho > 0$, the spatial kernel is centered on $\mathbf{0}$, and the orientation kernel is centered on the reference direction a .

Remark 9 (*Diffusion on \mathbb{M}_d*) It is important in the context of \mathbb{M}_d to choose diffusion that is isotropic spatially (with timescale ρ_s) and spherically (with timescale ρ_a) since this is the only diffusion that commutes with the left-invariant vector fields. Note that $G_{\rho_s}^{\mathbb{R}^d}(\mathbf{0}, \mathbf{y})$ depends only on $\|\mathbf{y}\|$ and $G_{\rho_a}^{S^{d-1}}(\mathbf{a}, \mathbf{m})$ depends only on $\arccos(\mathbf{a} \cdot \mathbf{m})$ making $G(\mathbf{y}, \mathbf{m})$ the heat kernel on the product manifold $\mathbb{R}^d \times S^{d-1}$. This smoothing method is a variant on the one used in [23].

The remaining basis vectors are determined by considering a rotation that maps $\mathcal{A}_d|_g$ to $\mathcal{B}_d|_g$ and then applying a specific rotation to the remaining $\mathcal{A}_i|_g$ that keeps the remaining spatial generators spatial. For an illustration see Fig. 4. How this rotation is chosen and applied is detailed in [23, App.B].

Having determined a data-adaptive frame $(\mathcal{B}_i)_{i=1}^{2d-1}$ (induced by \mathcal{B}_d , recall Fig. 4), we equip it with the following straightforward metric, where again we rely on the corresponding dual frame $(\beta^i)_{i=1}^{2d-1}$ given by

$$\langle \beta^i, \mathcal{B}_j \rangle = \delta_j^i. \tag{29}$$

Definition 4 (*Gauge metric tensor field*) We define the gauge metric tensor field $g \mapsto \mathcal{J}|_g(\cdot, \cdot)$ as

$$\mathcal{J} := \sum_{i=1}^{2d-1} \beta^i \otimes \beta^i, \tag{30}$$

which induces a norm on $\dot{g} \in T_g(\text{SE}(d))$:

$$\|\dot{g}\|_{\mathcal{J}} := \sqrt{\mathcal{J}|_g(\dot{g}, \dot{g})}, \tag{31}$$

a gradient on $\tilde{U} \in C^1(\text{SE}(d))$:

$$\nabla_{\mathcal{J}} \tilde{U} := \sum_{i=1}^{2d-1} (\mathcal{B}_i \tilde{U}) \mathcal{B}_i, \tag{32}$$

with norm

$$\|\nabla_{\mathcal{J}} \tilde{U}\|_{\mathcal{J}}^2 = \sum_{i=1}^{2d-1} (\mathcal{B}_i \tilde{U})^2, \tag{33}$$

and finally gives the divergence of a vector field as:

$$\text{div}_{\mathcal{J}} = \sum_{i=1}^{2d-1} \mathcal{B}_i \circ (g \mapsto \beta^i|_g), \tag{34}$$

which means that if we apply it to a vector field expressed in the gauge frame as $\mathbf{u} = \sum_{i=1}^{2d-1} u^i \mathcal{B}_i$ we have:

$$\text{div}_{\mathcal{J}} \mathbf{u} = \sum_{i=1}^{2d-1} \mathcal{B}_i u^i.$$

2.5 Coherence Enhancement Operator

Coherence-enhancing diffusion is a well-known technique for image enhancement [54]. It is intended for line amplification rather than strictly denoising. Crossing-preserving versions on \mathbb{M}_d have been developed [23] and evaluated for denoising. Here, crossing lines are well enhanced, but plateaus and boundaries of line structures are damaged. Therefore we propose to include the coherence enhancement technique into TV and MC flows.

Next we explain how this coherence enhancement operator is constructed from an orientation confidence.

In \mathbb{R}^3 , orientation confidence is calculated by the Laplacian in the subspace orthogonal to the line structure. We can take a similar approach in \mathbb{M}_d by taking the Laplacian in the space spanned by $(\mathcal{A}_i)_{i=1, i \neq d}^{2d-1}$. Recall that \mathcal{A}_d is implicitly aligned with the local line structure along \mathbf{n} . In the gauge-frame setting \mathcal{B}_d is explicitly aligned with the line structure (see Fig. 4), and therefore we take the Laplacian in the span of $(\mathcal{B}_i)_{i=1, i \neq d}^{2d-1}$.

In the sub-Riemannian case (i.e., $D_1 = \dots = D_{d-1} = 0$) this just reduces to the second derivatives in the $(d-1)$ -dimensional spaces spanned by $(\mathcal{A}_i)_{i=d+1}^{2d-1}$ and $(\mathcal{B}_i)_{i=d+1}^{2d-1}$, respectively. With that in mind we define orientation confidence in $\text{SE}(d)$ as follows.

Let $\tilde{U} : \text{SE}(d) \rightarrow \mathbb{R}$, then in the left-invariant case we define

$$C_{\tilde{U}, \text{LI}}(g) := \left| \sum_{i=1, i \neq d}^{2d-1} D_i \mathcal{A}_i|_g (\mathcal{A}_i \tilde{U}) \right|, \tag{35}$$

or

$$C_{\tilde{U}, \text{GF}}(g) := \left| \sum_{i=1, i \neq d}^{2d-1} \mathcal{B}_i|_g (\mathcal{B}_i \tilde{U}) \right|. \tag{36}$$

in the gauge-frame case.

Note that the \mathcal{B}_i 's are normalized with respect to old metric (12) and as such the parameters D_i are still included in (36). In the case that \mathcal{B}_d is aligned with \mathcal{A}_d we have $\mathcal{B}_i = \sqrt{D_i} \mathcal{A}_i$ and (35) and (36) coincide.

Definition 5 (*Isotropy factor*) Let $c > 0$ be a chosen scaling constant, then the isotropy factor is defined as:

$$\alpha_{\tilde{U}}(g) := e^{-\frac{C_{\tilde{U}}(g)}{c}},$$

with $C_{\tilde{U}}$ defined by (35), respectively (36).

What is convenient about this quantity is that it gives a number in the range $(0, 1]$ with a number close to zero indicating a high degree of anisotropy and a 1 indicating perfect isotropy. This is the quantity that we can use to steer flow.

The choice of c controls how steep the decline of the isotropy factor is. Its appropriate value depends on the application and on exactly how the data are represented numerically (normalized to $[0, 1]$ in our case) and are best determined heuristically or by histogram. For our experiments we have used $c = 0.2$.

Using this scalar function $\alpha_{\tilde{U}}$ on the group $SE(d)$, we can locally modify vectors based on how certain we are the data are locally aligned. We refer to this modification of vector fields as *coherence enhancement* (as in coherence enhancing diffusion [30]). Tangent vectors (such as the gradient as we will see) are modified as follows. Let \mathbf{v} be a vector field on $SE(d)$. Then the coherence-enhanced vector field is given as

$$E_G(\mathbf{v}) := \langle \omega^d, \mathbf{v} \rangle \mathcal{A}_d + \alpha_{\tilde{U}} \sum_{\substack{i=1 \\ i \neq d}}^{2d-1} \langle \omega^d, \mathbf{v} \rangle \mathcal{A}_i, \tag{37}$$

for the left-invariant geometry and as

$$E_J(\mathbf{v}) := \langle \beta^d, \mathbf{v} \rangle \mathcal{B}_d + \alpha_{\tilde{U}} \sum_{\substack{i=1 \\ i \neq d}}^{2d-1} \langle \beta^d, \mathbf{v} \rangle \mathcal{B}_i. \tag{38}$$

and for the gauge geometry.

Intuitively, these linear operators $E_G, E_J : T(G) \rightarrow T(G)$ preserve the magnitude of the vector in the main direction and weaken it orthogonal to the main direction if we are certain the data are locally aligned to the main direction.

2.6 Descending to the Homogeneous Space

So far we have developed two distinct geometries on the group $SE(d)$ that are summarized in Table 1. We can bring these geometries down to the homogeneous space \mathbb{M}_d by considering the natural extension of functions and vector fields on \mathbb{M}_d to $SE(d)$. Consider a function $U \in C^\infty(\mathbb{M}_d)$, then the function \tilde{U} , given by

$$\tilde{U}(g) = \tilde{U}(\mathbf{x}, \mathbf{R}) := U(\mathbf{x}, \mathbf{R}\mathbf{a}) \tag{39}$$

for all $g = (\mathbf{x}, \mathbf{R}) \in SE(d)$, is its natural extension to $SE(d)$ and is clearly also smooth. Similarly, a tangent vector field (recall that we understand these as differential operators acting on scalar functions) \mathbf{v} on \mathbb{M}_d can be extended as follows:

$$\left(\tilde{\mathbf{v}}\tilde{U} \right) (g) = \left(\tilde{\mathbf{v}}\tilde{U} \right) (\mathbf{x}, \mathbf{R}) := (\mathbf{v}U)(\mathbf{x}, \mathbf{R}\mathbf{a}), \tag{40}$$

under the additional constraint that $\tilde{\mathbf{v}}$ vanishes in the direction induced by the subgroup H_d (i.e., for all $i \geq 2d$ we have $\langle \rightarrow^i, \tilde{\mathbf{v}} \rangle = 0$) this extension is unique.

Having extended functions and vector fields upward to the group, we can apply the tools from Table 1 to them and subsequently project the results back to the homogeneous space by the mapping $(\mathbf{x}, \mathbf{R}) \mapsto (\mathbf{x}, \mathbf{R}\mathbf{a})$. This mapping is not injective. Nevertheless, thanks to metrics (12), (30) being laterally and spherically isotropic and the way we extend functions to the group by (39), (40), all the tools we list in Table 2 are well defined on \mathbb{M}_d .

Remark 10 (Choice of \mathbf{R}_n) While the choice of mapping $\mathbf{n} \mapsto \mathbf{R}_n$ does not matter for the final result, a choice does have to be made for an implementation when $d \geq 3$. The most straightforward manner is selecting that \mathbf{R}_n which is an in-plane rotation, meaning the plane of rotation is spanned by \mathbf{a} and \mathbf{n} . In the two cases where this is not possible (i.e., $\mathbf{n} = \pm\mathbf{a}$) we pick $\mathbf{R}_a = \mathbf{0}$ and $\mathbf{R}_{-a} = \mathbf{R}_{\mathbf{e}_z, \pi}$, where $\mathbf{R}_{\mathbf{e}_y, \pi}$ denotes the rotation around the axis \mathbf{e}_y by an angle π . Concretely the in-plane rotation in 3D is given in terms of the ZYZ-Euler angles α, β, γ by requiring that $\alpha = -\gamma$, which gives the mapping

$$\mathbf{n} \mapsto \mathbf{R}_{\mathbf{e}_z, -\alpha} \mathbf{R}_{\mathbf{e}_y, \beta} \mathbf{R}_{\mathbf{e}_z, \alpha}, \tag{41}$$

for the unique $\alpha \in [0, 2\pi)$ and $\beta \in (0, \pi)$ so that the resulting rotation maps \mathbf{a} to \mathbf{n} .

