



Currents and Finite Elements as Tools for Shape Space

James Benn¹ · Stephen Marsland² · Robert I. McLachlan³ · Klas Modin⁴ · Olivier Verdier^{5,6}

Received: 7 August 2017 / Accepted: 6 August 2019 / Published online: 19 August 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

The nonlinear spaces of shapes (unparameterized immersed curves or submanifolds) are of interest for many applications in image analysis, such as the identification of shapes that are similar modulo the action of some group. In this paper, we study a general representation of shapes as currents, which are based on linear spaces and are suitable for numerical discretization, being robust to noise. We develop the theory of currents for shape spaces by considering both the analytic and numerical aspects of the problem. In particular, we study the analytical properties of the current map and the H^{-s} norm that it induces on shapes. We determine the conditions under which the current determines the shape. We then provide a finite element-based discretization of the currents that is a practical computational tool for shapes. Finally, we demonstrate this approach on a variety of examples.

Keywords Currents · Finite elements · Shape space · Image analysis

Mathematics Subject Classification 32U40m · 62M40 · 65D18 · 74S05

1 Introduction

‘Shape’, wrote David Mumford in his 2002 ICM address [23], ‘is the ultimate nonlinear thing’. In the study of shapes,

one seeks analytic and numerical methods to compare them, to classify them, to recognize them, to evolve them—to understand them. As is usual in mathematics, nonlinear objects are constructed out of linear ones—vector spaces—using maps, quotients, and open subsets. Embeddings of an m -dimensional manifold M into an n -dimensional manifold N , which are equivalent up to reparametrization by diffeomorphisms, crystallize our intuition of shape. A *shape space* $\mathcal{S}$, then, is realized as the complete set of embeddings quotiented by diffeomorphisms of the domain in which each coset represents a particular shape: $\mathcal{S} = \text{Emb}(M, N)/\text{Diff}(M)$. For example, the set of shapes $\mathcal{S}_{sc}$ corresponding to simple closed curves is constructed as follows: one first considers the set of smooth embeddings $\text{Emb}(S^1, \mathbb{R}^2)$ of S^1 into $\mathbb{R}^2$. A smooth reparameterization of an embedding ϕ corresponds to composition on the right by a diffeomorphism of S^1 , i.e. a map $\eta \in \text{Diff}(S^1)$. Two embeddings ϕ and ψ are equivalent if and only if they differ by such a reparameterization ($\psi = \phi \circ \eta$) so that $\mathcal{S}_{sc}$ is the quotient space $\text{Emb}(S^1, \mathbb{R}^2)/\text{Diff}(S^1)$. In addition to embeddings, we shall also consider the wider class of immersions so that a general shape space is $\mathcal{S} = \text{Imm}(M, N)/\text{Diff}(M)$.

We wish to recognize examples of the same shape: those that are identical up to the action of diffeomorphisms and possibly noise or other obfuscation. This can be studied by

✉ Stephen Marsland
stephen.marsland@vuw.ac.nz

James Benn
jbenn2@alumni.nd.edu

Robert I. McLachlan
r.mclachlan@massey.ac.nz

Klas Modin
klas.modin@chalmers.se

Olivier Verdier
olivier.verdier@hvl.no; olivier@kth.se

¹ Palmerston North, New Zealand

² School of Mathematics and Statistics, Victoria University of Wellington, Wellington, New Zealand

³ Institute of Fundamental Sciences, Massey University, Palmerston North, New Zealand

⁴ Department of Mathematical Sciences, Chalmers University of Technology and University of Gothenburg, Gothenburg, Sweden

⁵ Department of Computing, Mathematics and Physics, Western Norway University of Applied Sciences, Bergen, Norway

⁶ Department of Mathematics, KTH, Stockholm, Sweden

finding a metric that is blind to changes inside the conjugacy classes of $\mathcal{S}$, or by finding a representation of the shapes that is invariant under the action of $\text{Diff}(M)$. An example of the latter for $\mathcal{S}_{\text{sc}}$ is given by the *differential invariant signature*—an influential new paradigm introduced by Calabi et al. [10]. To recognize elements of $\mathcal{S}_{\text{sc}}$ up to Euclidean transformations, each element of $\mathcal{S}_{\text{sc}}$ is represented by its Euclidean curvature κ and corresponding derivative with respect to arclength κ_s , which we call a signature curve $(\kappa, \kappa_s)(S^1) \subset \mathbb{R}^2$. The signature curve is clearly invariant under Euclidean motions *and* under reparameterizations by diffeomorphisms of S^1 . As Calabi et al. comment, ‘The recognition problem includes a comparison principle that would be able to tell whether two signature curves are close in some sense. Thus, we effectively reduce the group-invariant recognition problem to the problem of imposing a “metric” on the space of shapes but now by “shape” we mean the signature curve, not the original object’. This observation forms the motivating perspective for the present work in the context of functional analysis and geometric measure theory, specifically the method of currents.

The method of shape currents, first suggested by Glaunès et al. [17], is based on embedding the nonlinear shape space into a linear vector space endowed with a metric, thereby allowing the construction of flexible families of metrics on shape spaces that are easy to compute. It is robust to noise and provides for control of the resolution and accuracy of the representation. A shape current is a map that takes an immersion $\phi: M \rightarrow N$ and associates it with an element in the topological dual (denoted by $*$) of differential m -forms α over N via

$$[\phi](\alpha) := \int_M \phi^* \alpha. \quad (1)$$

Observe that this map is invariant under orientation-preserving reparameterizations (see Proposition 1). In this way, nonlinear shape space is mapped to a *linear* space of functionals on forms.

The main ingredient in the approach taken by Glaunès et al. [17] is that of the reproducing kernel Hilbert space (RKHS). The RKHS is obtained by first selecting a kernel K , typically a Gaussian kernel, then constructing a Hilbert space so that K is the reproducing kernel. The kernel K contains a parameter that allows one to control the metric on the space of currents and hence the measure of the distance between shapes. This construction has been developed further to allow metrics on shape currents to be combined with optimization routines for registration, typically by deforming the shape through the left action of a diffeomorphism group (e.g. [19]; see Sect. 1.1 below for more details of the use of currents in shape analysis). Here, we take the viewpoint of Calabi et al. that left symmetries (registration) are

taken care of by computing a signature curve, so the only remaining step is to impose a metric on shape space. In this way, we obtain a computable geometry on shapes modulo any classical transformation group; see Example 10.

Our approach differs from that of Glaunès et al. in that first a particular Hilbert space structure for the space of differential forms is selected, and then an isomorphism between the space of currents and differential forms is constructed via the Riesz representation theorem. The Green’s function corresponding to an elliptic partial differential operator defining the Hilbert structure is used to obtain a metric on currents. This is carried out in the context of immersed curves as follows: take $M = S^1$ and $N = \mathbb{T}^2$ —the flat 2-torus, which is constructed by taking a sufficiently large square in $\mathbb{R}^2$ and identifying opposite sides. Choosing N in this way (compact and without boundary) allows us to circumvent certain technical analytic issues; however, one could just as well consider N as an open, bounded subset of $\mathbb{R}^2$ since the geometries are identical. We endow the class of immersions and embeddings with the Lipschitz topology, while for the differential forms Λ we use the Sobolev H^s topology, where s is a non-negative integer chosen so as to guarantee that the current map (see Definition 4),

$$[\cdot]: \text{Imm}(S^1, \mathbb{T}^2) \rightarrow H^s(\Lambda^1(\mathbb{T}^2))^* \quad (2)$$

defined through Eq. (2) with $M = S^1$, is continuous. By altering the above construction slightly, we are able to show that the current map is also differentiable. In particular, embedding $\mathbb{T}^2$ in the flat 3-torus $\mathbb{T}^3$ and choosing s to be slightly larger, the current map

$$[\cdot]: \text{Imm}(S^1, \mathbb{T}^2) \rightarrow H^s(\Lambda^1(\mathbb{T}^3))^*, \quad (3)$$

is differentiable. This regularity is important: to use invariants to recognize shapes, they must have the property that nearby shapes have nearby invariants. A consequence of differentiability is that we obtain a strong Riemannian metric on the space of shape currents via restriction of the Sobolev metric on the ambient space of 1-forms so that the tools of differential geometry may be brought to bear on the problem of shape recognition, comparison, classification, and evolution.

Although subtle, the differences in our approach to shape currents are material and consequential. The metric we construct is of Sobolev type and different to the metrics that have previously been studied for shape currents. Sobolev metrics tend to have nice geometric properties due to the ellipticity of the differential operators defining them, while Sobolev spaces allow one to study objects of low regularity, as are often encountered in practice; in our setting, we are able to obtain analytic results for shapes that are Lipschitz immersions. Sobolev spaces have very useful com-




Fig. 1 A simple example of finite element currents using piecewise constant elements. An oriented shape (the image of a simple closed curve ϕ) is given in red on the left. A triangular mesh is laid over the domain. For each triangle T in the mesh, the finite element currents $\int_{\phi(S^1) \cap T} dx$ and $\int_{\phi(S^1) \cap T} dy$ are computed numerically. These determine the triangles (shown as solid black lines on the right) that intersect the shape,

and the x - and y -extents of the shape on those triangles. These data determine a piecewise linear approximant that interpolates the shape and the element edges. Higher-order currents, such as $\int_{\phi(S^1) \cap T} y dx$, are also possible and give a more accurate representation of the shape. See also Figs. 8 and 9 (Color figure online)

pactness properties, which can have useful consequences for the geometry of shape space. The RKHS property is not explicitly used in our approach, and moreover the Sobolev space H^1 is not an RKHS; we retain local information on the shape through Sobolev trace theorems (see Remark 5).

In this paper, we present both an analytic study of currents and a numerical approach based on a finite element discretization. The analytic study, Sect. 2, is structured as follows: in Sect. 2.1 we introduce, first, the Sobolev topology on the space of differential k -forms and, then, the spaces of Lipschitz embeddings and immersions; in Sect. 2.2 we define the current map on the Sobolev H^s differential forms and collect several of its invariance properties. Continuity and differentiability properties of the current map are stated and proved in Sect. 2.3. Specifically, we show that:

- the image of a Lipschitz immersion ϕ under the current map (2) lies in $H^s(\Lambda^1(\mathbb{T}^2))^*$, which generalizes [17, Proposition 1],
- the current map (2) is Hölder-continuous for $s \geq 1$,
- the current map (3) is differentiable for $s \geq 2$;

Finally, in Sect. 2.4, we describe the construction of the metric on currents via the Riesz representation theorem and, for a given subspace of functions $S^1 \rightarrow \mathbb{T}^2$ and a given linear space of 1-forms on $\mathbb{T}^2$, we show that the current map $\phi \mapsto [\phi]$ essentially determines the shape of ϕ .

Our numerical study of currents (Sect. 3) consists of discretizing our analytic constructions. The currents can be evaluated for all α in a finite element space, which gives a flexible and general representation of shapes. Specifically, we introduce spaces of finite elements V on S^1 and W on $\Lambda^1(\mathbb{T}^2)$ and evaluate the current $[\phi_V]_W$, where ϕ_V is the approximation of ϕ in V . A simple example is shown in Fig. 1. In this case, piecewise constant elements determine a piece-

wise constant approximant that interpolates the shape at the element edges. The discrete shape ϕ_V is mapped to the finite-dimensional vector space W , which inherits an inner product metric from $\Lambda^1(\mathbb{T}^2)$. The approximation $[\phi_V]_W$ contains more information than just its norm, namely its coordinates in W ; and the approximations in W of a set of shapes contain more information than just their pairwise distances. It is generally convenient to work with data in a normed vector space, as expected by the majority of statistical and learning algorithms.

In Sect. 3, we show the following results:

- In Sect. 3.2, we demonstrate that the quadrature errors in evaluating the currents are typically small and that the method is robust in the presence of noise. This robustness stems from the cancellation property of oscillatory integrals. The currents can be accurately computed even for very rough shapes (not even Lipschitz) and for noisy shapes.
- In Sect. 3.3, we study how accurately the discretized currents determine the shapes. This is a question in approximation theory.
- In Propositions 8 and 9, we show that the order of approximation is 2 for piecewise constant, 3 for piecewise linear, and 5 for piecewise quadratic elements in W .
- In Sect. 3.4, we introduce a discretization of the metric on shapes so that each shape is approximated by a point in a vector space equipped with an Euclidean metric.

Finally, we give several numerical examples of the geometry of shape space induced by the discretization in Sect. 4. The method not only compares pairs of shapes, it provides a direct approximation of shape space and its geometry: we present numerical experiments (e.g. Examples 9 and 14) applying

principal components analysis directly to the current representation in order to successfully separate classes of shapes.