3 Total Variation, Mean Curvature PDE Flows on \mathbb{M}_d

3.1 PDE System

On \mathbb{R}^n the formulation of total variation is built on the identity $\operatorname{div}(f\mathbf{v}) = f \operatorname{div}(\mathbf{v}) + \nabla f \cdot \mathbf{v}$. Similarly on \mathbb{M}_d we have:

$$\begin{aligned} \operatorname{div}(U\mathbf{v}) &= U \operatorname{div}(\mathbf{v}) + dU(\mathbf{v}) \\ &= U \operatorname{div}(\mathbf{v}) + (\nabla U, \mathbf{v})_{\mathbb{M}_d}. \end{aligned} \tag{42}$$

From the last equation we deduce the following integration by parts formula:

$$\int_{\Omega} U(p) \operatorname{div}(\mathbf{v}(p)) d\mu(p) = \int_{\Omega} (\nabla U(p), \mathbf{v}(p))_{\mathbb{M}_d} d\mu(p), \tag{43}$$

for all $U \in C^1(\Omega)$ and all smooth vector fields \mathbf{v} vanishing at the boundary $\partial\Omega$. This formula allows us to build a weak formulation of TVF on \mathbb{M}_d starting from functions of bounded variation (BV) [1].

Definition 6 (Weak formulation of TVF on \mathbb{M}_d) Let $\epsilon \geq 0$. Let $U \in \mathbb{L}_1(\Omega)$ (recall (1)). Let $\chi_0(\Omega)$ denote the vector space of smooth vector fields that vanish at the boundary $\partial\Omega$ and let $\epsilon \geq 0$. Then we define

$$TV_\epsilon(U) := \sup_{\substack{\psi \in C_c^\infty(\Omega) \\ \mathbf{v} \in \chi_0(\Omega) \\ \|\mathbf{v}(p)\|_\epsilon^2 + |\psi(p)|^2 \leq 1}} \int_\Omega \begin{pmatrix} \epsilon \\ U(p) \end{pmatrix} \cdot \begin{pmatrix} \psi(p) \\ \operatorname{div} \mathbf{v}(p) \end{pmatrix} d\mu(p). \tag{44}$$

If $TV_0(U) < \infty$ we say that $U \in BV(\Omega)$.

For all $U \in BV(\Omega)$ we have

$$TV_0(U) \leq TV_\epsilon(U) \leq TV_0(U) + \epsilon|\Omega|.$$

Recall Remark 7 about the sub-Riemannian setting, and recall the notion of horizontal tangent vectors (13). So (44) also covers the sub-Riemannian setting (i.e., $\epsilon = 0$) when setting

$$\|\mathbf{v}(p)\|_{\epsilon=0} = \begin{cases} \mathcal{G}|_p(\mathbf{v}(p), \mathbf{v}(p)) & \text{if } \mathbf{v}(p) \text{ is horizontal,} \\ \infty & \text{else,} \end{cases} \tag{45}$$

for tangent vector $\mathbf{v}(p) = (\dot{\mathbf{x}}(p), \dot{\mathbf{n}}(p)) \in T_p(\mathbb{M}_d)$ at base point $p = (\mathbf{x}, \mathbf{n}) \in \mathbb{M}_d$.

Lemma 1 Let $\epsilon \geq 0$. For $U \in C^1(\Omega, \mathbb{R})$ we have

$$TV_\epsilon(U) = \int_\Omega \sqrt{\|\nabla U(p)\|_{\mathbb{M}_d}^2 + \epsilon^2} d\mu(p). \tag{46}$$

Furthermore for $U \in C^2(\mathbb{M}_d, \mathbb{R})$ and $\epsilon, \epsilon > 0$ we have that

$$\partial TV_\epsilon(U) = \operatorname{div} \circ \left(\frac{\nabla U}{\sqrt{\|\nabla U\|_{\mathbb{M}_d}^2 + \epsilon^2}} \right). \tag{47}$$

Proof First we substitute (43) into (44), then we apply Gauss theorem and use $U\mathbf{v}|_{\partial\Omega} = 0$. Then we apply Cauchy–Schwarz on $V_p := \mathbb{R} \times T_p\mathbb{M}_d$ for each $p \in \mathbb{M}_d$, with inner product

$$(\epsilon_1, \mathbf{v}_1) \cdot (\epsilon_2, \mathbf{v}_2) = \epsilon_1\epsilon_2 + (\mathbf{v}_1, \mathbf{v}_2)_{\mathbb{M}_d},$$

which holds with equality iff the vectors are linearly dependent. Therefore we smoothly approximate

$$\frac{1}{\sqrt{\epsilon^2 + \|\nabla U\|_{\mathbb{M}_d}^2}}(\epsilon, \nabla U)$$

by (ψ, \mathbf{v}) one obtains (46).

For $U \in C^2(\Omega, \mathbb{R})$, $\delta \in C_c^\infty(\Omega, \mathbb{R})$ we obtain

$$\begin{aligned} (\partial TV_\epsilon(U), \delta)_{\mathbb{L}_2(\Omega)} &= \lim_{h \downarrow 0} \frac{TV_\epsilon(U + h\delta) - TV_\epsilon(U)}{h} \\ &\stackrel{(43)}{=} \left(\operatorname{div} \circ \left(\frac{\nabla U}{\sqrt{\|\nabla U\|_{\mathbb{M}_d}^2 + \epsilon^2}} \right), \delta \right)_{\mathbb{L}_2(\Omega)}, \end{aligned}$$

and the result follows. \square

For vector fields \mathbf{v} on \mathbb{M}_d define the regularized norm:

$$\|\mathbf{v}\|_\epsilon := \sqrt{\|\mathbf{v}\|_{\mathbb{M}_d}^2 + \epsilon^2}. \tag{48}$$

This is a common way to regularize denominators, and we will use Sect. 4 to prove that this approach converges for $\epsilon \rightarrow 0$.

Now we propose the following roto-translation equivariant enhancement PDE on $\Omega \subset \mathbb{M}_d$, recall (1).

Definition 7 (Equivariant enhancement PDE) Given $U \in BV(\Omega)$ with $\epsilon > 0$ then we call $W^\epsilon : \mathbb{M}_d \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ (being smooth) obeying

$$\begin{cases} \frac{\partial W^\epsilon}{\partial t}(p, t) = \|\nabla W^\epsilon(p, t)\|_\epsilon^a \operatorname{div} \left(\frac{\nabla W^\epsilon(\cdot, t)}{\|\nabla W^\epsilon(\cdot, t)\|_\epsilon^b} \right)(p), \\ \nabla_{\mathbb{R}^d} W^\epsilon(p, t) \cdot \mathbf{N}(\mathbf{x}) = 0 & \text{for } p = (\mathbf{x}, \mathbf{n}) \in \partial\Omega, t \geq 0, \\ W^\epsilon(p, 0) = U(p) & \text{for } p \in \Omega, \end{cases} \tag{49}$$

the gradient flow started at U with evolution time $t \geq 0$ and parameters $a, b \in \{0, 1\}$. Here we use Neumann boundary conditions with $\mathbf{N}(\mathbf{x})$ as the normal to the spatial boundary at $\mathbf{x} \in \Omega_f$.

The coherence enhancement version of this PDE is given by replacing div by $\operatorname{div} \circ E$ (recall (37) and (38)):

$$\frac{\partial W^\epsilon}{\partial t}(p, t) = \|\nabla W^\epsilon(p, t)\|_\epsilon^a \operatorname{div} \circ E \left(\frac{\nabla W^\epsilon(\cdot, t)}{\|\nabla W^\epsilon(\cdot, t)\|_\epsilon^b} \right)(p).$$

Remark 11 (Two versions of the PDE) This PDE system on the quotient \mathbb{M}_d has two versions depending on whether one chooses the left invariant or gauge geometry as outlined in Tables 1 and 2.

We then have the following cases:

- For $(a, b) = (1, 1)$ we arrive at mean curvature flow (MCF), studied for $d = 2$ in [13].
- For $(a, b) = (0, 1)$ we arrive at total variation flow, studied for $d = 2$ in [11].
- For $(a, b) = (0, 0)$ we arrive at a linear diffusion for which exact smooth solutions exist for both $d = 2$ and $d = 3$ [43].

Remark 12 (*PDE in terms of curvature*) By product rule (42) the right-hand side of (49) for $\varepsilon \downarrow 0$ becomes

$$\frac{\partial W^0}{\partial t} = \|\nabla W^0\|^{a-b} \Delta W^0 + 2b \bar{\kappa}_I \|\nabla W^0\|^a, \tag{50}$$

with the mean curvature $\bar{\kappa}_I(p, t)$ of level set

$$\{q \in \mathbb{M}_d \mid W^0(q, t) = W^0(p, t)\},$$

akin to [49, ch:3.2], and with (possibly hypo-elliptic) Laplacian $\Delta = \text{div} \circ \nabla$.

Remark 13 (*Lack of regularity and weak solutions*) For MCF and TVF smooth solutions to PDE (49) exist only under special circumstances. This lack of regularity is an advantage in image processing to preserve step edges and plateaus in images, yet it forces us to define a concept of weak solutions. Here, we distinguish between MCF and TVF.

For MCF one relies on viscosity solution theory developed by Evans–Spruck [26], see also [32,50] for the case of MCF with Neumann boundary conditions. In [13, Thm 3.6] existence of C^1 -viscosity solutions is shown for $d = 2$.

For TVF we will rely on gradient flow theory by Brezis–Komura [2,9].

Remark 14 In this article we do not address convergence of our PDE solutions toward the sub-Riemannian setting $\varepsilon \downarrow 0$, and we only focus on convergence results for $\varepsilon \downarrow 0$. In previous work (by others) convergence to the sub-Riemannian setting is addressed for special cases. For the special case $(a, b) = (0, 0)$ convergence of the solutions with respect to $\varepsilon \downarrow 0$ is clear from the exact solutions see [43, ch:2.7]. For such convergence in the challenging case $(a, b) = (1, 1)$ (MCF), see Citti et al. [4,13]. For Eikonal PDEs convergence of viscosity solutions toward the sub-Riemannian setting holds as well, see [24, Thm.2]. It is therefore interesting to see whether convergence results toward the sub-Riemannian setting hold for the general case including the TVF case, but this falls outside the scope of this article. In Sect. 4 we only focus on convergence results for $\varepsilon \downarrow 0$ for $\varepsilon > 0$ fixed.

3.2 Numerics

We implemented PDE system (49) by Euler forward time discretization, relying on standard B-spline or linear interpolation techniques for derivatives in the underlying tools on \mathbb{M}_d given in Table 2. For details see [17,30]. Also, the explicit upperbounds for stable choices of stepsizes can be derived by the Gershgorin circle theorem [17,30].

For $d = 2$ the discretization is straightforward [30], for $d = 3$ we discretized per [17] in the software package *Lie Analysis* for Mathematica developed by Martin et al. [37], to our PDEs of interest (49) on \mathbb{M}_3 .

Table 3 Relative computational time needed by our methods relative to spatial Perona–Malik for the denoising results in Fig. 11

Method	Relative computational time
Perona–Malik 2D	1.0
Left-invariant MCF	88.8
Left-invariant TVF	60.4
Gauge MCF	149.1
Gauge TVF	123.6

The Euler-forward discretizations are not unconditionally stable. For $a = b = 0$, the Gershgorin circle theorem [17, ch.4.2] gives the stability bound

$$\Delta t \leq (\Delta t)_{\text{crit}} := \left(\frac{(d-1)D_A + D_S}{2h^2} + \frac{(d-1)D_A}{2h_a^2} \right)^{-1},$$

when using linear interpolation with spatial stepsize h and angular stepsize h_a . In our experiments, for $d = 2$ we set $h = 1$ and for $d = 3$ we took $h_a = \frac{\pi}{25}$ using an almost uniform spherical sampling from a tessellated icosahedron with $N_A = 162$ points. TVF required smaller time steps when ε decreases. Keeping in mind (50) but then applying product rule (42) to the case $0 < \varepsilon \ll 1$, we concentrate on the first term as it is of order ε^{-1} when the gradient vanishes. Then we find $\Delta t \leq \varepsilon \cdot (\Delta t)_{\text{crit}}$ for TVF. For MCF we do not have this limitation.

While both the quantitative and qualitative results of our proposed methods are encouraging we have to end our numerics section mentioning the computational cost. In Table 3 we summarize the relative computational time of our methods versus spatial Perona–Malik, and this summary shows our methods being several orders of magnitude slower. In the interest of fairness we need to add that benchmarking our prototype Mathematica implementation against the built-in Perona–Malik implementation is not a fair comparison, we are confident that an optimized native implementation would fare much better in a performance comparison.

4 Gradient Flow and Convergence

In this section we provide a gradient flow formulation that we will use to prove the convergence of our regularization scheme for TVF. The reader who is more interested in the experimental results than the technical convergence results can safely choose to skip this section and continue reading Sect. 5.

4.1 Preliminaries

The total variation flow can be seen as a gradient flow of a lower-semicontinuous, convex functional in a Hilbert space, as we explain next.

If $F : H \rightarrow [0, \infty]$ is a proper (i.e., not everywhere equal to infinity), lower semicontinuous, convex functional on a Hilbert space H (not to be confused with the subgroup H above, as we will not need the subgroup anymore we will stick with convention and use H for the Hilbert space from now on), the subdifferential of F in a point u in the finiteness domain of F is defined as

$$\partial F(u) := \{z \in H \mid (z, v - u) \leq F(v) - F(u) \text{ for all } v \in H\}.$$

The subdifferential is closed and convex, and thereby it has an element of minimal norm, called “the gradient of F in u ” denoted by $\text{grad}F(u)$. Let u_0 be in the closure of the finiteness domain of F . By Brezis–Komura theory, [9], [2, Thm 2.4.15] there is a unique locally absolutely continuous curve $u : [0, \infty) \rightarrow H$ such that

$$-u'(t) = \text{grad}F(u(t)) \text{ for a.e. } t > 0 \text{ and } \lim_{t \downarrow 0} u(t) = u_0.$$

We call $u : [0, \infty) \rightarrow H$ the gradient flow of F starting at u_0 .

Recall the definition of Ω in (1), then the function $\text{TV}_\epsilon : \mathbb{L}_2(\Omega) \rightarrow [0, \infty]$ is lower-semicontinuous and convex for every $\epsilon \geq 0$. This allows us to generalize solutions to PDE (49) as follows:

Definition 8 Let $U \in \Xi := BV(\Omega) \cap \mathbb{L}_2(\Omega)$. We define by $t \mapsto W^\epsilon(\cdot, t)$ the gradient flow of TV_ϵ starting at U .

Remark 15 (Smooth solutions) A smooth solution W^ϵ to (49) with $(a, b) = (0, 1)$ is a gradient flow.

A functional $\Phi : H \rightarrow (-\infty, \infty]$ is said to be λ -convex for some $\lambda \in \mathbb{R}$ if

$$u \mapsto \Phi(u) - \frac{\lambda}{2} \|u\|^2$$

is convex. In that case, the functional

$$u \mapsto \Phi(u) - \frac{\lambda}{2} \|u - v\|^2$$

is convex as well, for arbitrary $v \in H$, because the latter functional deviates from the first by an affine functional.

We first prove a stability estimate for the minimization of $1/\tau$ -convex functionals.

Lemma 2 Let $\tau > 0$. If a functional $\Phi : H \rightarrow (-\infty, \infty]$ on H is $1/\tau$ -convex, and u^* is its unique minimizer, then for all $u \in H$,

$$\frac{1}{2\tau} \|u - u^*\|^2 \leq \Phi(u) - \Phi(u^*).$$

This lemma is an extension of a standard result regarding strongly convex functionals (see, e.g., [41, Thm. 2.1.7]) but with no assumptions on differentiability. We include the proof in “Appendix B.”

For a proper (i.e., not everywhere equal to ∞), lower semicontinuous, convex functional F , and $\tau > 0$, define the so-called [48] proximal operator $J_\tau^F : H \rightarrow H$ by

$$J_\tau^F [u_0] := \operatorname{argmin}_{u \in H} \left(\frac{1}{2\tau} \|u - u_0\|^2 + F(u) \right).$$

Proposition 1 Let $\delta > 0$ and let $F, G : H \rightarrow [0, \infty]$ be two nonnegative, proper, lower semicontinuous, convex functionals on a Hilbert space H , such that for all $u \in H$,

$$F(u) - \delta \leq G(u) \leq F(u) + \delta. \tag{51}$$

Let $u_0, v_0 \in H$, such that

$$|\partial F|(u_0) \leq L \quad \text{and} \quad |\partial G|(v_0) \leq L. \tag{52}$$

Then, we have the following estimate for the gradient flow $u : [0, \infty) \rightarrow H$ of F starting at u_0 and the gradient flow $v : [0, \infty) \rightarrow H$ of G starting at v_0 :

$$\|u(t) - v(t)\| \leq \begin{cases} 4\sqrt{\delta t} + \|u_0 - v_0\| & \text{for } 0 \leq t \leq \frac{\delta}{L^2} \\ 8\sqrt[3]{L\delta t^2} + \|u_0 - v_0\| & \text{for } t > \frac{\delta}{L^2}. \end{cases}$$

We provide the proof in “Appendix C.” The idea is that the stability estimate in Lemma 2 will allow us to conclude that $J_\tau^F [u_0]$ and $J_\tau^G [v_0]$ are close when u_0 and v_0 are close. By iterating the operators J_τ^F and J_τ^G , we approximate the gradient flows of F and G , respectively, and from slope estimate (51) we will derive that this approximation is uniform. This will allow us to derive bounds for the gradient flows from the bounds for J_τ^F and J_τ^G .

We now know that the gradient flows of F and G are close when the slopes $|\partial F|(u_0)$ and $|\partial G|(v_0)$ are bounded. This assumption can be rather stringent. We will relax it and merely require that $F(u_0)$ and $G(v_0)$ are bounded by some constant $E > 0$, in exchange for a bound between gradient flows that is slightly worse. Our approach will be to run the gradient flow for a small time s from u_0 and v_0 and use the regularizing property of the gradient flow to conclude a slope bound. On the other hand, if s is small, $u(s)$ and $v(s)$ will be close to u_0 and v_0 . We will then choose s (almost) optimally to derive a bound between the gradient flows.

4.2 Strong \mathbb{L}_2 -Convergence of TV Flows

We prove the convergence, stability, and accuracy of TV flows by considering them as the gradient flows of the family of functionals TV_ε . The theory of contraction semigroups [2, Ch. 4] will allow us to show that as $\varepsilon \rightarrow 0$ the gradient flow of TV_ε converges to the gradient flow of TV_0 in the \mathbb{L}_2 sense.

Just like Lemma 1 and Proposition 1 the following theorem was already announced in [25] but similarly lacked a proof, which we now include.

Theorem 1 *Let $F : H \rightarrow [0, \infty]$ and $G : H \rightarrow [0, \infty]$ be two proper, lower semicontinuous, convex functionals on a Hilbert space H , such that*

$$F(u) - \delta \leq G(u) \leq F(u) + \delta$$

for all $u \in H$. Let $u_0, v_0 \in H$ be such that $F(u_0) \leq E$ and $G(v_0) \leq E$ and $\|u_0 - u^*\| \leq M$ and $\|v_0 - v^*\| \leq M$, for some constants $E, M > 0$, where u^* and v^* minimize F and G , respectively.

Then the gradient flow $u : [0, \infty) \rightarrow H$ of F starting at u_0 and the gradient flow $v : [0, \infty) \rightarrow H$ of G starting at v_0 satisfy

$$\|u(t) - v(t)\| \leq 16(ME\delta t^2)^{1/5} + \|u_0 - v_0\|$$

for all $0 \leq t < E^6 M^6 / \delta^9$.

Proof By the evolution variational inequality [2, Theorem 4.0.4, (iii)], we know that for all $s > 0$

$$\|u(s) - u_0\| \leq \sqrt{2sF(u_0)} \tag{53a}$$

and

$$\|v(s) - v_0\| \leq \sqrt{2sG(v_0)}. \tag{53b}$$

By the regularizing property [2, Theorem 4.0.4, (ii)],

$$|\partial F|(u(s)) \leq \frac{1}{s} \|u_0 - u^*\| \leq \frac{M}{s} \tag{54a}$$

and

$$|\partial G|(v(s)) \leq \frac{1}{s} \|v_0 - v^*\| \leq \frac{M}{s} \tag{54b}$$

where u^* minimizes F and v^* minimizes G .

Because the gradient flow is a nonexpansive semigroup [2, Theorem 4.0.4, (iv)], we obtain

$$\|u(t) - v(t)\| \leq \|u(t+s) - v(t+s)\| + \|u(t+s) - u(t)\| + \|v(t+s) - v(t)\|$$




Fig. 6 Isotropic ($\epsilon = 1$) evolution of relative \mathbb{L}_1 and \mathbb{L}_2 errors over time of the collagen image with Gaussian noise benchmarked against spatial Perona–Malik. Remaining parameters are $D_S = 1$, $D_A = .01$ and $\varepsilon = .001$

$$\leq \|u(t+s) - v(t+s)\| + \|u(s) - u_0\| + \|v(s) - v_0\|.$$

Now assume $t < E^6 M^6 / \delta^9$. We will want to choose s (almost) optimally, depending on t . We choose

$$s = \frac{M^{2/5} \delta^{2/5} t^{4/5}}{E^{3/5}}$$

and note that with $L := M/s$, we have

$$t \geq \frac{\delta}{L^2}.$$

By the slope estimates (54) we can apply Proposition 1 to the gradient flows starting at $u(s)$ and $v(s)$, to obtain

$$\begin{aligned} \|u(t) - v(t)\| &\leq 8M^{1/3} s^{-1/3} \delta^{1/3} t^{2/3} + \|u(s) - v(s)\| \\ &\quad + \|u(s) - u_0\| + \|v(s) - v_0\| \\ &\leq 8M^{1/3} s^{-1/3} \delta^{1/3} t^{2/3} + 2\|u(s) - u_0\| \\ &\quad + 2\|v(s) - v_0\| + \|u_0 - v_0\| \end{aligned}$$