1.1 Related Literature

As was mentioned previously, currents have already been used in shape analysis, primarily for curve and surface matching. In [27], a surface in 3D was represented with currents defined on a surface mesh, and a norm was computed to enable the matching of two surfaces using the currents. This allows the calculation of the pairwise distances between a set of shapes, with each distance evaluated in an RKHS metric as a double integral over all pairs (a_1, a_2) of points or elements in the shapes with a_i in shape $i = 1, 2$. Fast algorithms based on multipole expansions are available to evaluate this integral [8]. This type of metric can lie in the same class as that considered here (but see Sect. 2.3), but to our knowledge the direct representation of shapes in a linear space, i.e. the direct approximation of the currents themselves, has not previously been considered. The authors developed a way to match shape currents in the spirit of the large deformation diffeomorphic metric mapping (LDDMM) framework [6], and the method was demonstrated on surfaces representing shapes and hippocampi. A variation on this approach for curves rather than surfaces was developed in [17] (where the currents are referred to as vector-valued measures). Again, the aim is a matching algorithm where curves can be deformed onto each other. LDDMM approaches can map *pairs* of shapes into a vector space (e.g. of autonomous or non-autonomous vector fields), but not individual shapes; the geometry induced in that vector space is different from the one considered here.

In [15], a different benefit of the linear representation provided by currents is recognized, which is that they provide a useful space in which to perform statistical analysis of the deformations between shapes. This was originally considered in [16], but there is a difficulty that the mean (template) shape has to be defined in such a way that both the shape and deformation are amenable to statistical analysis. A matching pursuit algorithm is defined in [15] to provide a computationally tractable representation of shape currents as the number of shapes grows.

The use of currents for matching was extended by Charon and Trounev in [13] as functional currents, where a function is added to the current representation so that the deformation of a shape and some function defined on its surface can be considered simultaneously. The same authors also considered how to deal with cases where the orientation property of currents is undesirable. The fact that a large spike in the appearance of the shape cancels out in the current representation as the positive and negative contributions are virtually identical is a benefit when considering the currents for dealing with noisy representations of shapes. However, in cases where these spikes can truly exist, or where there is orien-

tation information in the image, but it can differ arbitrarily by sign, such as in diffusion tensor images of the brain, the oriented manifold is a disadvantage. This leads to the consideration of varifolds in [12], where the registration of some directed surfaces is demonstrated. Currents and varifolds can be combined by using RKHS kernels that depend on both position and derivative information [19,24].

While there are many other computational approaches to shape space, as far as we are aware they all involve determining a point correspondence between shapes via optimization and/or working directly in the nonlinear shape space; see, for example, [5,11].

In contrast to the aforementioned work, we focus here solely on shape currents as a way to represent shapes in a linear space, to induce a geometry on shapes, and so compute statistics; we assume that registration (left-matching) has already been taken care of, for example through the signature curve, as suggested by Calabi et al. [10].

2 Currents and Their Induced Metric on Shapes

For any natural number k , we denote by $\Lambda^k(N)$ the space of smooth k -forms on a manifold N . We will introduce the Sobolev topology on $\Lambda^k(N)$ below and call an element of the corresponding topological dual a k -current.

Currents were introduced by de Rham [14]. They are natural generalizations (or completions) of the pairing by integration between a k dimensional submanifold $M \subset N$ and smooth k -forms on N . Currents are instrumental in geometric measure theory, where they are used to study a very wide class of (not necessarily smooth) subsets of $\mathbb{R}^n$, for example in minimal surface problems [22]. In that field, the functions ϕ are typically Lipschitz and the differential forms are smooth. We allow non-smooth shapes such as those represented by Lipschitz functions, but as we want to discretize the differential forms by finite elements, we will allow the forms to be non-smooth as well.

2.1 A Hilbert Space Structure for Differential Forms

When introducing the Sobolev topology on differential forms, it will be convenient for us to take $M = S^1$ and $N = \mathbb{T}^n$ (the flat n -torus, constructed by choosing a sufficiently large box in $\mathbb{R}^n$ and identifying opposite sides). In later sections, we will restrict our attention to dimensions $n = 2$ and $n = 3$. The advantage of considering N in this way (compact and without boundary) is that we are able to avoid certain technical analytic complications involved with open sets or sets with boundaries in $\mathbb{R}^n$ while retaining the same Euclidean geometry.

Let g be the flat Euclidean metric on $\mathbb{T}^n$. If α is a 1-form, then $\alpha^\sharp$ is the vector field such that $\alpha(v) = g(\alpha^\sharp, v)$ for any vector field v . In this way, we obtain an L^2 metric on $\Lambda^1(\mathbb{T}^n)$ defined by

$$\langle \alpha, \beta \rangle = \int_{\mathbb{T}^n} g(\alpha^\sharp, \beta^\sharp) d\mu,$$

where μ is the standard Lebesgue measure on $\mathbb{T}^n$. For forms of degree k , we first construct a pointwise inner product on $\Lambda^k(T^*\mathbb{T}^n)$ via

$$\begin{aligned} &(\alpha_1 \wedge \cdots \wedge \alpha_k, \beta_1 \wedge \cdots \wedge \beta_k)(x) \\ &= \sum_{\pi} \text{sign}(\pi) g(\alpha_1^\sharp, \beta_{\pi(1)}^\sharp)(x) \cdots g(\alpha_k^\sharp, \beta_{\pi(k)}^\sharp)(x), \end{aligned}$$

where π ranges over the set of permutations of $\{1, \dots, k\}$, and then define an L^2 metric on k -forms $\Lambda^k(\mathbb{T}^n)$ (that is, sections of the k th-fold wedge product of $T^*\mathbb{T}^n$) by

$$\langle \alpha, \beta \rangle_{L^2(\mathbb{T}^n)} = \int_{\mathbb{T}^n} (\alpha, \beta) d\mu. \quad (4)$$

We have the operator d , which takes k -forms to $(k+1)$ -forms, and its formal dual δ , which takes $(k+1)$ -forms to k -forms. In particular, for any k -form α and any $(k+1)$ -form β , we have

$$\int_{\mathbb{T}^n} (d\alpha, \beta) d\mu = \int_{\mathbb{T}^n} (\alpha, \delta\beta) d\mu.$$

The Hodge star operator $\star$ maps k -forms to $(n-k)$ -forms and is defined by $\langle \alpha, \beta \rangle \mu = \alpha \wedge \star \beta$, which relates d and δ by

$$\delta = (-1)^{nk-n-1} \star d \star.$$

Let $\Delta := d\delta + \delta d$ denote the Hodge Laplacian which is positive-definite and self-adjoint in the L^2 metric on k -forms.

Definition 1 For any non-negative integer s and real number σ , we define the Sobolev space $H_\sigma^s(\Lambda^k(\mathbb{T}^n))$ to be the completion of $\Lambda^k(\mathbb{T}^n)$ with respect to the inner product

$$\langle \alpha, \beta \rangle_{H_\sigma^s} := \int_{\mathbb{T}^n} (\alpha, (I + \sigma^{2s} \Delta^s) \beta) d\mu. \quad (5)$$

Formally, one can think of $H_\sigma^s(\Lambda^k(\mathbb{T}^n))$ as the Hilbert space of forms whose weighted derivatives up to order s are bounded in the L^2 norm defined by (4).

Remark 1 Observe that the operator d takes k -forms of class H^s to $(k+1)$ -forms of class H^{s-1} , the operator δ takes $(k+1)$ -forms of class H^s to k -forms of class H^{s-1} , and the operator Δ takes k -forms of class H^{s+1} to k -forms of H^{s-1} .

Remark 2 In view of the compactness of $\mathbb{T}^n$, we have the following string of inequalities

$$\begin{aligned} &\|\cdot\|_{L^2(\mathbb{T}^n)} \\ &\lesssim \|\cdot\|_{H^1(\mathbb{T}^n)} \lesssim \cdots \lesssim \|\cdot\|_{H^{s-1}(\mathbb{T}^n)} \lesssim \|\cdot\|_{H^s(\mathbb{T}^n)} \lesssim \cdots, \end{aligned}$$

and consequent nesting of Sobolev spaces

$$\begin{aligned} &\cdots \subset H_\sigma^s(\Lambda^1(\mathbb{T}^n)) \subset H_\sigma^{s-1}(\Lambda^1(\mathbb{T}^n)) \subset \\ &\cdots \subset H_\sigma^1(\Lambda^1(\mathbb{T}^n)) \subset L^2(\Lambda^1(\mathbb{T}^n)). \end{aligned}$$

These two facts will be implicitly used throughout Sect. 2.3.

Remark 3 Since the Hodge Laplacian is an unbounded operator that is self-adjoint in the L^2 metric (4), we may use the functional calculus (see Schmudgen [25]) to define Δ^t for any real number t . In particular, we may extend the definition of the Sobolev space $H_\sigma^s(\Lambda^k(\mathbb{T}^n))$ to any real number s .

The dual of $H_\sigma^s(\Lambda^k(\mathbb{T}^n))$ is $H_\sigma^{-s}(\Lambda^k(\mathbb{T}^n))$ and will be called the space of H_σ^{-s} currents; see Taylor [26] for a formal introduction to Sobolev spaces and their duals.

We will often omit the scale parameter σ in the sequel for the sake of readability. To understand why σ is called a scale parameter, consider the scaling function $\lambda \cdot x := \lambda x$, for any positive real number λ . This function $(\lambda \cdot)$ acts on 1-forms by pull-back, that is (with a slight abuse of notation) $\lambda \cdot \alpha := (\lambda^{-1} \cdot)^* \alpha$. The resulting form is $\lambda \cdot \alpha(x) = \lambda^{-1} \sum_{i=1,2} \alpha_i(\lambda^{-1} x_1, \lambda^{-1} x_2) dx_i$, so we see that $(\lambda \cdot \alpha, \lambda \cdot \beta)_{H_\sigma^s} = \lambda^{-1} (\alpha, \beta)_{H_{\sigma/\lambda}^s}$. We will also see in Sect. 2.4 that σ determines the length scale at which distances between shapes are measured.

The properties of the map from curves to H^{-s} currents depend not only on the topology of the currents, but also on that of the curves. We wish to allow very large classes of curves—indeed, one of the strengths of currents is that they allow this. We will define two spaces of curves: the first being the space of Lipschitz immersions and the second being the space of Lipschitz embeddings.

Definition 2 The space of *Lipschitz immersions* is the set of curves $\phi: S^1 \rightarrow \mathbb{T}^n$ whose components are Lipschitz functions and whose tangent vector ϕ' , wherever it is defined (which is almost everywhere), is nonzero. We call the space of such curves $\text{LipImm}(S^1, \mathbb{T}^n)$.

Definition 3 The space of *Lipschitz embeddings* ($\text{LipEmb}(S^1, \mathbb{T}^n)$) is the set of curves $\phi: S^1 \rightarrow \mathbb{T}^n$ that are injective Lipschitz immersions and homeomorphisms onto their images in $\mathbb{T}^n$.

Remark 4 A useful fact that we shall employ throughout the first half of this work is Rademacher's theorem which shows that any Lipschitz function on a compact manifold M belongs

to the Sobolev space $H^1(M)$ (see Hebey [18] for an excellent presentation of this and related results). For a curve $\phi = (\phi_1, \phi_2) \in \text{LipImm}(S^1, \mathbb{T}^n)$, the component functions ϕ_1 and ϕ_2 are Lipschitz functions and we shall measure these functions in a coordinate chart $V = (a, b) \subset \mathbb{R}$ on S^1 with the H^1 norm defined by

$$\|\phi_i\|_{H^1}^2 = \int_V |\phi_i(t)|^2 + |\phi_i'(t)|^2 dt. \quad (6)$$

Suppose $\psi(t)$ is a C^1 curve in $\text{LipImm}(S^1, \mathbb{T}^n)$ with $\psi(0) = \phi$ and $\partial_t|_{t=0}\psi(t) = v$. Then, v is a Lipschitz vector field along $\phi(S^1)$ and belongs to the Sobolev space $H^1(\phi(S^1))$.

2.2 De Rham Currents and Their Invariance

The current map (given below in Definition 4) is motivated by the signed area $\int_{\phi(S^1)} y dx$ enclosed by a closed curve. This is induced from a 1-form ($y dx$) on $\mathbb{R}^2$.¹ Currents are invariant under orientation-preserving diffeomorphisms of M —sense-preserving reparameterizations in the case of curves—and hence can be used to factor out the group $\text{Diff}^+(S^1)$, the subgroup of orientation-preserving diffeomorphisms. As can be seen by considering the length and area of a noisy curve (see the example in Fig. 7), currents are very robust to noise. Although we are in the particular setting $M = S^1$ and $N = \mathbb{T}^n$, the following definition and Proposition 1 hold for arbitrary manifolds M and N with appropriate notational changes.

Definition 4 Let $\phi \in \text{LipImm}(S^1, \mathbb{T}^n)$. The *current map* $[\cdot]$ is a mapping

$$[\cdot] : \text{LipImm}(S^1, \mathbb{T}^n) \rightarrow H^{-s}(\Lambda^1(\mathbb{T}^n)), \quad (7)$$

so that $[\phi]$ is a linear functional on forms defined by

$$[\phi](\alpha) := \int_{S^1} \phi^* \alpha. \quad (8)$$

Proposition 1 *The current map defined in Definition 4 is invariant with respect to orientation-preserving reparameterizations, that is $[\phi \circ \psi] = [\phi]$ for all orientation-preserving diffeomorphisms ψ of S^1 . In other words, the current map $\phi \mapsto [\phi]$ is $\text{Diff}^+(S^1)$ -invariant (where $\text{Diff}^+(S^1)$ denotes the space of orientation-preserving diffeomorphisms on S^1), and it induces a map from $\text{Imm}(S^1, \mathbb{T}^n)/\text{Diff}^+(S^1)$ to $H^{-s}(\Lambda^1(\mathbb{T}^n))$.*

Proof This is just the statement that integration of forms is well-defined, i.e. independent of the choice of coordinates. Specifically,

$$\int_{S^1} (\phi \circ \psi)^* \alpha = \int_{S^1} \psi^* \phi^* \alpha = \int_{\psi(S^1)} \phi^* \alpha = \int_{S^1} \phi^* \alpha$$

for orientation-preserving ψ . $\square$

Remark 5 For the current map (Definition 4) to exist, we need to be able to take *traces* of a form on any curve; in particular, we want this operation to be continuous (see Proposition 2). Given a bounded domain Ω in $\mathbb{R}^n$, with Lipschitz boundary $\partial\Omega$, a typical function f in $H^s(\Omega)$, with $s \leq \frac{n}{2}$, is not continuous and is only defined almost everywhere in Ω . Moreover, the boundary $\partial\Omega$ has n -dimensional Lebesgue measure zero and hence there is no direct meaning we can give to the expression ‘ f restricted to $\partial\Omega$ ’. The notion of a *trace operator* and the trace theorem resolves this issue for us: if a function f lies in $H^s(\Omega)$, then the restriction of f to a surface S of codimension 1 lies in $H^{s-\frac{1}{2}}(S)$, see [1] and Remark 3. It is important that we pick a Sobolev index s (necessarily depending on n) that is sufficiently high as to guarantee that the current map is defined. In Sect. 2.3, we shall choose an index that not only guarantees that the map is defined, but also continuous and, if chosen high enough, is differentiable.

Currents measure some particular aspects of curves. Consider the case of planar curves, and in particular, the curve $\phi(t) = (\cos(t), 0)$ for $0 \leq t \leq 2\pi$. This shape retraces itself in opposite directions, so $[\phi](\alpha) = 0$ for all α ; currents cannot distinguish this curve from the 0 curve. In the case that $\phi(S^1)$ is a submanifold of $\mathbb{T}^n$, $[\phi]$ is the m -current of integration on $\phi(S^1)$.

2.3 Properties of the Current Map

Having defined the Sobolev topology on differential forms as well as the topology on immersions and embeddings, we proceed to collect several results on the continuity and differentiability of the current map. The following result generalizes [17, Proposition 1] from piecewise C^1 to Lipschitz immersions.

Proposition 2 *Let $s \geq 1$ and let $\phi: S^1 \rightarrow \mathbb{T}^2$ be a Lipschitz immersion. Then, $[\phi]: H^s(\Lambda^1(\mathbb{T}^2)) \rightarrow \mathbb{R}$ is a bounded linear operator. In particular, for a Lipschitz immersion ϕ the corresponding current $[\phi]$ is an element of the Sobolev dual $H^{-s}(\Lambda^1(\mathbb{T}^2))$ and satisfies*

$$\|[\phi]\|_{H^{-s}(\Lambda^1(\mathbb{T}^2))} \leq C \quad (9)$$

for some positive, finite constant C .

¹ This also happens to be Euclidean-invariant, but this is *not* relevant to the sequel.




Fig. 2 Schematic representation of the induced geometry on shapes. The left-hand side shows the vector space of Lipschitz immersed curves. It is partitioned into equivalence classes such as A , B , and C , where curves are equivalent if they are related by a sense-preserving reparameterization. Each equivalence class maps under the current map $[\phi]$ to a single point on the right-hand side, which shows the vector space of

linear forms on 1-forms equipped with the operator norm induced by the H^s -metric on 1-forms. The set of all Lipschitz immersed curves maps into a very small subset of $H^{-s}(\Lambda^1(\mathbb{T}^3))$. This subset is labelled $[\text{LipImm}(S^1, \Omega)]$ on the right. The distance between two shapes is measured by the ‘straight line distance’ in the normed vector space $H^{-s}(\Lambda^1(\mathbb{T}^3))$

Proof Let $\phi \in \text{LipImm}(S^1, \mathbb{T}^2)$ and choose any $\alpha \in H^1(\Lambda^1(\mathbb{T}^2))$. We may restrict α to $\phi(S^1)$ so that the components of α are L^2 by the trace theorem. Then, using Hölder’s inequality we have

$$\begin{aligned} |[\phi](\alpha)| &= \left| \int_{\phi(S^1)} \alpha \right| \\ &\leq \int_{\phi(S^1)} |\alpha| \\ &\leq \text{Length}(\phi(S^1)) \cdot \|\alpha|_{\phi(S^1)}\|_{L^2} \\ &\leq \text{Length}(\phi(S^1)) \cdot \|\alpha\|_{H^1(\mathbb{T}^2)}. \end{aligned}$$

Taking the supremum over $\|\alpha\|_{H^1(\mathbb{T}^2)} = 1$, we obtain

$$\|[\phi]\|_{H^{-s}(\Lambda^1(\mathbb{T}^2))} := \sup_{\substack{\alpha \in H^s(\Lambda^1(\mathbb{T}^2)) \\ \|\alpha\|_{H^s} = 1}} |[\phi](\alpha)| \leq \text{Length}(\phi(S^1)). \quad (10)$$

□

Proposition 2 shows that currents map shapes into the Hilbert space of dual Sobolev differential forms (in a highly nonlinear way) and we can measure the similarity of two shapes by measuring distance between their currents using the norm on that Hilbert space: for two shapes ϕ_1, ϕ_2 , their distance is (see Fig. 2 for a pictorial version of this):

$$d(\phi_1, \phi_2) := \|[\phi_1] - [\phi_2]\|_{H^{-s}(\Lambda^1(\mathbb{T}^2))}.$$

At first sight, this appears to be a wasteful representation of shapes, as it uses forms on the higher-dimensional space $\mathbb{T}^2$ instead of on S^1 . However, in finite-dimensional examples in which quotient spaces are represented using invariants, it is common that large numbers of invariants are needed.² Furthermore, this approach allows one to represent much larger classes of objects than the smooth embeddings, such as weighted, non-smooth, and immersed shapes.

Proposition 3 *Let $s \geq 1$. Then, the current map $[\cdot]: \text{LipImm}(S^1, \mathbb{T}^2) \rightarrow H^{-s}(\Lambda^1(\mathbb{T}^2))$ is Hölder-continuous with exponent $1/2$; in particular, it is continuous.*

Proof It suffices to prove the proposition for $s = 1$, for if two linear functionals are close in H^{-1} , then they are close in H^{-s} for $s > 1$. First, consider the case of immersions that are also embeddings. Let $\phi = (\phi_1, \phi_2)$ be a fixed Lipschitz embedding and let $\psi = (\psi_1, \psi_2)$ be any Lipschitz embedding with $\|\phi_i - \psi_i\|_{H^1(S^1)} < \delta$. Let $\alpha \in H^1(\Lambda^1(\mathbb{T}^2))$. We need to estimate

$$A := |[\phi](\alpha) - [\psi](\alpha)|.$$

² For example, when S^1 acts on $\mathbb{C}^n$ by $z_i \mapsto e^{i\theta} z_i$, the set of invariants $\bar{z}_i z_j$, $1 \leq i, j \leq n$ — n^2 real invariants in all—is complete, and n^2 is much larger than $\dim(\mathbb{C}^n/S^1) = 2n - 1$. One can find smaller complete sets, limited by the dimension of the smallest Euclidean space into which $\mathbb{C}^n/S^1$ can be embedded. However, in such sets the individual invariants are more complicated [4], so that the total complexity is $O(n^2)$.




Fig. 3 Constructions used in the proof of Proposition 3. Here, the red curve is the reference curve ϕ and the nearby (blue) curve ψ is shown with reversed orientation. Left: The line integral $\int_{\phi} \alpha + \int_{\psi} \alpha$ is equal, by Stokes' theorem, to $\int_A \alpha - \int_B \alpha + \int_C \alpha - \int_D \alpha$, so its magnitude is bounded by $\int_{A \cup B \cup C \cup D} |\alpha|$. Middle: When the reference curve (in red) intersects itself in the sense shown, $\int_{\phi} \alpha - \int_{\psi} \alpha = \int_A \alpha - \int_B \alpha +$

$\int_C \alpha - \int_D \alpha$, so $\int_E \alpha$ does not appear, i.e. it has weight $d = 0$. Right: When the reference curve (in red) intersects itself in the other sense, we have $\int_{\phi} \alpha - \int_{\psi} \alpha = -\int_A \alpha - \int_B \alpha - \int_C \alpha - \int_D \alpha - 2 \int_E \alpha$, i.e. $\int_E \alpha$ appears with weight $d = 2$. A similar construction applies if the red curve takes on the same value 3 or more times (Color figure online)

First, we have:

$$A = \left| \int_{\phi(S^1)} \alpha - \int_{\psi(S^1)} \alpha \right|.$$

Let R be the region in $\mathbb{T}^2$ enclosed between $\phi(S^1)$ and $\psi(S^1)$ (see Fig. 3). From Stokes's theorem:

$$A = \left| \int_R d\alpha \right|. \quad (11)$$

In coordinates, $\alpha = \alpha_1 dx + \alpha_2 dy$, which gives $d\alpha = (\partial_x \alpha_2 - \partial_y \alpha_1) dx \wedge dy$. Defining $f = \partial_x \alpha_2 - \partial_y \alpha_1$ we have $d\alpha = f dx \wedge dy$, so that

$$A = \left| \int_R f(x, y) dx \wedge dy \right|. \quad (12)$$

Now, applying the Cauchy–Schwarz inequality with the integrand $1 \cdot f$ we obtain

$$\begin{aligned} A &\leq \|f\|_{L^2} (\text{area}(R))^{\frac{1}{2}} \\ &\leq \|\alpha\|_{H^1} (\text{area}(R))^{\frac{1}{2}} \\ &\leq \|\alpha\|_{H^1} (C\delta)^{\frac{1}{2}}, \end{aligned}$$

where C is a constant depending on ϕ (approximately equal to the length of ϕ). Therefore, for all such ψ we have

$$\begin{aligned} \|[\phi] - [\psi]\|_{H^{-1}(\Lambda^1(\mathbb{T}^2))} \\ = \sup_{\|\alpha\|_{H^1(\Lambda^1(\mathbb{T}^2))}=1} |[\phi](\alpha) - [\psi](\alpha)| \leq (C\delta)^{\frac{1}{2}}, \end{aligned}$$

establishing the claim.

For immersions that are not embeddings, Eq. (11) is modified to take into account any intersections. Let the region R between ϕ and ψ be $R = \cup_{i=0}^n R_i$, where R_0 is the non-overlapping part and $R_1, \dots, R_n$ are the overlapping parts. Then

$$A \leq \sum_{i=0}^n d_i \|f\|_{L^2} \text{area}(R_i)^{\frac{1}{2}}$$

where $d_0 = 1$ and each d_i is either 0 or 2, depending on the orientation of the boundary curves of each R_i (see Fig. 3). Applying the Cauchy–Schwarz inequality with $g = \sum_{i=0}^n d_i \chi(R_i)$ gives

$$A \leq \|\alpha\|_{H^1} \left(\sum_{i=0}^n d_i^2 \text{area}(R_i) \right)^{\frac{1}{2}}.$$

As the number of intersections is fixed by the choice of ϕ , again we have $A \leq (C\delta)^{\frac{1}{2}}$, establishing the claim. $\square$

In order to obtain differentiability results on the current map, we shall work with slightly different spaces of Sobolev 1-forms, depending on whether we are considering immersions or embeddings. For the embeddings $\text{LipEmb}(S^1, \mathbb{T}^2)$, we shall consider the space $H_{\sigma}^s(\Lambda^1(\mathbb{T}^2))$ with $s \geq 2$; if $s = 1$, then Proposition 7 shows that the map $[\cdot]$ is not differentiable. For immersions, we first embed $\mathbb{T}^2$ (considered as a sufficiently large square in $\mathbb{R}^2$ with opposite sides identified) in a cube of similar dimensions in $\mathbb{R}^3$, with opposite ends identified to create the flat 3-torus $\mathbb{T}^3$. If $s \geq 2$, then Propositions 2 and 3 remain true for the current map as a map into $H_{\sigma}^{-s}(\Lambda^1(\mathbb{T}^3))$, with appropriate modifications to the proofs; specifically, embed $\mathbb{T}^2$ in the plane $z = 0$ so that all terms in

the forms α and $d\alpha$ involving dz integrate to zero in the current map. We now consider the differentiability of the current map as a linear functional on $H^s_\sigma(\Lambda^1(\mathbb{T}^3))$ with $s \geq 2$.