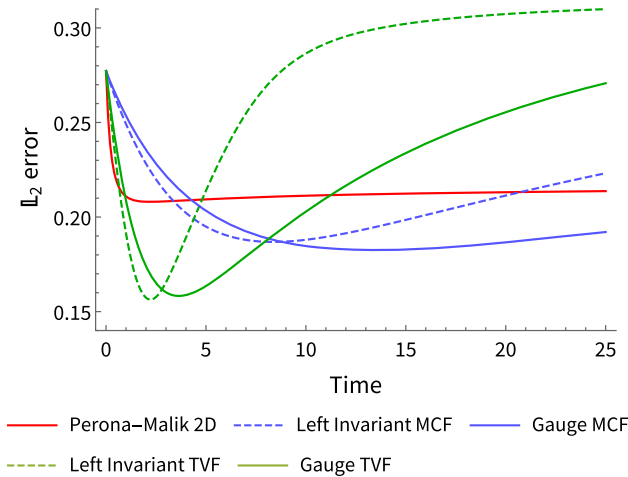



Fig. 7 Anisotropic ($\epsilon = .25$) evolution of relative \mathbb{L}_1 and \mathbb{L}_2 errors over time of the collagen image with Gaussian noise benchmarked against spatial Perona–Malik. Run with parameters $D_S = 1$, $D_A = .01$ and $\epsilon = .001$. The resulting errors are slightly worse than the isotropic setup from Fig. 6

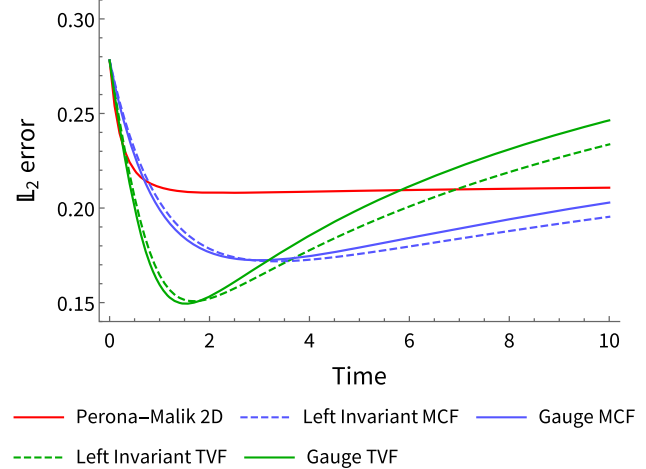



Fig. 8 Isotropic with coherence enhancement evolution of relative \mathbb{L}_1 and \mathbb{L}_2 errors over time of the collagen image with Gaussian noise benchmarked against spatial Perona–Malik. Run with parameters $D_S = 1$, $D_A = .01$, $\epsilon = .001$ and $c = .2$

$$\begin{aligned} &\stackrel{(53)}{\leq} 8M^{1/3}s^{-1/3}\delta^{1/3}t^{2/3} + \sqrt{32sE} + \|u_0 - v_0\| \\ &= 16M^{1/5}\delta^{1/5}t^{2/5}E^{1/5} + \|u_0 - v_0\|. \end{aligned}$$

□

If, for the general result of Theorem 1, we take $F = TV_0$, $G = TV_\epsilon$ and $\delta = \epsilon |\Omega|$ we obtain the following result.

Corollary 1 (Strong \mathbb{L}_2 -convergence, stability and accuracy of TV flows) *Let $U \in \mathbb{L}_2(\Omega)$ and let W^ϵ be the gradient flow of TV_ϵ starting at U and $\epsilon, \epsilon \geq 0$. Let $t \geq 0$. Let $\delta = \epsilon |\Omega|$. Then*

$$\lim_{\epsilon \downarrow 0} W^\epsilon(\cdot, t) = W^0(\cdot, t) \quad \text{in } \mathbb{L}_2(\Omega).$$

More precisely, for all $U \in BV(\Omega)$, $E \geq TV_0(U) + \delta$, $M \geq \|U\|$ and $0 \leq t < E^6 M^6 / \delta^9$ we have that

$$\|W^\epsilon(\cdot, t) - W^0(\cdot, t)\|_{\mathbb{L}_2(\Omega)} \leq 16(ME\delta t^2)^{\frac{1}{5}}.$$

5 Experiments

In our experiments, we aim to enhance contour and fiber trajectories in medical images and to remove noise. Lifting the image $f : \mathbb{R}^d \rightarrow \mathbb{R}$ toward its orientation lift $U : \mathbb{M}_d \rightarrow \mathbb{R}$ defined on the space of positions and orientations $\mathbb{M} = \mathbb{R}^d \times S^{d-1}$ preserves crossings [30] and avoids leakage of wavefronts [24].

For our experiments for $d = 3$ the initial condition $U : \mathbb{M}_3 \rightarrow \mathbb{R}^+$ is a fiber orientation density function (FODF) obtained from DW-MRI data [44].

For our experiments for $d = 2$ the initial condition U is an invertible orientation score (OS) that we sampled on 8 equidistant orientations.

For both $d = 2$ (Sect. 5.1) and $d = 3$ (Sect. 5.4), we show advantages of TVF and MCF over crossing-preserving diffusion flows [17,30] on \mathbb{M}_d .

Finally, we include denoising experiments where we show qualitative and quantitative results where comparison with

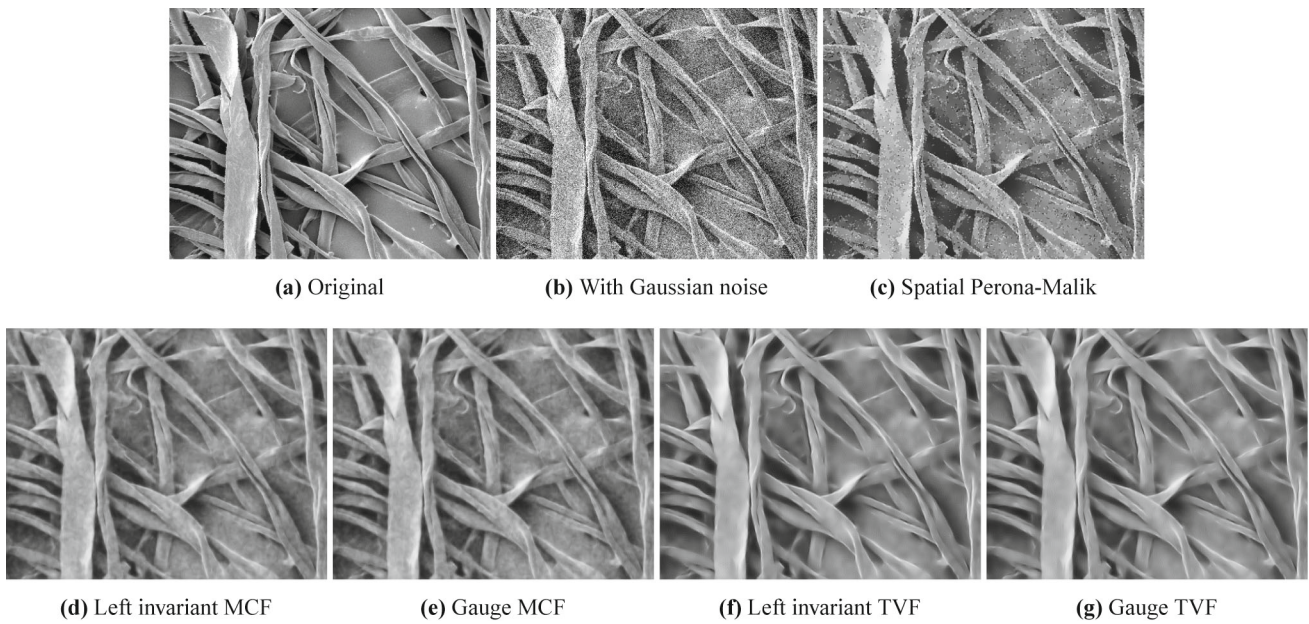


Fig. 9 Qualitative comparison of over-smoothed collagen images starting from a noisy image with $\sigma = .2$ Gaussian noise. Images are smoothed for twice the time needed to reach their minimal \mathbb{L}_2 error per Fig. 8. We use isotropic processing with coherence enhancement

the well-known denoising technique BM3D [18] shows advantages and good results.

5.1 Image Enhancement/Denoising

In accordance with the workflow in Fig. 1 we go through the following steps:

$$\begin{aligned}
 f &\mapsto \mathcal{W}_\psi f \mapsto \Phi_t \circ \mathcal{W}_\psi f \mapsto f_t^a := \mathcal{W}_\psi^* \circ \Phi_t \circ \mathcal{W}_\psi f \\
 &\approx \int_{S^1} \Phi_t(\mathcal{W}_\psi f)(\cdot, \mathbf{n}) d\mu_{S^1}(\mathbf{n}). \tag{55}
 \end{aligned}$$

for $t \geq 0$. With respect to the final step we recall that we use cake wavelets that allow for sharp approximate reconstruction by integration over angles only. Here $U \mapsto W(\cdot, t) = \Phi_t(U)$ denotes the flow operator on \mathbb{M}_2 (49). Hence the initial condition for our TVF/MCF-PDE (49) is set by an orientation score of image $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by (2).

By the invertibility of the orientation score one has $f = f_0^a$, so all flows depart from the original image.

We refer to the different methods we experimented with, by the following terms.

- MCF: we set $(a, b) = (1, 1)$.
- TVF: we set $(a, b) = (0, 1)$.
- Left invariant: we use the left-invariant geometry per the first column of Table 1.
- Gauge: we use the locally adaptive-frames geometry per the second column of Table 1.
- Isotropic: we set $\epsilon = 1$.




Fig. 10 Isotropic ($\epsilon = 1$) evolution of relative \mathbb{L}_1 and \mathbb{L}_2 errors over time of the collagen image with correlated noise ($\sigma = 0.2$ Gaussian noise with $r = 2$ and $\sigma = 1$ Gaussian filter) benchmarked against spatial Perona–Malik. Remaining parameters are $D_S = 1$, $D_A = .01$ and $\epsilon = .001$




Fig. 11 Isotropic with coherence enhancement evolution of relative \mathbb{L}_1 and \mathbb{L}_2 errors over time of the collagen image with correlated noise ($\sigma = 0.2$ Gaussian noise with $r = 2$ Gaussian filter) benchmarked against spatial Perona–Malik. Run with parameters $D_S = 1$, $D_A = .01$, $\varepsilon = .001$ and $c = .2$

- Anisotropic: we set $\varepsilon = 0.25$.
- With coherence enhancement: we use the PDE with the E operator per (49).

So that *Isotropic Gauge TVF with coherence enhancement* for example would equate to setting $(a, b) = (0, 1)$, $\varepsilon = 1$, includes the E operator in the PDE and uses the second column of Table 1 to define our geometric objects.

For quantitative comparison we will look at relative \mathbb{L}_1 and \mathbb{L}_2 errors, meaning if we have a (clean) source image f_{source} and a denoised image f_t that has been processed to time t we calculate the relative error as:

$$\frac{\|f_t - f_{source}\|}{\|f_{source}\|}, \tag{56}$$

with the corresponding \mathbb{L}_1 or \mathbb{L}_2 norm.

We will test against two types of noise: Gaussian and correlated.