Proposition 4 (1) *If $s \geq 2$, then the current map $[\cdot]: \text{LipEmb}(S^1, \mathbb{T}^2) \rightarrow H^{-s}(\Lambda^1(\mathbb{T}^2))$ is differentiable. In particular, let $\psi(t)$ be a C^1 curve in $\text{LipEmb}(S^1, \mathbb{T}^2)$ with $\psi(0) = \phi$ and $\partial_t|_{t=0} = v$; then, for $\alpha \in H^s(\Lambda^1(\mathbb{T}^2))$, the derivative of the current map at ϕ , evaluated on α , is given by*

$$(D[\phi] \cdot v)(\alpha) = \int_{\phi(S^1)} i_v d\alpha \quad (13)$$

and is a bounded linear operator in the dual Sobolev topology.

(2) *Let $s \geq 2$ and consider $\mathbb{T}^2$ as an embedded submanifold of $\mathbb{T}^3$. Then, the current map $[\cdot]: \text{LipImm}(S^1, \mathbb{T}^2) \rightarrow H^{-s}(\Lambda^1(\mathbb{T}^3))$ is differentiable. In particular, let $\psi(t)$ be a C^1 curve in $\text{LipImm}(S^1, \mathbb{T}^2)$ with $\psi(0) = \phi$ and $\partial_t|_{t=0} = v$. For $\alpha \in H^s(\Lambda^1(\mathbb{T}^3))$, the derivative of the current map at ϕ , evaluated on α , is a bounded linear operator in the dual Sobolev topology, whose formula is also given by (13).*

Proof To prove the statement for embeddings, we shall compute the derivative formally then justify the expression. Let $\psi(t)$ be a C^1 curve in the space of Lipschitz embeddings with $\psi(0) = \phi$, $\partial_t|_{t=0} = v \in H^1(\phi(S^1))$ (see Remark 4), and let α be as in the statement of the proposition. Then,

$$\begin{aligned} (D[\phi] \cdot v)(\alpha) &= \partial_t|_{t=0}[\psi(t)](\alpha) \\ &= \partial_t|_{t=0} \int_{S^1} \psi(t)^* \alpha \\ &= \int_{S^1} \partial_t|_{t=0} \psi(t)^* \alpha \\ &= \int_{S^1} \phi^* \mathcal{L}_v \alpha \\ &= \int_{\phi(S^1)} di_v \alpha + i_v d\alpha \\ &= \int_{\phi(S^1)} i_v d\alpha, \end{aligned}$$

where the fourth equality follows from the definition of the Lie derivative, the fifth from Cartan's magic formula (also known as the homotopy formula [3]), and the sixth from an application of Stokes' theorem. Since $\alpha \in H^s(\Lambda^1(\mathbb{T}^2))$ and $s \geq \frac{3}{2}$, it follows from the trace theorem (see Remark 5) that the L^2 norm of $d\alpha|_{\phi(S^1)}$ is bounded. Using Hölder's inequality, we estimate

$$\begin{aligned} |(D[\phi] \cdot v)(\alpha)| &= \left| \int_{\phi(S^1)} i_v d\alpha \right| \\ &\leq C \|d\alpha|_{\phi(S^1)}\|_{L^2} \cdot \|v\|_{L^2(\phi(S^1))} \\ &\leq C \|\alpha\|_{H^s(\Lambda^1(\mathbb{T}^2))} \cdot \|v\|_{H^1(\phi(S^1))}. \end{aligned}$$

Taking the supremum over $\|\alpha\|_{H^s(\mathbb{T}^2)} = 1$, we see that

$$\|D[\phi] \cdot v\|_{H^{-s}(\Lambda^1(\mathbb{T}^2))} \leq C \|v\|_{H^1(\phi(S^1))},$$

which justifies the ansatz and proves that the derivative exists as a bounded linear operator.

If $\psi(t)$ is an immersion, then the range of ψ may be lifted from $\mathbb{T}^2$ to $\mathbb{T}^3$ and perturbed slightly at the crossings so that it becomes an embedding $\tilde{\psi}$ in $\mathbb{T}^3$. The same formula for the derivative now holds in $\mathbb{T}^3$ for $\tilde{\psi}$. Since $\alpha \in H^s(\Lambda^1(\mathbb{T}^3))$ and $s \geq 2$, it follows from the trace theorem (see Remark 5) that the L^2 norm of $d\alpha|_{\tilde{\psi}(S^1)}$ is bounded. This can be seen by first restricting $d\alpha$ to the surface S in $\mathbb{T}^3$ bounded by $\tilde{\psi}$ so that $\frac{1}{2}$ an order of Sobolev differentiability is lost using the trace theorem, then restricting $d\alpha|_S$ to $\tilde{\psi}$ using the trace theorem once more so that a total of 1 order of Sobolev differentiability is lost. Estimating as above and taking the supremum over $\|\alpha\|_{H^s(\Lambda^1(\mathbb{T}^3))} = 1$, we obtain

$$\|D[\tilde{\phi}] \cdot \tilde{v}\|_{H^{-s}(\Lambda^1(\mathbb{T}^3))} \leq C \|\tilde{v}\|_{H^1(\phi(S^1))}.$$

Letting the perturbation tend to zero, it follows that the derivative exists as a bounded linear operator. $\square$

Remark 6 When $s \geq 2$, one can define a continuous Riemannian metric on shapes as the restriction of $H^{-s}(\Lambda^1(\mathbb{T}^3))$ to the currents of immersed shapes. While many families of Riemannian metrics on shapes have been studied [21], this one appears to be new. However, in this paper we do not use the induced Riemannian metric, but rather the ('straight line') subset metric illustrated in Fig. 2, which is far easier to compute. Riemannian properties of currents in this setting are left for future work.

2.4 Representers of Shapes

In this section, we return to the current map on immersions as a map into $H^{-s}(\Lambda^1(\mathbb{T}^2))$ and derive an expression for the representer of a shape; that is, a PDE on differential 1-forms whose solution represents the immersed shape under consideration. In view of Proposition 2, the Riesz representation theorem guarantees the existence of a unique $\beta \in H^s(\Lambda^1(\mathbb{T}^2))$ that represents the current $[\phi]$; we call this β the (Riesz) *representer* of $[\phi]$; it satisfies:

$$[\phi](\alpha) = \langle \alpha, \beta \rangle_{H^s} \quad (14)$$

for all $\alpha \in H^s(\Lambda^1(\mathbb{T}^2))$. We can also write

$$\|[\phi]\|_{H^{-s}(S^1)}^2 := \|\beta\|_{H^s(\Lambda^1(\mathbb{T}^2))}^2 = [\phi](\beta). \quad (15)$$

That is, two shapes are close in H^{-s} if the representers of their currents are close in H^s .

The definition of the representer, Eq. (14), is an elliptic PDE in weak form. We now derive an explicit expression for it. Let $A_{\phi(S^1)}$ be the closure of the region bounded by $\phi(S^1)$ and let $\chi_{A_{\phi(S^1)}}$ be the characteristic function of this closed set. Rewriting the left-hand side of (14) in terms of the L^2 norm and using, in order, Stokes' theorem, the Leibniz rule, and Stokes' theorem once more, we obtain

$$\begin{aligned} [\phi](\alpha) &= \int_{\phi(S^1)} \alpha \\ &= \int_{A_{\phi(S^1)}} d\alpha \\ &= \int_{\mathbb{T}^2} \chi_{A_{\phi(S^1)}} d\alpha \\ &= \int_{\mathbb{T}^2} d(\chi_{A_{\phi(S^1)}} \alpha) - \int_{\mathbb{T}^2} d\chi_{A_{\phi(S^1)}} \wedge \alpha \\ &= \int_{\mathbb{T}^2} \alpha \wedge d\chi_{A_{\phi(S^1)}} \\ &= \int_{\mathbb{T}^2} \alpha \wedge -\star d\chi_{A_{\phi(S^1)}} \\ &= (\alpha, -\star d\chi_{A_{\phi(S^1)}})_{L^2(\mathbb{T}^2)}. \end{aligned}$$

From (14) and the definition of the Sobolev inner product, we obtain

$$\langle \alpha, -\star d\chi_{A_{\phi(S^1)}} \rangle_{L^2(\mathbb{T}^2)} = \langle \alpha, (I + \sigma^{2|s|} \Delta^s) \beta \rangle_{L^2(\mathbb{T}^2)}. \quad (16)$$

Since this holds for any form α , the representer β is the solution of the PDE

$$(I + \sigma^{2|s|} \Delta^s) \beta = -\star d\chi_{A_{\phi(S^1)}}, \quad (17)$$

where the right-hand side is to be interpreted in the sense of distributions. For example, when $s = 1$, β is given by

$$(I + \sigma^2 \Delta) \beta = -\star d\chi_{A_{\phi(S^1)}}.$$

Since $\chi_{A_{\phi(S^1)}}$ belongs to $L^2(\mathbb{T}^2)$, the derivative $d\chi_{A_{\phi(S^1)}}$ belongs to $H^{-1}(\Lambda^1(\mathbb{T}^2))$. It follows from the standard theory on elliptic differential operators (see Chapter 5, Theorem 1.3, of Taylor [26]) that $\beta \in H^1(\Lambda^1(\mathbb{T}^2))$ and satisfies the estimate

$$\|\beta\|_{H^1(\Lambda^1(\mathbb{T}^2))} \leq \|(I + \sigma^2 \Delta) \beta\|_{H^{-1}(\Lambda^1(\mathbb{T}^2))} + \|\beta\|_{L^2(\Lambda^1(\mathbb{T}^2))}.$$

When $s = 2$, we consider $\mathbb{T}^2$ as an embedded submanifold of $\mathbb{T}^3$ and the current map into $H^2(\Lambda^1(\mathbb{T}^3))$ is differentiable. The Riesz representation theorem guarantees the existence of a unique representer β in H^2 . As before, $d\chi_{A_{\phi(S^1)}} \in H^{-1}(\Lambda^1(\mathbb{T}^3))$ but now $\beta \in H^3(\Lambda^1(\mathbb{T}^3)) \subset H^2(\Lambda^1(\mathbb{T}^2))$. Moreover, $H^3(\Lambda^1(\mathbb{T}^3))$ embeds compactly in $H^2(\Lambda^1(\mathbb{T}^2))$

by the Rellich–Kondrachov Lemma so that closed and bounded sets of representers are compact in $H^2(\Lambda^1(\mathbb{T}^2))$.

As remarked earlier, for immersions, the currents do not determine the shape, as parts of the shapes that retrace themselves are invisible to currents. For embeddings, the situation is better:

Proposition 5 *Let $\phi_1, \phi_2: S^1 \rightarrow \Omega$ be two Lipschitz embeddings of S^1 and let $s = 1$. If $[\phi_1] = [\phi_2]$, then the curves represent the same oriented shape.*

Proof Suppose that a point x in $\phi_1(S^1)$ is not in $\phi_2(S^1)$. Then, there is a neighbourhood of x which does not intersect $\phi_2(S^1)$ either. Choose a 1-form α that has support in this neighbourhood. As $\phi_1(S^1)$ is an embedding, the form α can be chosen such that $[\phi_1](\alpha) \neq 0$. However, $[\phi_2](\alpha) = 0$, which gives a contradiction. $\square$

However, currents do determine the shape of immersions if there is some control over the self-intersections:

Proposition 6 *Let $\phi_1, \phi_2: S^1 \rightarrow \Omega$ be two Lipschitz immersions, each with a finite number of self-intersections such that at each self-intersection the tangent vectors are continuous and distinct. Let $s = 1$. If $[\phi_1] = [\phi_2]$, then the curves represent the same oriented shape.*

Proof As in the proof of Proposition 5, the currents determine the images of ϕ_1 and ϕ_2 away from self-intersections. At the self-intersections, the hypothesis on continuity and distinctness of the tangent vectors allows the non-self-intersecting pieces of the shapes to be joined together in a unique way, thus determining the same oriented shape. $\square$

The assumption on the self-intersections is necessary. Even without parts of curves that retrace themselves, self-intersections—either with equal or discontinuous tangent vectors—prevent the current from recognizing the shape up to $\text{Diff}(S^1)$ reparameterizations (see Fig. 4). The current sees only the image of the curve.

Insight into the H^{-s} shape metric is obtained by considering the target domain $\mathbb{R}^2$ and shapes consisting of vertical lines with periodic boundary conditions in y . Proposition 7 illustrates both the non-differentiability of the current map $[\cdot]$ in the $s = 1$ case and the difference between the $s = 1$ and $s = 2$ metrics. The $s = 1$ metric weights nearby portions of the shapes more heavily than the $s = 2$ metric does. It also shows the role of the length scale σ ; roughly, all curves more than Euclidean distance σ apart are an equal distance apart in the H^{-s} metrics.

We first prove an elementary lemma.

Lemma 1 *Suppose that H is a reproducing kernel Hilbert space of functions on $\mathbb{R}$. Let $\delta_x \in H^*$ denote the evaluation at the point $x \in \mathbb{R}$. Suppose further that the kernel is translation*




Fig. 4 Limitations of currents. Currents see only the image of the curve and thus cannot distinguish some shapes. In these shapes, the numbers indicate the order in which the curves are traversed. The shapes in each row have the same currents. In the top row, the curves are smooth, but the curves intersect tangentially, and the current does not determine the

order of traversal. In the bottom row, the current determines the shape only if it is known that the curve is smooth at its self-intersections (bottom left); otherwise, there are four different orderings; that is, four distinct elements of $\text{LipImm}(S^1, \mathbb{R}^2)/\text{Diff}^+(\mathbb{R})$ have the same current

invariant, i.e. the representer for δ_x takes the form $x' \mapsto K(x' - x)$. Then, the distance between δ_0 and δ_ϵ for any $\epsilon \in \mathbb{R}$ is

$$\|\delta_0 - \delta_\epsilon\| = \sqrt{2(K(0) - K(\epsilon))}. \quad (18)$$

Proof In general, if K_x denotes the representer of δ_x , we have

$$\begin{aligned} \|\delta_0 - \delta_\epsilon\|^2 &= (K_0 - K_\epsilon, K_0 - K_\epsilon) \\ &= \langle \delta_0 - \delta_\epsilon, K_0 - K_\epsilon \rangle \\ &= K_0(0) - K_\epsilon(0) - K_0(\epsilon) + K_\epsilon(\epsilon) \\ &= K_0(0) + K_\epsilon(\epsilon) - 2K_0(\epsilon) \end{aligned}$$

Using the translation invariance of the kernel, we have $K_\epsilon(\epsilon) = K_0(0)$, which finishes the proof. $\square$

Assuming periodic boundary conditions in y , the representer of a vertical line at $x = x_0$ is $K((x - x_0)/\sigma)$ where K is the corresponding one-dimensional kernel. We define the *distance per unit length* to be the distance in one dimension, with the corresponding one-dimensional kernel.

Proposition 7 The distance per unit length between two straight lines a (Euclidean) distance ϵ apart is $\mathcal{O}(\epsilon^{1/2})$ when $s = 1$ and $\mathcal{O}(\epsilon)$ when $s = 2$.

Proof Using Lemma 1, the distance by unit length is $\sqrt{2(K(0) - K(\epsilon))}$.

For $s = 1$, $K(x) = \frac{1}{2}e^{-|x|}$ and the distance per unit length between the lines is $(1 - e^{-\epsilon/\sigma})^{1/2} = \mathcal{O}(\epsilon^{1/2})$ as $\epsilon \rightarrow 0$.

For $s = 2$, $K(x) = \frac{1}{4}e^{-|x|}(1 + |x|)$ and the distance per unit length between the lines is $(\frac{1}{2}(1 - e^{-(\epsilon/\sigma)})(1 + \epsilon/\sigma))^{1/2} = \mathcal{O}(\epsilon)$ as $\epsilon \rightarrow 0$. $\square$

3 Discretization of Currents by Finite Elements

The discretization of the representation of shapes by currents proceeds in two steps:

- (i) Discretization of the space of currents and approximation of the currents of shapes; and
- (ii) Discretization of the metric on currents.

We shall consider these separately, as their numerical properties are somewhat independent; item (i) determines how accurately the shapes themselves are represented, while item (ii) determines the geometry of the induced discrete shape space. Although we mostly work with finite elements, the ‘finite’ part is not essential; in principle, any linear space of functions will do. Examples 10–12 use a spectral method, i.e. high-degree polynomials on a single element, while Example 13 uses a space of radial basis functions, the kernel of the metric. The computational advantages and disadvantages of each approach are very similar to the analogous questions in the numerical solution of elliptic PDEs.

3.1 Discretization of Currents

We first return to the general setting in which we work with immersions $\phi: M \rightarrow N$, where M and N are manifolds, and with their currents $[\phi](\alpha) = \int_M \phi^* \alpha$ for $\alpha \in H^s \Lambda^m(N)$. The discretization of currents requires the choice of three things:

- (i) A space V of finite elements on M ;
- (ii) A space W of finite elements on $\Lambda^m(N)$;
- (iii) A method of evaluating or approximating $[\phi_V]_W$, where ϕ_V is the finite element representation of ϕ in V .

We will explore the ability of $[\phi_V]_W$ to represent shapes in different settings analytically and numerically.

To illustrate the ideas, we first consider an extremely simple example, namely sets of n points on a line.

Example 1 Let $M = \{1, \dots, n\}$ and $N = \mathbb{R}$. A shape is then an unordered set of n points in $\mathbb{R}$ with isotropy the group of permutations. In this case, M does not need to be discretized. We choose a uniform mesh with spacing Δx on $\mathbb{R}$ and let W be the piecewise polynomials of degree at most d on the mesh. Let $w(x)dx \in W$. Then

$$[\phi](w(x)dx) = \int_M \phi^*(w(x)dx) = \sum_{j=1}^n w(x_j).$$

The simplest case is $d = 0$, i.e. piecewise constant elements. For these, the currents count how many points are in each cell. Therefore, these elements represent the shape with an accuracy of $\mathcal{O}(\Delta x)$, as there is no way to tell where in each cell the points are located.

The next case is $d = 1$, i.e. piecewise linear elements. The piecewise constants determine how many points are in each cell, and the linear elements determine $\sum_{j: x_j \in W_i} x_j$, that is, the mean location of the points in cell W_i . If there is at most 1 point in each cell, then the representation is perfect, as the points are located exactly. If there is more than 1 point in a cell, then these elements represent the shape with an accuracy of $\mathcal{O}(\Delta x)$. Notice how the use of currents factors out the parameterization of the shape, i.e. it is invariant under permutations of M .

With piecewise elements of degree d , the number of points in each cell and their first d moments are determined; if there are at most d points in each cell, then the points are located exactly. This is because the moment equations on $[0, \Delta x]$ (for example) are $\sum_{j: x_j \in W_i} (x_j)^i = c_i$. For given values of the moments c_i , this set of polynomial equations has total degree $d!$, hence at most $d!$ solutions; but if there is one real solution, then any permutation of the x_j is a solution; hence, the moments determine the points up to ordering.

We now consider our main example of oriented closed planar curves. Let $M = S^1$ and let $N = \Omega$, a domain in $\mathbb{R}^2$. We

will choose V to be the continuous piecewise polynomials of a given degree on a uniform mesh on S^1 and W to be either the discontinuous or the continuous piecewise polynomials of a given degree on a fixed mesh on Ω . We will take the finite element representation ϕ_V of ϕ to be the element of V that interpolates ϕ at the finite element nodes. The current $[\phi_V]_W$ is then given by the integral of a piecewise polynomial function. This can be evaluated exactly; however, in this study we evaluate $[\phi_V]_W$ by quadrature, either the midpoint rule for piecewise linear elements or Simpson's rule for piecewise quadratics. As we are interested in fairly compact representations of shapes, we will pick the mesh size of V to be much smaller than the mesh size of W .

In the following subsections, we study the behaviour of finite element currents with regard to (i) quadrature errors, with and without noise; (ii) accuracy of representation of shapes; and (iii) accuracy of the induced metric on shapes.

3.2 Quadrature Errors and Robustness of Currents

One strong motivation for considering currents, as opposed to other possible shape invariants such as those based on arclength (cf. Proposition 1), is that—essentially because they are signed—currents are expected to be robust against noise and to function well on quite rough shapes. We present three examples measuring different aspects of robustness.

Example 2 In this example, the shapes are rough, but there is no noise. Figure 5 shows the quadrature errors for 3 rough shapes as a function of the number of points used to discretize them. In this and the following example, the H^{-1} shape norm is discretized using polynomials of degree 9 on $[-1, 1]^2$, although only the quadrature error is reported. Shapes (a) and (b) are continuous, but not Lipschitz continuous, and yet the quadrature errors are well under control, and for shape (b) are even $\mathcal{O}((\Delta s)^2)$ (where Δs is the mean mesh spacing).

Example 3 In this example, shown in Fig. 6, the shape is smooth, but different levels of noise are added. Thus, the exact value to which we compare is the zero noise, zero mesh spacing limit. Two different quadratures (the midpoint and Simpson's rule) are compared. Although high levels of noise can dominate the quadrature error, they do not prevent the computation of highly accurate values for the currents until the noise level is actually greater than the mesh spacing Δs .

Example 4 We compare the sensitivity of currents to that of arclength-based currents in Fig. 7. The reference shape is a unit line segment, discretized with uniform mesh spacing Δs , and independent normally distributed noise of mean 0 and standard deviation ϵ is added to the x and y components of each point except the endpoints. The current $\int y dx$ is computed by the midpoint approximation $\sum (y_{i+1} - y_i)(x_{i+1} - x_i)$. For $\epsilon \ll \Delta s$, the errors accumulate like a sum of random




Fig. 5 Example 2: Discretized currents can describe rough shapes robustly. Here, 3 shapes of different smoothness are captured at different resolutions and the quadrature error in their norms (as computed with the midpoint rule) compared. Shape (a) has Fourier coefficients $\hat{\phi}_k = \mathcal{O}(|k|^{-3/2})$; it is in H^p for $p < 1$ and is continuous, but not differentiable. Shape (b) has Fourier coefficients $\hat{\phi}_k = \mathcal{O}(|k|^{-2})$; it is in H^p for $p < 3/2$ and is continuous, but just fails to be differentiable. Shape (c) has Fourier coefficients $\hat{\phi}_k = \mathcal{O}(|k|^{-3})$; it is nearly twice differentiable. Even the very rough shape (a) can be represented with an error of 1% using 256 points

variables, the error being $N(0, \epsilon(\Delta s)^{\frac{1}{2}})$. For $\epsilon \gtrsim \Delta s$, nonlinear effects create a larger error of size $\mathcal{O}(\epsilon(\Delta s)^{-\frac{1}{2}})$. Overall, good results are obtained even with noise levels $\epsilon = \mathcal{O}(\Delta s)$. In contrast, the positive nature of length means that the arclength approximation $\sum((x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2)^{\frac{1}{2}}$ accumulates errors of size $\mathcal{O}((1 + (\epsilon/\Delta s)^2)^{\frac{1}{2}} - 1)$. These are $\mathcal{O}(\epsilon^2/(\Delta s)^2)$ for small ϵ and $\mathcal{O}(\epsilon/\Delta s)$ for larger ϵ . Overall, $\epsilon = o(\Delta s)$ is necessary for acceptable results.

3.3 Accuracy of Representation of Shapes

In this section, we consider simple closed planar shapes and study how accurately they are represented by discontinuous finite elements. Let $M = S^1$ and let $N = \Omega$, a domain




Fig. 6 Example 3: Errors in the shape norm $\|\phi\|_*$ calculated for different resolutions and noise levels. The shape to be measured is smooth and has 13 random nonzero Fourier coefficients. The relative error in its shape is shown as calculated by the midpoint rule (thick lines) and by Simpson's rule (thin lines), as a function of the mean mesh spacing $\Delta s := \|\phi(t_{i+1}) - \phi(t_i)\|$. Independent normally distributed noise with standard deviation ϵ is added to each point. When $\epsilon = 0$, the error is $\mathcal{O}((\Delta s)^2)$ for the midpoint rule and $\mathcal{O}((\Delta s)^4)$ for Simpson's rule, as expected. The shape can be represented with an error of 5% with 16 points and a noise of 0.05 (5%), or with an error of 0.1% with 256 points and noise of 0.1%

in $\mathbb{R}^2$. Let $\phi: S^1 \rightarrow \Omega$ be a smooth embedding. Take a triangular mesh on Ω consisting of triangles $\mathcal{T}$ of maximum diameter Δx . Let W be the space of discontinuous piecewise polynomial finite elements of degree $\leq d$. We are studying the effect of the discretization of the ambient space N , so we do not discretize M ; we assume that all currents are evaluated exactly. We assume that the currents $[\phi]_W$ are given, and we want to know how accurately they determine the shape.