5.2 Gaussian Noise

We apply Gaussian noise with standard deviation 0.2 to our normalized (to $[0, 1]$) source image; the original and noisy images are shown in Fig. 9a, respectively, Fig. 9b.

In Fig. 6 we show how the errors progress with $t \geq 0$ for the isotropic ($\varepsilon = 1$) case without coherence enhancement (i.e., without E). For comparison we plot the same error with spatial Perona–Malik. While Perona–Malik is clearly more stable and resilient to over-smoothing, both our proposed methods have much smaller minimal errors.

Remark 16 (*Interpretation of timescales*) The different methods work on different timescales, we scale these to be able to plot the results together, but no meaning should be attributed to one method obtaining a minimum earlier than another. The error graphs just show:

- how large the minimal error is and
- how fast the image deteriorates after this minimum has been reached.

In our next experiment we increase anisotropy by setting $\varepsilon = .25$, the resulting errors are plotted in Fig. 7. We gain no improvement in minimal error while requiring more computational cycles to reach the minimum, from which we conclude that for this application isotropic processing is more desirable.

In Fig. 8 we show the errors for the isotropic setup with coherence enhancement included (for $c = .2$). We get a very minor improvement in minimal errors and a decent improvement in over-smoothing stability, although still not on the level of Perona–Malik. It is remarkable that with coherence enhancement included the data-adaptive geometry is less stable than the left-invariant geometry, we observe that combining two different methods of adapting to the data is counter-productive in this instance.

For a qualitative comparison of the different isotropic methods with coherence enhancement we over-smooth the collagen image past the time of its lowest \mathbb{L}_2 error with a factor of two, the corresponding qualitative results are shown in Fig. 9.

5.3 Correlated Noise

For correlated noise we apply a Gaussian filter with $\sigma = 1.0$ to Gaussian noise with $\sigma = 0.2$. The error evolution for the isotropic methods is plotted in Fig. 10. We observe that MCF performs worse in this setting against correlated noise; both in minimal error and in stability it does not do as well as spatial Perona–Malik. TVF on the other hand has a better minimal error than Perona–Malik at the cost of stability. The

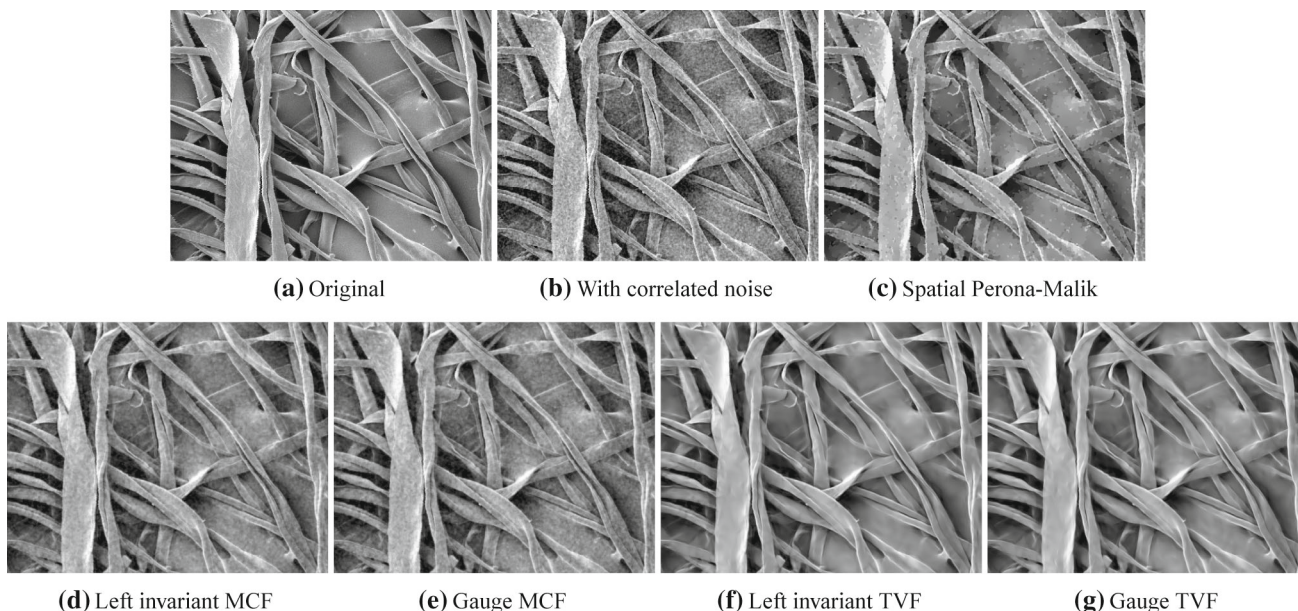


Fig. 12 Qualitative comparison of over-smoothed collagen images starting from a noisy image with correlated noise ($\sigma = 0.2$ Gaussian noise with $r = 2$ Gaussian filter). Images are smoothed for twice the

time needed to reach their minimal \mathbb{L}_2 error per Fig. 11. We use isotropic processing with coherence enhancement




Fig. 13 Comparison of left-invariant TVF with coherence enhancement against BM3D. The BM3D image is denoised at $\sigma = 0.12$ ($1.5 \times$ optimal \mathbb{L}_2) and the left-invariant TVF denoised for $1.5 \times$ optimal \mathbb{L}_2 time. Pay particular attention to the diagonal edges where TVF exhibits superior performance



Fig. 14 A spiral test image and a monochrome Mona Lisa that were used for the PSNR experiments in Table 4

stability does somewhat improve if we turn on the locally adaptive frames.

The error evolution of the experiment including the use of coherence enhancement is displayed in Fig. 11. Overall this improves the results, but MCF still exhibits an inferior performance in this setting. TVF sees both an improvement in minimal error and stability. As with the Gaussian noise experiment we see that turning on both coherence enhancement and locally adaptive frames is counterproductive.

A qualitative comparison of the methods against correlated noise is shown in Fig. 12, where again we smooth twice the optimal time. We observe the same general trend as in Fig. 9; all methods do a good job of preserving edges, but TVF stands out in clearing the plateaus.

As a final method to compare against we look at BM3D [18]. Both BM3D and our methods share a dependence on some prior knowledge for optimal performance: BM3D

Table 4 Comparing peak signal-to-noise ratio (dB) for the coherence-enhanced methods against BM3D and spatial Perona–Malik at the minimal \mathbb{L}_2 error

	Collagen	Spiral	Mona Lisa
<i>Gaussian noise</i>			
Noisy image	14.1	15.5	14.1
Perona–Malik	20.1	19.5	20.5
BM3D	23.1	21.4	23.9
Left inv. MCF	21.7	19.5	23.3
Gauge MCF	21.7	19.8	23.7
Left inv. TVF	22.4	20.9	26.0
Gauge TVF	23.0	21.2	26.1
<i>Correlated noise</i>			
Noisy image	23.9	25.3	23.9
Perona–Malik	24.2	25.3	25.1
BM3D	24.0	21.4	26.3
Left inv. MCF	23.8	24.3	26.2
Gauge MCF	23.9	24.2	26.2
Left inv. TVF	24.7	25.7	26.8
Gauge TVF	24.9	25.9	26.9

Bold values indicate the best result per image/noise combination. Three images were tested, the first being the collagen image we have been using, second a spiral test image and third a grayscale Mona Lisa. Note that all these measurements depend on some ground-truth knowledge, knowing the standard deviation of the noise in the case of BM3D and knowing the optimal processing time in case of the others.

requires us to know the standard deviation of the noise, and our method requires us to know the optimal processing time.

We make a qualitative comparison of removing correlated noise between BM3D and left-invariant TVF by smoothing $1.5\times$ past the optimal \mathbb{L}_2 error: $1.5\times$ the optimal standard deviation in case of BM3D and $1.5\times$ the processing time in case of TVF. The resulting images are shown in Fig. 13.

For a broader comparison, we compute peak signal-to-noise ratios for the collagen image we already saw and for two additional images of different styles shown in Fig. 14. Results are summarized in Table 4.

As the gauge TVF method compares favorably against correlated noise according to the PSNR value in Table 4, we will look at its qualitative result and compare it against BM3D in Fig. 15.

5.4 Denoising and Fiber Enhancement on FODFs in DW-MRI

In DW-MRI image processing one obtains a field of angular diffusivity profiles (orientation density function) of water molecules. A high diffusivity in particular orientation correlates to biological fibers structure, in brain white matter, along that same direction. Crossing-preserving enhancement of FODF fields $U : \mathbb{M}_3 \rightarrow \mathbb{R}^+$ helps to better identify struc-

tural pathways in brain white matter, which is relevant for surgery planning, see for example [38,44].

For a quantitative comparison we applied TVF, MCF, and PM diffusion [17] to denoise a popular synthetic FODF $U : \mathbb{M}_3 \rightarrow \mathbb{R}^+$ from the Fiberfox Tractometer challenge with realistic noise profiles [40]. In Fig. 16, we can observe the many crossing fibers in the dataset. Furthermore, we depicted the absolute \mathbb{L}_2 -error $t \mapsto \|U - \Phi_t(U)\|_{\mathbb{L}_2(\mathbb{M}_3)}$ as a function of the evolution parameter t , where $\Phi_t(U) = W_\varepsilon(\cdot, t)$ with optimized $\varepsilon = 0.02$ in the case of TVF (in green), and MCF (in blue), and where Φ_t is the PM diffusion evolution [17] on \mathbb{M}_3 with optimized PM parameter $K = 0.2$ (in red). We also depict results for $K = 0.1, 0.4$ (with the dashed lines) and $\varepsilon = 0.01, 0.04$. We see that the other parameter settings provide on average worse results, justifying our optimized parameter settings. We set $D_S = 1.0$, $D_A = 0.001$, $\Delta t = 0.01$. We observe that:

- TVF can reach lower error values than MC flow with adequate $\Delta t = 0.01$,
- MCF provides more stable errors for all $t > 0$ than TV flow with respect to $\varepsilon > 0$,
- TVF and MCF produce lower error values than PM diffusion,
- PM diffusion provides the most variable results for all $t > 0$.

For a qualitative comparison we applied TVF, MCF, PM diffusion and linear diffusion to a FODF field $U : \mathbb{M}_3 \rightarrow \mathbb{R}^+$ obtained from a standard DW-MRI dataset (with $b = 1000$ s/mm², 54 gradient directions) via constrained spherical deconvolution (CSD) [19,52]. See Fig. 17, where for each method, we used the optimal parameter settings with the artificial dataset. We see that

- all methods perform well on the real datasets. Contextual alignment of the angular profiles better reflects the anatomical fiber bundles,
- MCF and TVF better preserve boundaries and angular sharpness,
- MCF better preserves the amplitude at crossings at longer times.

6 Conclusion

We have proposed a PDE system on the homogeneous space $\mathbb{M}_d = \mathbb{R}^d \times S^{d-1}$ of positions and orientations, for crossing-preserving denoising and enhancement of (lifted) images containing both complex-elongated structures and plateaus.