First we consider piecewise constant elements, that is, $d = 0$.

Proposition 8 For smooth embeddings $\phi: S^1 \rightarrow \Omega \subset \mathbb{R}^2$, and a triangular mesh $\mathcal{T}$ of sufficiently small mesh size Δx , the currents of ϕ evaluated on 1-forms constant on each tri-




Fig. 7 Example 4: Comparison of the robustness of currents versus arclength-based invariants. Here, the reference shape is the unit line segment $([0, 1], 0)$ discretized with mesh spacing Δs . The current $\int y \, dx$, the signed area under the curve, is calculated using the midpoint rule, and the arclength is calculated from a piecewise linear approximation. The data x_i and y_i , except the endpoints, are perturbed with many repeated trials of independent Gaussian noise of mean 0 and standard deviation ϵ . The 2-standard-deviation limits of the error in the results are shown. For the current, the errors (dashed lines) are $\mathcal{O}(\epsilon(\Delta s)^{1/2} + \epsilon(\Delta s)^{-1/2})$ so that excellent results are obtained even with $\epsilon \sim \Delta s$; for the arclength (solid lines), the errors are $\mathcal{O}(\epsilon/\Delta s)^2$ for small ϵ , so that $\epsilon = o(\Delta s)^2$ is required for acceptable results

angle $T \in \mathcal{T}$ determine a piecewise linear approximation $\hat{\phi}$ of ϕ of pointwise second-order accuracy.

Proof The currents are

$$\left\{ \left(\int_{\phi(S^1) \cap T} dx, \int_{\phi(S^1) \cap T} dy \right) : T \in \mathcal{T} \right\}.$$

If the mesh is sufficiently fine, then these currents are either zero (if the curve does not intersect T), or they record the jumps in x and y of that part of the curve that lies in T . The elements on which the currents are nonzero therefore determine the set of elements whose interiors intersect the curve ϕ . If the mesh is sufficiently fine and the curve is in general position, then these elements form a ‘discrete topological circle’, a set of triangles each sharing an edge with exactly two others (see Fig. 9).

We now consider finding a continuous shape $\hat{\phi}$ with the same currents as ϕ . The currents have two degrees of freedom




Fig. 8 Any nonzero vector can be placed in any triangle such that the vector’s head and tail lie on the sides of the triangle. If there are two such positions, then they are incident on different pairs of sides. There is one such position when the vector is parallel to one of the sides




Fig. 9 Approximation of a shape by piecewise constant currents (built up from Fig. 1). A smooth curve is shown in red, and a triangular finite element mesh in black. The unique approximating curve that is continuous, is linear on each triangle and interpolates the shape at the element edges is shown in blue. On each triangle, the currents evaluated on the constant 1-forms dx and dy determine the jumps in x and y of the shape across the triangle. These jumps determine (i) if the shape intersects the triangle; and (ii) at most two line segments with the given jumps. From these two, the line segment that meets other active edges is selected, yielding the second-order accurate approximant shown (Color figure online)

per triangle, as do shapes that are linear on each triangle; we therefore seek such an approximant $\hat{\phi}$. For any values of the currents of ϕ on piecewise constants, i.e. for any values of the jumps in x and y , there are at most two line segments in T with endpoints on the edges of T whose currents take on these values, see Fig. 8. If there are two such line segments, then they join different pairs of edges. The line segment that joins two edges that are part of the known discrete topological circle can then be chosen as shown in Fig. 9.

This piecewise linear approximation $\hat{\phi}$ to ϕ interpolates ϕ at the edges of the elements, and, as ϕ is assumed to be smooth, obeys

$$\max_{s \in S^1} \min_{t \in S^1} \|\phi(t) - \hat{\phi}(s)\| = \mathcal{O}((\Delta x)^2)$$

on each triangle T . It therefore determines the shape to second-order accuracy. $\square$

Next we consider the improvement that can be obtained using discontinuous piecewise linear or quadratic elements, i.e. $d = 1, 2$.

Proposition 9 *Let $\mathcal{T}$ be a triangular planar mesh. Let*

$$V_x := \{dx, dy, y dx, xy dx, y^2 dx\},$$

$$V_y := \{dx, dy, x dy, x^2 dy, xy dy\}.$$

Then, for sufficiently smooth $\phi: S^1 \rightarrow \mathbb{R}^2$ in general position and for $\mathcal{T}$ sufficiently fine, the integrals of the first k 1-forms in V_x and V_y , $2 \leq k \leq 5$, over $\phi(S^1) \cap T$ for each $T \in \mathcal{T}$ determine an $\mathcal{O}((\Delta x)^k)$ -accurate approximation of $\phi(S^1)$.

Proof From Proposition 8, the piecewise constant currents determine the occupied triangles. Suppose that the shape can be written on an occupied triangle T in the form $y = g(x)$. If the derivatives of g are bounded, we use the 1-forms in V_x only. Otherwise, the curve can be written in the form $x = f(y)$ where f has bounded derivatives and we use the 1-forms in V_y only. Without loss of generality, we consider the first case.

The case $k = 2$ is covered in Proposition 8.

For $k = 3$, we know the intersection points of the curve with the edges of each occupied triangle and, in addition, the value of the current $\int_U y dx$ for each $U = \phi(S^1) \cap T$, $T \in \mathcal{T}$. We take coordinates in which the shape is $y = g(x)$ and the x -intersections are at $x = 0$ and $x = h$, where $h = \mathcal{O}(\Delta x)$; this fixes the parameterization. We consider the quadratic approximation $y = \hat{g}(x)$ to $g(x)$ that interpolates $g(x)$ at $x = 0$ and at $x = h$ and has the same value of $\int_0^h y dx$. This yields the approximation

$$\hat{g}(x) = g(0) \left(\frac{h-x}{h} \right) + g(h) \frac{x}{h} + a_0 x(h-x)$$

where

$$a_0 = \frac{6}{h^3} \left(\int_0^h g(x) dx - \frac{g(0) + g(h)}{2} h \right).$$

At $x = sh$, $0 \leq s \leq 1$, the approximation error is

$$g(sh) - \hat{g}(sh) = \frac{h^3}{12} s(1-s)(2s-1)g'''(0) + \mathcal{O}(h^4).$$

That is, piecewise linear elements determine the shape to third order accuracy.

For $k = 4$, we know, in addition, the current $\int_U xy dx$, and we choose a cubic approximation. There is a unique such

cubic, and it yields a fourth-order approximation to $g(sh)$ with leading order error

$$\frac{h^4}{120} s(s-1)(5s^2-5s+1)g''''(0).$$

For $k = 5$, we know, in addition, the current $\int_U y^2 dx$, and we seek a degree 4 approximant. The equations for the coefficients of the approximant are now nonlinear, of total degree 2. They have two real solutions. One has leading order error less than $0.00003h^5|g^{(5)}(0)|$, and the other has leading order error less than $0.003h^5|g^{(5)}(0)|$. This establishes fifth-order accuracy for $k = 5$. $\square$

Note that in practice, the integrals of both V_x and V_y would be used, but on most triangles they do not provide independent information.

We briefly discuss several things we learn from this proposition:

1. The currents determine smooth shapes very accurately on sufficiently fine meshes. The errors (less than $0.008h^3|g^{(3)}|$ for $k = 3$, $0.0006h^4|g^{(4)}|$ for $k = 4$, and $0.00003h^5|g^{(5)}|$ for $k = 5$), are less than twice that of Chebyshev interpolation.
2. The results suggest that with order d elements with k degrees of freedom per triangle (so that $k = \mathcal{O}(d^2)$), the order of approximation will be $\mathcal{O}(k)$ rather than $\mathcal{O}(d)$. That is, the higher dimensionality of the target manifold N does not seem to be important.
3. The inherent nonlinearity of the approximation need not be an obstacle. Despite the nonlinearity of the approximation, leading to quadratic equations when $k = 5$, the best approximant can be chosen systematically. Nevertheless, we anticipate that at very high degrees d the nonlinearity may render it difficult to reconstruct an accurate approximation.
4. The situation here is an example of a *moment problem*. If we choose coordinates in which the base of the triangle T is at $y = 0$, and consider the domain:

$$E_i = \{(x, y) : 0 \leq x \leq h, 0 \leq y \leq g(x)\},$$

then we are being given the moments $\iint_{E_i} x^m y^n dx dy$ (for some set of values of m, n) and are asking how accurately we can reconstruct $g(x)$. When $n > 1$, we have a nonlinear approximation problem about which, as far as we know, little is known.

5. The moment problem considered in this section for discontinuous piecewise polynomial finite elements is nonlinear, but at least it is entirely local. We anticipate that analogous results to Proposition 9 hold for the approximation of shapes by the currents of continuous piecewise polynomial finite elements.

The computational cost of the finite element-based method is largely based on the required meshing of the domain and is hence driven by mesh size. For high-resolution meshes, particularly in higher dimensions, this could be an issue. However, we note that for the 2D meshes that we are using, the FEniCS implementation available at [7] runs in times well under a second for meshes up to size 64×64 on a standard PC.

3.4 Discretization of the Metric on Shapes

Recall that the norm of the shape metric is defined by Eqs. (14, 15):

$$\|[\phi]\|_{H^{-s}} := \|\beta\|_{H^s}, \quad \text{where } [\phi](\alpha) = (\beta, \alpha)_{H^s} \quad \forall \alpha, \quad \beta \in H^s \Lambda^1(\Omega).$$

If $W \subset H^s \Lambda^1(\Omega)$, then both $[\phi]$ and the H^s inner product may be restricted to W to yield a finite element approximation of the representer β and a metric on W^* . That is:

$$\|[\phi]\|_{W^*} := \|\beta\|_{H^s}, \quad \text{where } [\phi](\alpha) = (\beta, \alpha)_{H^s} \quad \forall \alpha, \quad \beta \in W.$$

In coordinates, let $w_1, \dots, w_k$ be a basis of W and let $\mathbf{b}, \mathbf{f}$ be the coordinate vectors of β and $[\phi]$, respectively; that is, $f_i = [\phi](w_i)$ and $\beta = \sum_{i=1}^k b_i w_i$. Then

$$G\mathbf{b} = \mathbf{f}, \quad (19)$$

where

$$G_{ij} = (w_i, w_j)_{H^s} \quad (20)$$

is the matrix of the H^s -metric restricted to W . For $s = 1$, G is a linear combination of the mass and stiffness matrices of W . Note that (19), (20) amount to a standard finite element solution of the inhomogeneous Helmholtz equation. Then, we have

$$\|[\phi]\|_{W^*} = \sqrt{\mathbf{b}^T G \mathbf{b}}.$$

In practice, we use the Cholesky decomposition of G to represent β in an orthonormal basis, i.e. we compute $\tilde{\mathbf{b}} = G^{\frac{1}{2}} \mathbf{b}$ so that

$$\|[\phi]\|_{W^*} = \|\tilde{\mathbf{b}}\|_2.$$

In this way, each shape maps to a point in a standard Euclidean space $\mathbb{R}^{\dim W}$ and standard techniques such as principal components analysis can be applied. We think of $\|\cdot\|_{W^*}$ as providing a highly compressed or approximate geometric representation of shape space.




Fig. 10 Illustration of the representer of a shape. The representer (blue vector field) of a randomly chosen shape (red) in $[-1, 1]^2$ with respect to the H^1 metric with length scale $\sigma = 1/\sqrt{10}$. Here, W is the space of polynomials on $[-1, 1]^2$ of degree less than 10, times dx or dy . The colour represents the size of the representer vector (Color figure online)

Although the choice $s = 0$ does not make sense at the continuous level—the ‘representer’ would be a delta function supported on the shape, which is not in $L^2 \Lambda^1(\Omega)$ —it does make sense at the discrete level. An example is provided by the piecewise constant finite elements considered in Sect. 3.3. The representer is nonzero on the triangles that intersect the shape; it is similar to a discrete greyscale drawing of the shape. As the currents of piecewise constant finite elements know the intersection of the shape with the edges of the mesh, they already provide a very sensitive discretization of the metric on shapes. However, it does not converge as $\Delta x \rightarrow 0$. Moreover, in this approximate metric, all shapes that intersect with different edges are seen as equally far away. For example, the computed distance between two straight segments located at $x = 0$ and $x = x^*$ is proportional to

$$\begin{cases} x^*, & x^* < \Delta x, \\ \Delta x, & x^* \geq \Delta x. \end{cases}$$

For these reasons, we do not consider discontinuous elements further.