It includes TVF, MCF and diffusion flows as special cases and includes (sub-)Riemannian geometry. Thereby we generalized recent related works by Citti et al. [13] and Chambolle




Fig. 15 Comparing gauge TVF with coherence enhancement and BM3D against correlated noise, the standard deviation for BM3D and the evolution time for TVF were set at 1.5 times the number needed to reach the optimal \mathbb{L}_2 error

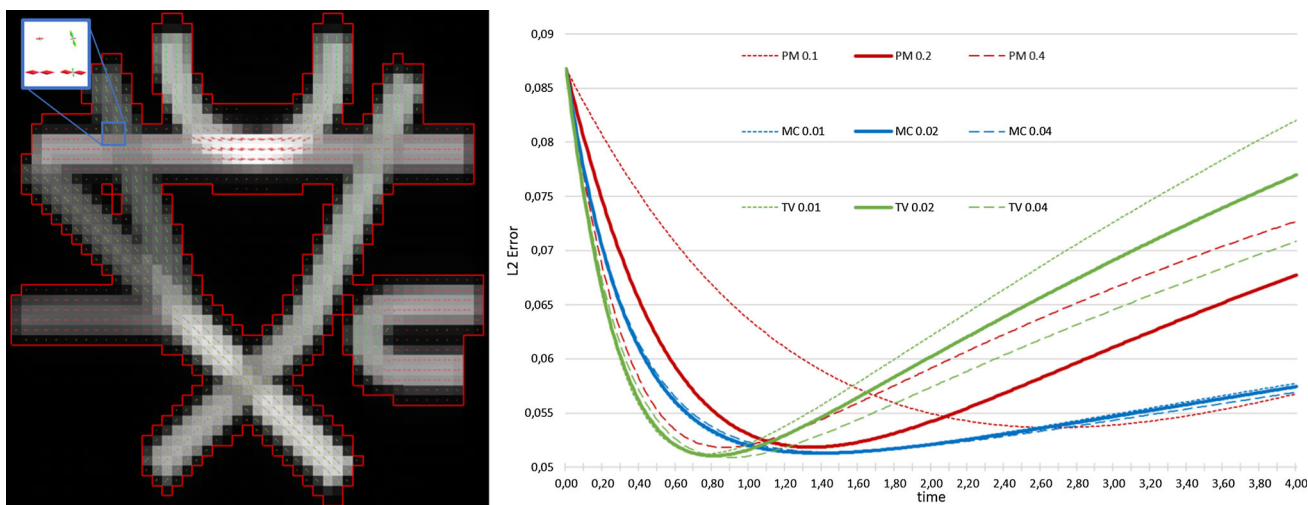


Fig. 16 Quantitative comparison of denoising a fiber orientation density function (FODF) obtained by (CSD) [19,52] from the benchmark DW-MRI dataset Fiberfox [40]

and Pock [11] from 2D to 3D using a different numerical scheme with new convergence results (Theorem 1) and stability bounds. We used the divergence and intrinsic gradient on a (sub)-Riemannian manifold above \mathbb{M}_d for a formal weak formulation of total variation flows, which simplifies if the lifted images are differentiable (Lemma 1).

For 2D image denoising and enhancement we have shown that in all cases TVF on \mathbb{M}_2 has a better minimal error than Perona–Malik and MCF at the cost of being more sensitive to

oversmoothing, recall Figs. 6, 7, 8, 9, 10, 11, and 12. The \mathbb{L}_1 , \mathbb{L}_2 and PSNR measures indicate the potential of our proposed methods for denoising, and we manage to improve PSNR results against methods such as BM3D against correlated noise on some images, recall Figs. 13 and 15. Qualitatively this is mainly reflected in better clearing of plateaus while still preserving hard edges and crossings.

In 3D we compared to previous nonlinear crossing-preserving diffusion methods on \mathbb{M}_3 ; we showed improve-

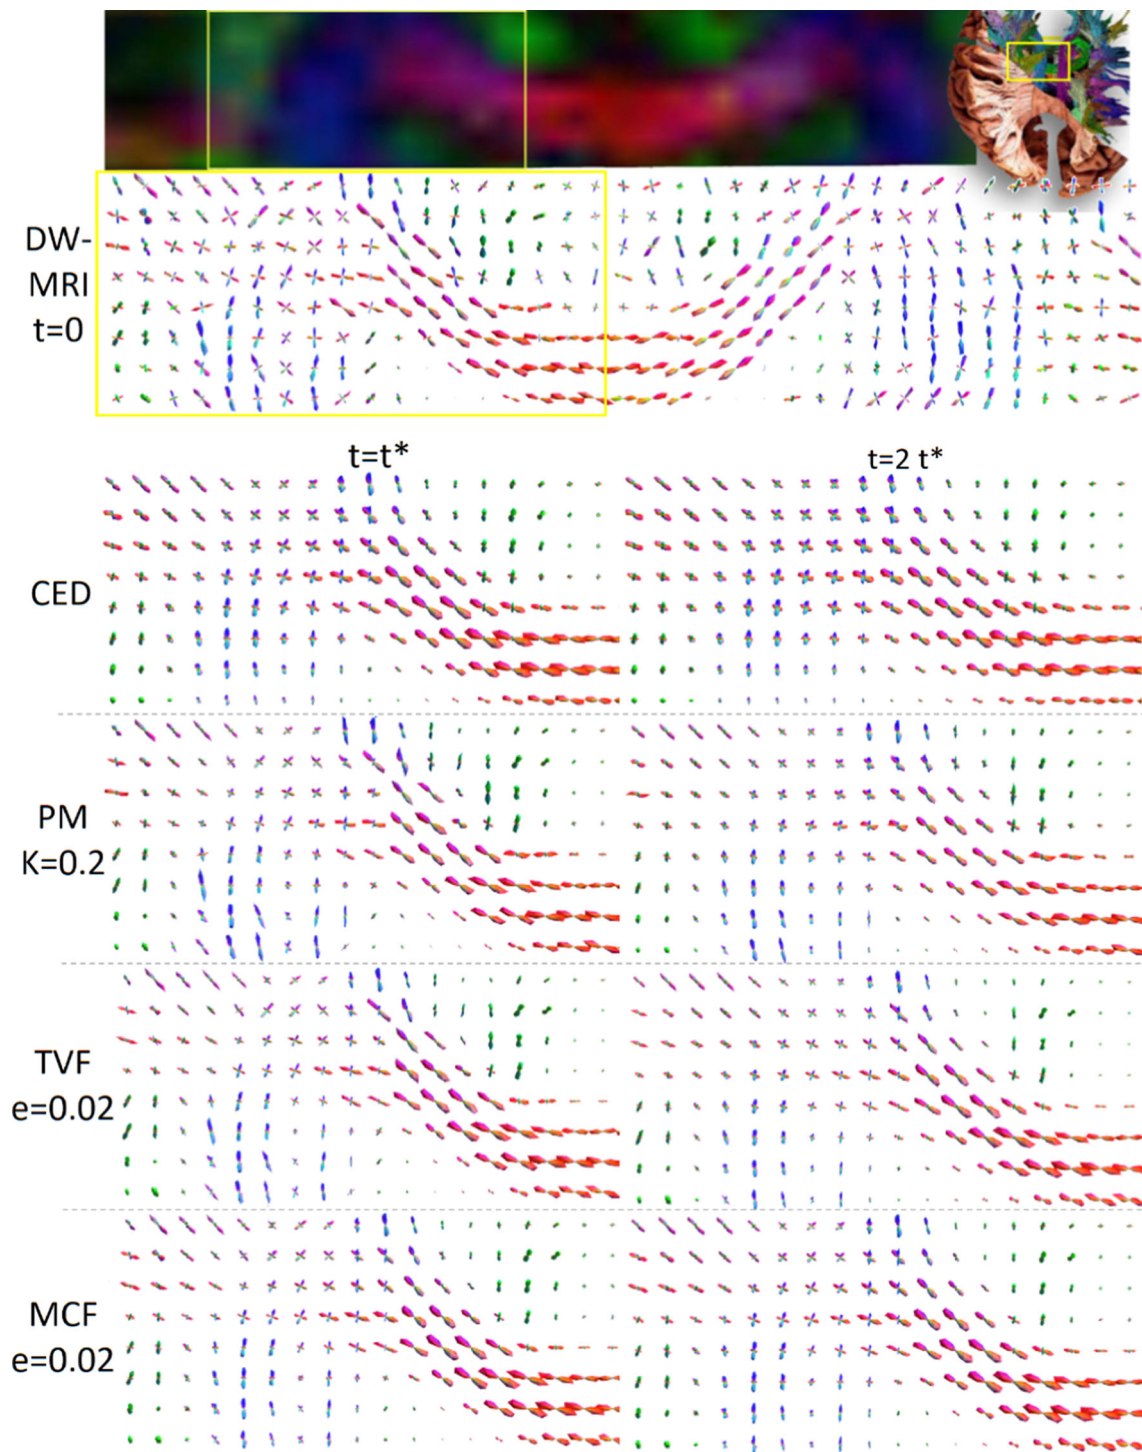


Fig. 17 Qualitative comparison of denoising a FODF obtained by (CSD) [19,52] from a standard DW-MRI dataset (with $b = 1000 \text{ s/mm}^2$ and 54 gradient directions). For the CSD we used up to 8th order

spherical harmonics, and the FODF is then spherically sampled on a tessellation of the icosahedron with 162 orientations

ments over Perona–Malik and improvements over contextual fiber enhancement methods in DW-MRI processing [17,21] on real medical image data. We observe that crossings and boundaries (of bundles and plateaus) are better preserved over time. We support this quantitatively by a denoising experiment on a benchmark DW-MRI dataset, where MCF performs better than TVF and both perform better than Perona–Malik diffusions, in view of error reduction and stability.

Altogether, we conclude that our TVF and MCF methods on \mathbb{M}_d work well for denoising and enhancement for both $d = 2$ and $d = 3$. In general we see clear benefits of the inclusion of locally adaptive frames and of limited inclusion of coherence enhancement. The code from our experiments is available as a Mathematica notebook at https://bmnsnets.com/files/tvf_mcf_denoising_jmiv.nb.

Future work While we have shown the potential of our PDE system on \mathbb{M}_d as a denoising/enhancement method some challenges remain for future work:

- Determining stopping time, our methods show good minimal errors but are prone to degrading the image if left running for too long. For general applications a robust automatic stopping method would be helpful. Spectral analysis of nonlinear operators [10,16] may apply here.
- Coherence enhancement [54] was not originally conceived for denoising. It is therefore interesting to see how edge enhancing diffusion [27] (EED) performs when generalized to \mathbb{M}_d , i.e., we would reformulate our enhancement operator E as:

$$E := \int_{S^d} \mathbf{c} \otimes \mathbf{c} e^{-\frac{|\nabla U \cdot \mathbf{c}|^2}{2\kappa^2}} d\mu(\mathbf{c}),$$

and test its performance.

- In this article we obtained convergence results of our PDE solutions for $\epsilon \downarrow 0$ while keeping $\epsilon > 0$ fixed. It is interesting to study the full limiting case $(\epsilon, \epsilon) \rightarrow (0, 0)$, for the general setting covering total variation flow.