An example of an $s = 1$ representer is shown in Fig. 10, calculated using a single square element together with polynomial currents of degree less than 10. The Helmholtz equation smears out the shape over a length scale σ ($1/\sqrt{10}$ in this example). This allows shapes to be sensitive to each other’s positions over lengths of order σ , which is typically many times the mesh spacing.

For any s , one possible choice of W is the set of all polynomials up to some chosen degree. This gives a spectral method. However, we do not expect spectral accuracy because the representers are not smooth, and they are only in $H^s \Lambda^1(\Omega)$. As

in finite element solutions of PDEs, the mesh size and order of the elements can be adjusted depending on the application. But unlike PDEs, in most shape applications it is not necessary to have small errors with respect to the continuous problem. Rather, each choice of W provides a different approximation or description of shape space. In some applications, W may be relatively low dimensional.

For $s = 1$, we can take any triangulation of Ω with W the continuous piecewise polynomials of degree $d \geq 1$; these lie in $H^1 \Lambda^1(\Omega)$, so the above construction applies directly. For $s > 1$, we use the same elements and form the same metric G associated with $s = 1$, but determine $\mathbf{b}$ by solving

$$G^s \mathbf{b} = \mathbf{f}. \quad (21)$$

Such representers satisfy higher-order Helmholtz equations with slightly different boundary conditions than those implied by Eqs. (14) and (15). In practice, this difference is immaterial as long as the curve is located away from the boundaries, since the representer tends to zero far from the curve. Moreover, this approximation error is outweighed by the ease of solving (21) and of using standard finite elements.

The H^2 metric we have chosen for the numerical calculations is different from the Sobolev metric (see Definition 1) with $s = 2$. This poses no problems because the norms defined by these two metrics are equivalent in the sense that there exist finite constants C_1 and C_2 such that:

$$C_1 \|\cdot\|_{\text{numerics}} \leq \|\cdot\|_{H^2_\sigma} \leq C_2 \|\cdot\|_{\text{numerics}}.$$

They therefore generate the same topology and the current map retains all the properties given in Propositions 2–7. The metric in Definition 1 was selected for clarity of proof, while this equivalent Sobolev metric was selected for ease of implementation.

4 Examples

We have implemented our approach to currents using finite elements in Python using FEniCS and Dofin [2, 20]. The implementation is available at [7]. We define a rectangular mesh with continuous Galerkin elements of predefined order. To compute the invariants, we use the `tree.compute_entity_collisions()` function to identify which mesh cells intersect with the curve. We then evaluate the basis functions of those cells and update the computation of the currents as the sum of the basis elements of intersecting mesh elements weighted by the size of the cell. Solving for the representer in the relevant norm is based on a matrix solve $\mathbf{G}b_x = f_x$, where $\mathbf{G}$ is the tensor representation of the norm (see Eq. (20)), b_x is a function on the finite element space, and f_x is the vector of x currents

(i.e. an integral of $w_i(x) dx$ over the curve), and similarly for b_y and f_y . We used a length scale of $\sigma = 1/\sqrt{10}$ throughout, so the H^{-1} norm that we used was computed as $m = 1./10 * \text{inner}(\text{grad}(u), \text{grad}(v)) * dx() + u * v * dx()$, where u and v were trial and test functions on the function space, respectively. In order to compute the H^{-2} norm, it is necessary to perform a second solve, based on the matrix built from $x * v * dx()$.

We present a series of experiments demonstrating the use of our implementation on a variety of test cases. The first few examples are chosen to test the robustness of the approach. They demonstrate the reparameterization of the shape, its numerical stability of the method in both norms, and robustness with curves that lose differentiability, for example at corners. These examples are followed by a set of experiments demonstrating that the representation of the curves that are produced is consistent with human perceptions of the shapes.

Example 5 In the first experiment, we demonstrate that the representation in the H^{-1} and H^{-2} norms is invariant to reparameterization of the curves. The curve is a bowtie shape discretized with 512 points. The equally spaced positions of these points were then perturbed by a Gaussian random variable of standard deviation 0.1 (recall that the radius of the circle was 0.5) and the points re-sorted into monotonically increasing order of arclength. Figure 11 shows the representers for the unperturbed shapes on the left and the perturbed ones on the right using piecewise linear finite elements and a mesh size of 10×10 . Visually, there seems to be no difference between them. The relative difference in the computed norms between the original shape and the reparameterized one was of the order of 10^{-4} for both metrics.

Example 6 We examine the convergence of the discretized metric with respect to the mesh size and the order of the finite elements. A circle of radius 0.5 in the domain $\Omega = [-1, 1]^2$ is discretized with 5000 equally spaced points so that quadrature errors are negligible. The grid is a uniform triangular mesh on an $M \times M$ grid. Using elements of order 1 (piecewise linear) up to order 4, we computed the difference in each of the H^1 and H^2 norms as M was increased in powers of 2 from 1 up to 128, i.e. we compute $\|[\phi]\|_{H^{-s}, 2M} - \|[\phi]\|_{H^{-s}, M}$. The results are shown in Fig. 12. Reference lines illustrating convergence of order 1 for H^{-1} and order 2.5 for H^{-2} are also shown. The improved convergence of the H^{-2} metric is clear; the benefit of increasing the order of the finite elements themselves is less clear.

Example 7 In this example, we study the sensitivity of the norms to small wiggly perturbations. The base shape is the same circle of radius 0.5 as in Example 6, and the domain remains $\Omega = [-1, 1]^2$. The circle is perturbed by scaling its radius by a factor $1 + \epsilon \cos(\omega\theta)$; see the illustration in Fig. 13. We take 5000 points on the shapes and compute the




Fig. 11 Example 5: Representer independent of shape parameterization. The representers for a bowtie shape with equally spaced points (left) and under a random reparameterization (right). Visually, there

seems to be no difference between the representers, and the numerical relative differences are around 10^{-4} with both metrics




Fig. 12 Example 6: Convergence of the discrete metric as a function of the mesh size and element order. The log difference between successive approximations of the H^{-1} norm (left) and H^{-2} norm (right) is shown


distances using piecewise linear finite elements and apply Richardson extrapolation to the results from $M = 80$, $M = 160$, and $M = 320$ to ensure that finite element discretization errors are negligible. The results for the two norms are shown in the table in Fig. 13. As expected from the analysis of nearby straight sections in Sect. 3.4, for fixed frequency ω the distances are $\mathcal{O}(\epsilon^{1/2})$ for the H^{-1} metric and $\mathcal{O}(\epsilon)$ for the H^{-2} metric. However, it appears that overall the distances are independent of frequency for H^{-1} and $\mathcal{O}(\omega^{-1/2})$ as $\omega \rightarrow \infty$ for H^{-2} . Overall, the H^{-2} metric is much less sensitive to small wiggles, especially small high-frequency wiggles, than the H^{-1} metric.

Example 8 In this experiment, we consider a family of shapes, the supercircles $x^r + y^r = (1/2)^r$. The shapes, together with their representers, are shown in Fig. 14 with 512 points on the curves, $M = 80$ gridpoints, piecewise linear

as the mesh size increases in powers of 2 for four different orders of finite elements. Reference lines are provided in each plot to illustrate the order of convergence

elements, and $r = 2^1, 2^{1.5}, 2^2, 2^{2.5}$. Note that although the difference between the *curves* is apparent, it is much harder to see a difference between the *representers*. Nevertheless, the norms do see a difference, with the H^{-1} norm decreasing from 4.67 to 3.61 over the four curves shown in the figure. The H^{-2} norm decreases from 1.86 to 1.26.

Example 9 In this example, we study the H^{-1} geometry of the family of supercircles. We take the curves $x^r + y^r = (1/2)^r$ for $0 < r < \infty$. This family ranges from a cross (equivalent to a null shape for currents), through an astroid, a circle, to a square. We compute the shapes for 13 exponents $r = 2^j$, $j = -3, -2.5, \dots, 3$. We use 512 points on each curve, $M = 10$, and piecewise linear elements. Thus, W is a normed vector space of dimension $2 \cdot (M + 1)^2 = 242$. It is transformed to a standard Euclidean $\mathbb{R}^{242}$ using the Cholesky factorization described in Sect. 3.4. The 12 data points in



ω	H^{-1}			H^{-2}		
	$\epsilon = 0.1$	$\epsilon = 0.05$	$\epsilon = 0.025$	$\epsilon = 0.1$	$\epsilon = 0.05$	$\epsilon = 0.025$
2	0.9817	0.6959	0.4868	0.1704	0.0871	0.0440
4	0.9876	0.7017	0.4875	0.1376	0.0709	0.0360
8	0.9905	0.7034	0.4903	0.0994	0.0520	0.0267
16	0.9969	0.7027	0.4868	0.0698	0.0363	0.0189
32	0.9967	0.7037	0.4886	0.0525	0.0256	0.0132
64	0.9991	0.7140	0.4881	0.0450	0.0195	0.0093

Fig. 13 Example 7: Sensitivity of norms to wiggles. The reference shape (red circle) and perturbation (wiggly blue curve) shown for amplitude $\epsilon = 0.05$ and frequency $\omega = 32$; the computed distances between the reference and perturbed shapes are shown at right and tabulated

below in the H^{-1} and H^{-2} norms. The norm perturbation due to the wiggles appears to be $\mathcal{O}(\epsilon^{-1/2})$ for the H^{-1} norm and $\mathcal{O}(\epsilon\omega^{-1/2})$ in the H^{-2} norm (Color figure online)




Fig. 14 Example 8: First look at a family of shapes. The representers for a set of shapes showing a circle deforming to get progressively more 'square-like'. With the chosen scale $\sigma = 1/\sqrt{10}$ on a 1×1 grid, the visual differences between the representers for the various shapes are very small

$\mathbb{R}^{242}$ are projected to $\mathbb{R}^2$ using standard PCA. The results are shown in Fig. 15. The bunching up of the points at either end is clear. The induced geometry of the family can be seen in the curvature of the family.

Example 10 In the previous example, the data set was intrinsically one-dimensional. We now consider a high-dimensional data set. We generate 32 random smooth shapes using random Fourier coefficients and then compare them in the H^{-1} metric using currents $\int x^m y^n dx dy$ for $0 \leq m + n < 10$. The Fourier coefficients $\tilde{z}_k$ of the shapes have

$\tilde{z}_0 = 0$, $\tilde{z}_1 = 0.5$, $\tilde{z}_{-1} = 0$ (so that they are all roughly centred and of the same size), with $\tilde{z}_k$ for $2 \leq k \leq 6$ being independent random normal variables with standard deviation $1/(1 + |k|^3)$. Thus, the data set is ten-dimensional. We then perform an optimization step that computes the planar embedding of the shape currents whose Euclidean distance matrix best approximates the H^{-1} distance matrix, using the best planar subspace of W as computed by PCA as the initial condition for the optimization. The results are shown in Fig. 16. The mean distance error of the embedding is 0.06. This gives a pictorial representation of a high-dimensional




Fig. 15 Example 9: Geometry of a smooth family of shapes. 2D embedding of a set of simple shapes that have a strong dependence on a single parameter, using the first two principal components. The curve of shapes seems to have constant curvature, but we do not know why this is so

data set lying within an infinite-dimensional nonlinear shape space. The triangular shapes appear on the edges, sorted by orientation, while the squarish shapes appear near the centre. Obviously, similar shapes, such as 5 and 16, and 2 and 4, are placed very close together.


Example 11 In order to compare our approach to the differential signature method, we took the same shapes as in Example 10 and compared them modulo the special Euclidean group of the plane by computing their Euclidean differential signatures (κ, κ_s) (where κ is the Euclidean curvature and s is arclength) using the Euclidean-invariant finite

differences introduced by Calabi et al. [10] and corrected by Boutin [9, Eq. (6)]. The signatures are mapped into $[-1, 1]^2$ as $(0.8 \arctan(\kappa/3), 0.8 \arctan(\kappa_s/150))$; this mapping controls the weighting of the extreme features of the signatures at which κ and (especially) κ_s are large. The currents using moments for order < 10 are computed as for Fig. 16, and the currents embedded in the plane to preserve as best as possible their pairwise distances by performing a least-squares optimization in this space to minimize the distance between the pairwise distances between the points in the original space and in this 2D version (Fig. 17).

There are some interesting differences between the two embeddings, such as the tight clustering of 14, 24, 26, 28 in this figure, of which only 24 and 26 are close in Fig. 16. Looking at the position of shape 10 in both figures, it is also clear that the differential signature is dominated by the slight kink on the right of the shape (see also shape 13).

Example 12 Our next example (Fig. 18) is a family of complicated shapes that have a linear dependence on a parameter, $\phi(t; a) = \phi_1(t) + a\phi_2(t)$. As the shapes are relatively complicated and the differences between them relatively small, the currents of the shapes (again, moments of order < 10) embed extremely well in a plane: mean distance error $< 10^{-3}$. However, in this embedding in the plane, significant geometry is seen within the family, with speed along the curve decelerating, changing direction, and accelerating. This example and the previous one illustrate that the method of finite element currents copes with complicated immersions.

Fig. 16 Example 10: Geometry of a set of random shapes. Here, 32 random shapes are created and compared using currents of order $N \leq 10$. The planar embedding of the shape currents that best preserves their pairwise distances is shown. One can see that shapes are grouped together according to their general shapes (triangular or square) and orientation. As the underlying distance is robust with respect to noise, apparently different shapes such as 15, 9 and 13 are grouped together, but their shape (triangular) and orientation are roughly the same






Fig. 17 Example 11 Geometry of a set of random Euclidean shapes. Here, the same 32 shapes as in Fig. 16 are compared modulo the special Euclidean group of the plane (top). The differential signature of each shape is plotted below

Example 13 For some metrics, the Green's function of the elliptic operator appearing in Eqs. (14, 17) can be found explicitly. This allows the direct evaluation of the representer at any point as a (discrete) line integral over the shape. Alternatively, the metric can be *defined* through the choice of a

convenient Green's function, such as a Gaussian, a common choice in other applications of currents in shape analysis. In this example, we illustrate this option.

Let the Green's function (kernel) be K and let Σ be a finite subset of the domain $\mathbb{R}^2$ (e.g. a regular grid). Then, the




Fig. 18 Example 12: Geometry of a smooth family of immersions. A family of shapes is chosen with linear dependence on a parameter. They embed extremely well into the plane, largely forming a curve that reflects the dependence upon the parameter

representer β of a shape $\phi: S^1 \rightarrow \mathbb{R}^2$ is given by (see [17, Lemma 1])

$$\beta(x) = \int K(\phi(s), x) \phi'(s) ds$$

for which the discrete (midpoint rule) approximation

$$\beta_{\Sigma}(x) = \sum_{j=1}^J K\left(\frac{\phi(s_j) + \phi(s_{j+1})}{2}, x\right) (\phi(s_{j+1}) - \phi(s_j))$$

is easily computed for each $x \in \Sigma$. Here, J is the number of points in the discretization of ϕ . In contrast to the finite element case, the representer can be computed without any error due to the discretization of the target domain. To determine the discrete metric, we first solve the linear system

$$\beta_{\Sigma}(x) = \sum_{y \in \Sigma} \gamma_y K(x, y), \quad \forall x \in \Sigma \quad (22)$$

for the coefficients γ_y . Noting that the (spatially continuous) inner product

$$(K(x, y), K(x, z))_{H^s} = K(y, z),$$

we have the discrete inner product

$$(\beta_{\Sigma}, \beta_{\Sigma}) = \sum_{y, z \in \Sigma} \gamma_y \gamma_z K(y, z).$$

That is, the matrix of the inner product in terms of β_{Σ} is $(K_{yz})^{-1}$. This approximates the operator of the elliptic PDE Eq. (14), (17).




Fig. 19 Example 13: Convergence of discrete metric calculated using Green's functions. The log difference between successive approximations of the H^{-2} norm for the shape shown in Fig. 11 is shown as the mesh size increases in powers of 2. The reference line has slope -4 . The norm is calculated from a representer evaluated exactly on a regular mesh using Green's functions. Compare with the right-hand plot of Fig. 12

We have applied this construction to Example 6 for the bowtie shape of Example 5. We have used an H^2 metric with operator $(1 - \sigma^2 \nabla^2)^2$, target domain $\mathbb{R}^2$, and Green's function $K(x, y) = \frac{1}{4\pi\sigma^3} \|x - y\| K_1(\|x - y\|/\sigma)$. We examine the convergence of the shape norm for Σ an $M \times M$ regular mesh in $[-1, 1]^2$ with $M = 2, 4, 8, 16, 32, 64$ (Fig. 19). The norm converges like M^{-4} , which is similar to a high-order finite element method. The errors are due to the mesh size and to boundary effects.

Our implementation of this example is limited by its use of a direct solve for the dense $|\Sigma| \times |\Sigma|$ system (22). This is slower than the solution of the sparse systems arising in finite elements, but it could be accelerated by multipole methods. Nor can it be used for the Gaussian kernel $K(x, y) = e^{-\|x - y\|^2}$, which (because it corresponds to an H^∞ metric) is ill-posed and requires regularization to invert. A possible alternative is to retain the representer β_{Σ} as the discrete shape—which is, after all, an element of a linear space as we demanded—but to adopt a different metric. However, this would take us outside the framework of this paper.

In some shape analyses, such as the classification of a set of shapes using PCA without shifting or scaling, it is possible to calculate the required inner products between pairs of shapes directly from the shapes themselves [17, Sec. 3.1], as implemented, e.g. in Deformetrica [8]. The trade-offs involved in the information that this provides, as the number and resolution of the shapes vary, require a separate study beyond the scope of this paper.

Example 14 In our final example, we investigate how currents can be used to separate populations of shapes. We

Fig. 20 Example 14: Separating populations of shapes: data set. Data drawn from three slightly different sets, where the difference is in the third and fourth Fourier components. It is hard to visually distinguish the three sets





Fig. 21 Example 14: Separating populations of shapes: results. The 2D embedding of the data set using the first two principal components. The first principal component completely fails to separate the data, but the second does it perfectly

generate a set of shapes using six Fourier components as previously, but we choose different values for the third and fourth Fourier components and add relatively small noise to them. The resulting shapes thus come from three different closely related sets. We generate a small data set comprising ten examples from each of these sets. These are shown in Fig. 20 in the three sets; it can be seen that it is hard to distinguish between the examples.

We then apply PCA to the (piecewise linear) currents and use PCA to project the shapes into 2D. Figure 21 shows the results of this PCA with a few examples of each shape plotted. The three classes are clearly visible in this embedding.


References

- Adams, R., Fournier, J.: Sobolev Spaces, Pure and Applied Mathematics. Elsevier Science, Amsterdam (2003)
- Alnaes, M., Blechta, J., Hake, J., Johansson, A., Kehlet, B., Logg, A., Richardson, C., Ring, J., Rognes, M.E., Wells, G.: The FEniCS project version 1.5. *Arch. Numer. Softw.* **3**, 9–23 (2015)
- Arnold, V.: *Mathematical Methods of Classical Mechanics*, 2nd edn. Springer, Berlin (1989)
- Bandeira, A.S., Cahill, J., Mixon, D.G., Nelson, A.A.: Saving phase: injectivity and stability for phase retrieval. *Appl. Comput. Harmon. Anal.* **37**, 106–125 (2014)
- Bauer, M., Bruveris, M., Marsland, S., Michor, P.W.: Constructing reparameterization invariant metrics on spaces of plane curves. *Differ. Geom. Appl.* **34**, 139–165 (2014)
- Beg, M.F., Miller, M.I., Trounev, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vis.* **61**, 139–157 (2005)
- Benn, J., Marsland, S., McLachlan, R.I., Modin, K., Verdier, O.: femshape: library for computing shape invariants of planar curves using the finite element method (2018). <https://github.com/olivierverdier/femshape>
- Bône, A., Louis, M., Martin, B., Durrleman, S.: Deformetrica 4: an open-source software for statistical shape analysis. In: *International Workshop on Shape in Medical Imaging*, pp. 3–13. Springer (2018)
- Boutin, M.: Numerically invariant signature curves. *Int. J. Comput. Vis.* **40**, 235–248 (2000)
- Calabi, E., Olver, P.J., Shakiban, C., Tannenbaum, A., Haker, S.: Differential and numerically invariant signature curves applied to object recognition. *Int. J. Comput. Vis.* **26**, 107–135 (1998)
- Celledoni, E., Eslitzbichler, M., Schmeding, A.: Shape analysis on Lie groups with applications in computer animation. *J. Geom. Mech. (JGM)* **8**, 273–304 (2016)
- Charon, N., Trounev, A.: The varifold representation of non-oriented shapes for diffeomorphic registration. *SIAM J. Imaging Sci.* **6**, 2547–2580 (2013)
- Charon, N., Trounev, A.: Functional currents: a new mathematical tool to model and analyse functional shapes. *J. Math. Imaging Vis.* **48**, 413–431 (2014)
- De Rham, G.: *Variétés différentiables: formes, courants, formes harmoniques*, vol. 3, Editions Hermann (1973)
- Durrleman, S., Pennec, X., Trounev, A., Ayache, N.: Statistical models of sets of curves and surfaces based on currents. *Med. Image Anal.* **13**, 793–808 (2009)
- Glaunès, J., Joshi, S.: Template estimation from unlabeled point set data and surfaces for computational anatomy. In: *Proceedings of the International Workshop on the Mathematical Foundations of Computational Anatomy (MFCA-2006)*, pp. 29–39 (2006)
- Glaunès, J., Qiu, A., Miller, M.I., Younes, L.: Large deformation diffeomorphic metric curve mapping. *Int. J. Comput. Vis.* **80**, 317–336 (2008)
- Hebey, E.: *Nonlinear Analysis on Manifolds: Sobolev Spaces and Inequalities*. American Mathematical Society, Providence (1999)
- Kaltenmark, I., Charlier, B., Charon, N.: A general framework for curve and surface comparison and registration with oriented varifolds. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3346–3355 (2017)
- Logg, A., Wells, G., Hake, J.: DOLFIN: a C++/python finite element library. In: Logg, A., Mardal, K.A., Wells, G. (eds.) *Automated Solution of Differential Equations by the Finite Element*

Method. Lecture Notes in Computational Science and Engineering, vol. 84. Springer, Berlin, Heidelberg (2012)

21. Michor, P.W., Mumford, D.: An overview of the Riemannian metrics on spaces of curves using the Hamiltonian approach. *Appl. Comput. Harmon. Anal.* **23**, 74–113 (2007)
22. Morgan, F.: *Geometric Measure Theory: A Beginner's Guide*, 4th edn. Academic Press, Boston (2009)
23. Mumford, D.: Pattern theory: the mathematics of perception. *Proc. Int. Congr. Math.* **III**, 1–21 (2002)
24. Roussillon, P., Glaunès, J.A.: Surface matching using normal cycles. In: *International Conference on Geometric Science of Information*, pp. 73–80. Springer (2017)
25. Schmudgen, K.: *Unbounded Self-Adjoint Operators on Hilbert Space*. Springer, Berlin (2012)
26. Taylor, M.: *Partial Differential Equations I: Basic Theory*. Springer, Berlin (2011)
27. Vaillant, M., Glaunès, J.: Surface matching via currents. *Inf. Process. Med. Imaging* **19**, 381–92 (2005)

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



James Benn has a Ph.D. from the University of Notre Dame and a masters and honours degrees from Massey University. He is interested in analysis and geometry.



Stephen Marsland is Professor of Mathematics and Data Science at Victoria University of Wellington, New Zealand. He studied at Oxford University and the University of Manchester, graduating with a Ph.D. in 2002. He moved to Massey University, New Zealand, in 2004, and then to Victoria University of Wellington in 2018. He is interested in geometry and its applications, particularly machine learning, signal processing (birdsong analysis), and image analysis.



Robert I. McLachlan FRSNZ is Distinguished Professor in Applied Mathematics at Massey University, New Zealand. He studied at Caltech with Herb Keller and worked at the University of Colorado, Boulder, and ETH Zurich, before joining Massey University in 1994. He is interested in applications of geometry to all areas of applied mathematics, including geometric numerical integration, PDEs, dynamical systems, and image processing.



Klas Modin is a mathematician and Associate Professor at Chalmers University of Technology and University of Gothenburg in Sweden. He obtained his Ph.D. from Lund University in 2010. After that, he did two post-docs, first at Massey University in New Zealand and then at the University of Toronto in Canada. He was recruited to Chalmers in 2013. Modin's research interests lie in the borderland between geometry and computational mathematics.



Olivier Verdier is associate professor of mathematics in the Information and Communication Technology group in the Department of Computing, Mathematics and Physics of the Western Norway University of Applied Sciences. He obtained his Ph.D. from Lund University in 2009. His main interests are Bayesian statistics and machine learning; functional programming, certified programming and type theory; and scientific computing and geometric integration.