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

Appendix A Left-Invariant Vector Field Basis for SE(3)

Two charts are required to cover SO(3), when using ZYZ-Euler angles:

$$g = (x, y, z, R_{e_z, \gamma} R_{e_y, \beta} R_{e_z, \alpha}), \tag{57}$$

where $\beta \in (0, \pi)$ and $\alpha, \gamma \in [0, 2\pi)$, the left-invariant vector field basis is given by:

$$\begin{aligned} \mathcal{A}_1|_g &= (\cos \alpha \cos \beta \cos \gamma - \sin \alpha \sin \gamma) \partial_x \\ &\quad + (\sin \alpha \cos \gamma + \cos \alpha \cos \beta \sin \gamma) \partial_y \\ &\quad - \cos \alpha \sin \beta \partial_z, \\ \mathcal{A}_2|_g &= (-\sin \alpha \cos \beta \cos \gamma - \cos \alpha \sin \gamma) \partial_x \\ &\quad + (\cos \alpha \cos \gamma - \sin \alpha \cos \beta \sin \gamma) \partial_y \\ &\quad + \sin \alpha \sin \beta \partial_z, \\ \mathcal{A}_3|_g &= \sin \beta \cos \gamma \partial_x + \sin \beta \sin \gamma \partial_y + \cos \beta \partial_z, \\ -\mathcal{A}_5|_g &= \cos \alpha \cot \beta \partial_\alpha + \sin \alpha \partial_\beta - \frac{\cos \alpha}{\sin \beta} \partial_\gamma, \\ \mathcal{A}_4|_g &= -\sin \alpha \cot \beta \partial_\alpha + \cos \alpha \partial_\beta + \frac{\sin \alpha}{\sin \beta} \partial_\gamma, \\ \mathcal{A}_6|_g &= \partial_\alpha. \end{aligned}$$

Recall that in this article we take the convention that A_{d+i} is the in-planar rotation from axis $\mathbf{a} = (0, 0, 1)$ to A_i for $i = 1, \dots, d - 1$. Thereby $-\mathcal{A}_5$ is the counterclockwise rotation about axis \mathcal{A}_1 , whereas \mathcal{A}_4 is the counterclockwise rotation about axis \mathcal{A}_2 . Recall Fig. 3.

The above set of expressions is not valid for $\beta = 0$ or $\beta = \pi$; in that case, we can switch to XYZ-Euler angles:

$$g = (x, y, z, R_{e_x, \gamma'} R_{e_y, \beta'} R_{e_z, \alpha'}), \tag{58}$$

with $\alpha' \in [0, 2\pi)$, $\beta' \in [-\pi, \pi)$ and $\gamma' \in (-\pi/2, \pi/2)$. The basis vector fields are then given by:

$$\begin{aligned} \mathcal{A}_1|_g &= \cos \alpha' \cos \beta' \partial_x \\ &\quad + (\cos \gamma' \sin \alpha' + \cos \alpha' \sin \beta' \sin \gamma') \partial_y \\ &\quad + (\sin \alpha' \sin \gamma' - \cos \alpha' \sin \beta' \cos \gamma') \partial_z, \\ \mathcal{A}_2|_g &= -\sin \alpha' \cos \beta' \partial_x \\ &\quad + (\cos \alpha' \cos \gamma' - \sin \alpha' \sin \beta' \sin \gamma') \partial_y \\ &\quad + (\sin \alpha' \sin \beta' \cos \gamma' + \cos \alpha' \sin \gamma') \partial_z, \\ \mathcal{A}_3|_g &= \sin \beta' \partial_x - \cos \beta' \sin \gamma' \partial_y + \cos \beta' \cos \gamma' \partial_z, \\ -\mathcal{A}_5|_g &= -\cos \alpha' \tan \beta' \partial_{\alpha'} + \sin \alpha' \partial_{\beta'} + \frac{\cos \alpha'}{\cos \beta'} \partial_{\gamma'}, \\ \mathcal{A}_4|_g &= \sin \alpha' \tan \beta' \partial_{\alpha'} + \cos \alpha' \partial_{\beta'} - \frac{\sin \alpha'}{\cos \beta'} \partial_{\gamma'}, \end{aligned}$$

$$\mathcal{A}_6|_g = \partial_{\alpha'},$$

for $|\beta'| \neq \pi/2$.

Appendix B Proof of Lemma 2

We recall the lemma and give a proof.

Lemma 2 *Let $\tau > 0$. If a functional $\Phi : H \rightarrow (-\infty, \infty]$ on H is $1/\tau$ -convex, and u^* is its unique minimizer, then for all $u \in H$,*

$$\frac{1}{2\tau} \|u - u^*\|^2 \leq \Phi(u) - \Phi(u^*).$$

Proof The functional $\Psi : H \rightarrow (-\infty, \infty]$ given by

$$\Psi(u) := \Phi(u) - \Phi(u^*) - \frac{1}{2\tau} \|u - u^*\|^2$$

is convex. It is sufficient to show that Ψ is nonnegative. If it were not, there would exist a $v \in H$ such that $\Psi(v) < 0$. We will show that then, for t small enough, $\Phi(tv + (1 - t)u^*) < \Phi(u^*)$, contradicting that u^* is a minimizer. We first have by definition that, for $t \in (0, 1)$,

$$\begin{aligned} \Phi(tv + (1 - t)u^*) - \Phi(u^*) - \frac{t^2}{2\tau} \|v - u^*\|^2 \\ = \Psi(tv + (1 - t)u^*). \end{aligned}$$

By the convexity of Ψ ,

$$\begin{aligned} \Psi(tv + (1 - t)u^*) &\leq t\Psi(v) + (1 - t)\Psi(u^*) \\ &= t\Psi(v). \end{aligned}$$

Combining the two inequalities, we find

$$\Phi(tv + (1 - t)u^*) - \Phi(u^*) \leq t\Psi(v) + O(t^2),$$

so that indeed, for t small enough, $\Phi(tv + (1 - t)u^*) < \Phi(u^*)$, leading to the announced contradiction.

Therefore, Ψ is nonnegative, which means that

$$\frac{1}{2\tau} \|u - u^*\|^2 \leq \Phi(u) - \Phi(u^*)$$

for all $u \in H$. □

Appendix C Proof of Proposition 1

We recall the proposition and give a proof.

Proposition 1 *Let $F, G : H \rightarrow [0, \infty]$ be two nonnegative, proper, lower semicontinuous, convex functionals on a Hilbert space H , such that for all $u \in H$,*

$$F(u) - \delta \leq G(u) \leq F(u) + \delta. \tag{59}$$

Let $u_0, v_0 \in H$, such that

$$|\partial F|(u_0) \leq L \quad \text{and} \quad |\partial G|(v_0) \leq L. \tag{60}$$

Then, we have the following estimate for the gradient flow $u : [0, \infty) \rightarrow H$ of F starting at u_0 and the gradient flow $v : [0, \infty) \rightarrow H$ of G starting at v_0 :

$$\|u(t) - v(t)\| \leq \begin{cases} 4\sqrt{\delta t} + \|u_0 - v_0\| & \text{for } 0 \leq t \leq \frac{\delta}{L^2} \\ 8\sqrt[3]{L\delta t^2} + \|u_0 - v_0\| & \text{for } t > \frac{\delta}{L^2}. \end{cases}$$

Proof Let $\tau > 0$ and let $u_1^F := J_\tau^F[u_0]$ and $v_1^G := J_\tau^G[v_0]$. Set also $v_1^F := J_\tau^F[v_0]$ and $u_1^G := J_\tau^G[u_0]$. Then, using the definition of v_1^F in the second inequality below, we find

$$\begin{aligned} \frac{1}{2\tau} \|v_1^F - v_0\|^2 + G(v_1^F) \\ \stackrel{(59)}{\leq} \frac{1}{2\tau} \|v_1^F - v_0\|^2 + F(v_1^F) + \delta \\ \leq \frac{1}{2\tau} \|v_1^G - v_0\|^2 + F(v_1^G) + \delta \\ \stackrel{(59)}{\leq} \frac{1}{2\tau} \|v_1^G - v_0\|^2 + G(v_1^G) + 2\delta. \end{aligned}$$

Because the functional

$$v \mapsto \frac{1}{2\tau} \|v - v_0\|^2 + G(v)$$

is $1/\tau$ -convex, it follows by Lemma 2 that

$$\frac{1}{2\tau} \|v_1^F - v_1^G\|^2 \leq 2\delta.$$

Now we use that J_τ^F is nonexpansive [2, Eq. (4.0.2)], so

$$\|u_1^F - v_1^F\| = \|J_\tau^F(u_0) - J_\tau^F(v_0)\| \leq \|u_0 - v_0\|.$$

We conclude that

$$\|u_1^F - v_1^G\| \leq \|u_0 - v_0\| + 2\sqrt{\delta\tau}.$$

By iterating this estimate, we derive

$$\|(J_\tau^F)^n[u_0] - (J_\tau^G)^n[v_0]\| \leq \|u_0 - v_0\| + 2n\sqrt{\delta\tau}. \tag{61}$$

The a priori estimate [2, Theorem 4.0.4, (v)] yields that the gradient flows u and v of F and G , respectively, are approximated well by $(J_{t/n}^F)^n[u_0]$ and $(J_{t/n}^G)^n[v_0]$. More precisely,

for $t > 0$ and $n > 0$, the a priori estimate gives

$$\|u(t) - (J_\tau^F)^n\| \leq \frac{Lt}{\sqrt{2n}} \quad \text{and} \quad \|v(t) - (J_\tau^G)^n\| \leq \frac{Lt}{\sqrt{2n}}.$$

By these a priori estimates and the estimate for discrete flows (61), we see that

$$\begin{aligned} \|u(t) - v(t)\| &\leq \|u(t) - (J_{t/n}^F)^n[u_0]\| + \|v(t) - (J_{t/n}^G)^n[v_0]\| \\ &\quad + \|(J_{t/n}^F)^n[u_0] - (J_{t/n}^G)^n[v_0]\| \\ &\leq \sqrt{2}L \frac{t}{n} + 2n\sqrt{\frac{\delta t}{n}} + \|u_0 - v_0\|. \end{aligned}$$

To derive the final estimates, we need to make good choices for n . If $0 \leq t \leq \delta/L^2$, we take $n = 1$ and obtain

$$\|u(t) - v(t)\| \leq \sqrt{2}Lt + 2\sqrt{\delta t} + \|u_0 - v_0\| \leq 4\sqrt{\delta t} + \|u_0 - v_0\|.$$

If $t > \delta/L^2$, we choose $n = \lceil L^{2/3}(t/\delta)^{1/3} \rceil$, which is larger than or equal to 2. In that case,

$$n/2 \leq n - 1 < L^{2/3}(t/\delta)^{1/3} \leq n.$$

We then obtain

$$\|u(t) - v(t)\| \leq 8L^{1/3}\delta^{1/3}t^{2/3} + \|u_0 - v_0\|.$$

□

References

- Ambrosio, L., Ghezzi, R., Magnani, V.: BV functions and sets of finite perimeter in sub-Riemannian manifolds. *Ann. Inst. Henri Poincaré Non Linear Anal.* **32**(3), 489–517 (2015)
- Ambrosio, L., Gigli, N., Savaré, G.: *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Birkhäuser, Basel (2005)
- Baspinar, E., Citti, G., Sarti, A.: A geometric model of multi-scale orientation preference maps via gabor functions. *JMIV* **60**(6), 900–912 (2018)
- Baspinar, E.: *Minimal Surfaces in Sub-Riemannian Structures and Functional Geometry of the Visual Cortex*. Ph.D. thesis, University of Bologna (2018)
- Bekkers, E.: *Retinal Image Analysis using Sub-Riemannian Geometry in $SE(2)$* . Ph.D. thesis, TU/e Eindhoven (2017)
- Bekkers, E., Duits, R., Mashatkov, A., Sanguinetti, G.: A PDE approach to data-driven sub-Riemannian geodesics in $SE(2)$. *SIIMS* **8**(4), 2740–2770 (2015)
- 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: *International Conference on Scale Space and Variational Methods in Computer Vision*, pp. 472–484. Springer, Cham (2019)
- Boscain, U., Chertovskih, R., Gauthier, J.P., Prandi, D., Remizov, A.: Highly corrupted image inpainting by hypoelliptic diffusion. *JMIV* **60**(8), 1231–1245 (2018)
- Brézis, H.: *Opérateurs maximaux monotones et semi-grope de contractions dans les espaces de Hilbert*, vol. 50. North-Holland Publishing Co., Amsterdam (1973)
- Bungert, L., Burger, M., Tenbrinck, D.: Computing nonlinear eigenfunctions via gradient flow extinction. In: *International Conference on Scale Space and Variational Methods in Computer Vision*, pp. 291–302. Springer, Cham (2019)
- Chambolle, A., Pock, T.: Total roto-translation variation. *Numer. Math.* **142**, 611–666 (2019)
- Chirikjian, G.S., Kyatkin, A.B., Buckingham, A.C.: *Engineering applications of noncommutative harmonic analysis: with emphasis on rotation and motion groups*. *Appl. Mech. Rev.* **54**(6), B97–B98 (2001)
- Citti, G., Franceschiello, B., Sanguinetti, G., Sarti, A.: Sub-Riemannian mean curvature flow for image processing. *SIIMS* **9**(1), 212–237 (2016)
- Citti, G., Sarti, A.: A cortical based model of perceptual completion in the roto-translation space. *JMIV* **24**(3), 307–326 (2006)
- Cohen, E., Deffieux, T., Demeñé, C., Cohen, L., Tanter, M.: 3d vessel extraction in the rat brain from ultrasensitive Doppler images. In: *Computer Methods in Biomechanics and Biomedical Engineering*. LNB, pp. 81–91 (2018)
- Cohen, I., Falik, A., Gilboa, G.: Stable explicit p-Laplacian flows based on nonlinear eigenvalue analysis. In: *International Conference on Scale Space and Variational Methods in Computer Vision*, pp. 315–327. Springer, Cham (2019)
- Creusen, E.J., Duits, R., Florack, L., Vilanova, A.: Numerical schemes for linear and non-linear enhancement of DW-MRI. *NM-TMA* **6**(3), 138–168 (2013)
- Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Trans. Image Process.* **16**(8), 2080–2095 (2007)
- Descoteaux, M., Deriche, R., Knosche, T.R., Anwander, A.: Deterministic and probabilistic tractography based on complex fibre orientation distributions. *IEEE Trans. Med. Imaging* **28**(2), 269–286 (2008)
- Duits, R.: *Perceptual organization in image analysis*. Ph.D. thesis, TU/e (2005)
- Duits, R., Creusen, E., Ghosh, A., Dela Haije, T.: Morphological and linear scale spaces for fiber enhancement in DW-MRI. *JMIV* **46**(3), 326–368 (2013)
- 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 $SE(2)$. *QAM-AMS* **68**, 255–292 (2010)
- Duits, R., Janssen, M., Hannink, J., Sanguinetti, G.: Locally adaptive frames in the roto-translation group and their applications in medical image processing. *JMIV* **56**(3), 367–402 (2016)
- Duits, R., Meesters, S., Mirebeau, J., Portegies, J.: Optimal paths for variants of the 2D and 3D Reeds–Shepp car with applications in image analysis. *JMIV* **60**, 816–848 (2018)
- Duits, R., St-Onge, E., Portegies, J., Smets, B.: Total variation and mean curvature PDEs on the space of positions and orientations. In: *International Conference on Scale Space and Variational Methods in Computer Vision*, pp. 211–223 (2019)
- Evans, L.C., Spruck, J.: Motion of level sets by mean curvature. *J. Differ. Geom.* **33**(3), 635–681 (1991)
- Fabbrini, L., et al.: Improved edge enhancing diffusion filter for speckle-corrupted images. *IEEE Geosci. Remote Sens. Lett.* **11**(1), 99–103 (2013)
- Felsberg, M., Forssen, P.E., Schar, H.: Channel smoothing: efficient robust smoothing of low-level signal features. In: *IEEE PAMI*, pp. 209–222 (2006)
- Franken, E.: *Enhancement of crossing elongated structures in images*. PhD thesis, Technical University Eindhoven (2008)

30. Franken, E.M., Duits, R.: Crossing preserving coherence-enhancing diffusion on invertible orientation scores. *IJCV* **85**(3), 253–278 (2009)
31. Ghimpeanu, G., Batard, T., Bertalmio, M., Levine, S.: A decomposition framework for image denoising algorithms. *IEEE Trans. Image Process.* **25**(1), 388–399 (2015)
32. Giga, Y., Sato, M.H.: Generalized interface evolution with the Neumann boundary condition. *Proc. Jpn. Acad. Ser. A Math. Sci.* **67**(8), 263–266 (1991)
33. Janssen, M.H.J., Janssen, A.J.E.M., Bekkers, E.J., Bescós, J.O., Duits, R.: Processing of invertible orientation scores of 3d images. *JMIV* **60**(9), 1427–1458 (2018)
34. Janssen, M.H.J., Dela Haije, T.C.J., Martin, F.C., Bekkers, E.J., Duits, R.: The Hessian of axially symmetric functions on $SE(3)$ and application in 3D image analysis. In: LNCS (Proceedings of SSVM), pp. 643–655 (2017)
35. Lebrun, M.: An analysis and implementation of the BM3D image denoising method. *Image Process. Line* **2**, 175–213 (2012)
36. Lee, J.M.: Smooth manifolds. In: *Introduction to Smooth Manifolds*, pp. 1–31. Springer, New York (2013)
37. Martin, F., Bekkers, E., Duits, R.: Lie analysis package. www.lieanalysis.nl/ (2017). Accessed 5 Oct 2019
38. Meesters, S., et al.: Stability metrics for optic radiation tractography: towards damage prediction after resective surgery. *J. Neurosci. Methods* **288**, 34–44 (2017)
39. MomayyezSiahkal, P., Siddiqi, K.: 3D stochastic completion fields for mapping connectivity in diffusion MRI. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(4), 983–995 (2012)
40. Neher, P.F., Laun, F.B., Stieltjes, B., Maier-Hein, K.H.: Fiberfox: facilitating the creation of realistic white matter software phantoms. *Magn. Reson. Med.* **72**(5), 1460–1470 (2014)
41. Nesterov, Y.: *Introductory Lectures on Convex Programming: A Basic Course*. Springer, Berlin (2004)
42. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.* **12**(7), 629–639 (1990)
43. Portegies, J.M., Duits, R.: New exact and numerical solutions of the (convection-) diffusion kernels on $SE(3)$. *DGA* **53**, 182–219 (2017)
44. Portegies, J.M., Fick, R., Sanguinetti, G.R., Meesters, S.P.L., Girard, G., Duits, R.: Improving fiber alignment in HARDI by combining contextual PDE flow with constrained spherical deconvolution. *PLoS ONE* **10**(10), e0138122 (2015)
45. Portegies, J.: PDEs on the Lie Group $SE(3)$ and their applications in diffusion-weighted MRI. Ph.D. thesis, Department of Mathematics, TU/e (2018)
46. Reisert, M., Burkhardt, H.: Efficient tensor voting with 3d tensorial harmonics. In: *IEEE Conference, CVPRW '08*, pp. 1–7 (2008)
47. Reisert, M., Kiselev, V.G.: Fiber continuity: an anisotropic prior for ODF estimation. *IEEE Trans. Med. Imaging* **30**(6), 1274–1283 (2011)
48. Rockafellar, R.T., Wets, R.J.B.: *Variational Analysis*, vol. 317. Springer, Berlin (2009)
49. Sapiro, G.: *Geometric Partial Differential Equations and Image Analysis*. Cambridge University Press, Cambridge (2006)
50. Sato, M.H.: Interface evolution with Neumann boundary condition. *Adv. Math. Sci. Appl.* **4**(1), 249–264 (1994)
51. Schmidt, M., Weickert, J.: Morphological counterparts of linear shift-invariant scale-spaces. *J. Math. Imaging Vis.* **56**(2), 352–366 (2016)
52. Tournier, J.D., Calamante, F., Connelly, A.: Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution. *Neuroimage* **35**(4), 1459–1472 (2007)
53. Vogt, T., Lellmann, J.: Measure-valued variational models with applications to diffusion-weighted imaging. *JMIV* **60**(9), 1482–1502 (2018)
54. Weickert, J.A.: Coherence-enhancing diffusion filtering. *Int. J. Comput. Vis.* **31**(2/3), 111–127 (1999)


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



Bart M. N. Smets received his M.Sc. degree in Industrial and Applied Mathematics in 2019 at the TU/e, The Netherlands. He is currently a Ph.D. student at the TU/e under the supervision of dr.ir. R. Duits at the Mathematics and Computer Science Department. His research interests include differential geometry, machine learning and their applications to image processing.



Jim Portegies obtained M.Sc. degrees in Industrial and Applied Mathematics and in Applied Physics from the TU/e in 2009. He received his Ph.D. in Mathematics from the Courant Institute of Mathematical Sciences in New York in 2014. He spent the Fall semester of 2013 at NYU Shanghai in China. From 2014 to 2016 he was a postdoc at the Max Planck Institute for Mathematics in the Sciences in Leipzig, Germany. Currently, he is an Assistant Professor at TU/e, specializing in Geometric Analysis and Artificial Intelligence.



Etienne St-Onge received his M.Sc. degree in Computer Science at the Université de Sherbrooke. Currently, he is a Ph.D. candidate in medical image analysis under Pr. Maxime Descoteaux in the Computer Science department, Université de Sherbrooke. As a member of the Sherbrooke Connectivity Imaging Lab (SCIL), his research focuses on discrete differential geometry applied to Diffusion-Weighted Magnetic Resonance Imaging (DW-MRI). For his graduate studies, he was awarded with a scholarship from the Fonds de Recherche du Québec–Nature et technologies (FRQNT) and the Natural Sciences and Engineering Research Council (NSERC) of Canada.



Remco Duits received his M.Sc. degree (cum laude) in Mathematics in 2001 at the TU/e, The Netherlands. He received his Ph.D.-degree (cum laude) at the Department of Biomedical Engineering at the TU/e. Now he is associate professor at the Department of Applied Mathematics and Computer Science and affiliated part-time at the Biomedical Image Analysis group at the Biomedical Engineering Department. He has received an ERC-StG 2013 personal grant (Lie analysis, no. 335555). His research

interests subsume group theory, functional analysis, differential geometry, PDE's, harmonic analysis, geometric control and their applications to biomedical imaging and vision.