

THESIS FOR THE DEGREE OF LICENTIATE OF ENGINEERING

Geometric discretization in shape analysis

Erik Jansson



CHALMERS
UNIVERSITY OF TECHNOLOGY



UNIVERSITY OF GOTHENBURG

Department of Mathematical Sciences
Division of Applied Mathematics and Statistics
Chalmers University of Technology and the University of Gothenburg
Göteborg, Sweden 2022

Geometric discretization in shape analysis

Erik Jansson

©Erik Jansson, 2022

Department of Mathematical Sciences

Division of Applied Mathematics and Statistics

Chalmers University of Technology and the University of Gothenburg

SE-412 96 Göteborg

Sweden

Telephone: +46 (0)31-772 1000

Typeset with L^AT_EX

Printed by Chalmers Reproservice

Göteborg, Sweden 2022

Geometric discretization in shape analysis

Erik Jansson

Department of Mathematical Sciences

Division of Applied Mathematics and Statistics

Chalmers University of Technology and the University of Gothenburg

Abstract

Discretizations in shape analysis is the main theme of this licentiate thesis, which comprises two papers. The first paper considers the problem of finding a parameterized time-dependent vector field that warps an initial set of points to a target set of points. The parametrization introduces a restriction on the number of available vector fields. It is shown that this changes the geometric setting of the matching problem and equations of motion in this new setting are derived. Computational algorithms are provided, together with numerical examples that emphasize the practical importance of regularization. Further, the modified problem is shown to have connections with residual neural networks, meaning that it is possible to study neural networks in terms of shape analysis. The second paper concerns a class of spherical partial differential equations, commonly found in mathematical physics, that describe the evolution of a time-dependent vector field. The flow of the vector field generates a diffeomorphism, for which a discretization method based on quantization theory is derived. The discretization method is geometric in the sense that it preserves the underlying Lie–Poisson structure of the original equations. Numerical examples are provided and potential use cases of the discretization method are discussed, ranging from compressible flows to shape matching.

Keywords: Shape analysis, diffeomorphisms, machine learning, residual neural networks, compressible fluids, quantization.

Contents

Abstract	i
List of Publications	iv
Acknowledgements	v
1 Preliminaries	1
1.1 Background	1
1.2 Differential geometry	2
1.3 Geometric mechanics	11
2 Shape Analysis	17
2.1 Diffeomorphic shape matching	17
2.2 The geometry of shape matching	18
3 Summary of Included Papers	21
3.1 Paper 1	21
3.2 Paper 2	22
Bibliography	23

List of publications

The following papers are included in this thesis:

- Paper 1:** Jansson, E. & Modin, K. (2022). Sub-Riemannian landmark matching and its interpretation as residual neural networks. arXiv:2204.09351
- Paper 2:** Jansson, E. & Modin, K. (2022). Geometric discretization of diffeomorphisms. arXiv:2210.02328

Author contribution:

- Paper 1:** I have developed most of the theoretical framework and drafted the manuscript under supervision and support of Klas Modin, with whom the final text was produced. I have implemented `Julia` code producing the plots and simulation results presented in the paper.
- Paper 2:** I have performed the necessary calculations to verify the properties of the extension considered in the paper. In close collaboration with Klas Modin, I have developed the remaining theory and drafted the first manuscript. We have together finalized the paper. I have produced the plots and the simulation results, of which some builds on code previously written by Klas Modin.

Acknowledgements

Let's keep things simple! I thank in particular my supervisor Klas Modin and my co-supervisor Annika Lang for their support during these years. Further, I would like to thank all coworkers, coauthors, family and friends.

Erik Jansson
Göteborg, 2022

Chapter 1: Preliminaries

1.1 Background

Shape analysis is the mathematical study of shape, which immediately begs the question: What does a mathematician mean by the word shape? While many of us have an intuitive understanding of what it is—something like a circle, square or dodecahedron—in shape analysis a shape is anything a so-called diffeomorphism can act on. It can be for instance a function, a density or a curve.

The central problem of this thesis is the shape matching problem. It is about finding the optimal way to *warp* a source shape to a target shape. Here, optimal means that the diffeomorphism should be the solution to a minimization problem—it should be the least complicated warp that gets the source shape the closest to the target shape while distorting the ambient space as little as possible. In Figure 1.1, we see a schematic illustration of the idea behind the matching problem.




Figure 1.1: The idea behind shape matching: move from A to B in the optimal way.

The study of how one shape can be deformed into another has proven to be an important problem with applications in for instance medical image analysis [Bistoquet et al., 2008, Bruveris and Holm, 2013, Ceritoglu et al., 2013, Risser et al., 2013, Qiu et al., 2009]. As a more concrete example, Bruveris and Holm [2013] describes how Qiu et al. [2009] uses shape analysis to compare MRI images of a number of patients with a common “template brain”. It was then seen that certain

changes in the shape of the brain are associated with Alzheimer’s disease and mild cognitive decline, an intermediate stage between normal ageing and Alzheimer’s disease.

A possible lens through which to view shape analysis is that of geometric mechanics, which allows for a “fluid interpretation” of the shape matching problem. By means of the Riemannian geometry of the manifold of diffeomorphisms, the matching problem can be reduced to a PDE called the EPDiff equation [Bruveris and Holm, 2013]. Many equations in hydrodynamics and mechanics, for instance the rigid body equations, Euler’s equation or the KdV equation, can be obtained using the same type of reduction used to arrive at the EPDiff equation [Arnold and Khesin, 1998]. In fact, there are connections between shape analysis and many areas of mathematics, computer science and physics: numerical analysis of partial differential equations [Azencot et al., 2018, Larsson et al., 2016], hydrodynamics [Mumford and Michor, 2013], optimal transport [Feydy et al., 2017], information theory [Bauer et al., 2015] and even artificial intelligence [Vialard et al., 2020, Younes et al., 2020].

In this thesis, two of the problems arising from connections between shape analysis and other areas of mathematics are investigated. The thesis begins with an introduction to some background and notation necessary for describing shape analysis and the methods used in the papers. For a more complete description of the preliminaries, the interested reader is referred to the sources on which the first chapter builds, Arnold [1978], Lee [2012, 2018], Marsden and Ratiu [1999], da Silva [2008].

1.2 Differential geometry

Our approach to shape analysis is to use *geometric mechanics*. As its name suggests, this approach to mechanics builds on the language of geometry. In this section we therefore introduce several concepts needed to do geometric mechanics. The idea is to start with some basics of differential geometry, such as manifolds of different kinds, differential forms and Lie groups.

The basic geometric concept necessary is that of a *manifold*. Intuitively, an n -dimensional manifold M is a space locally resembling $\mathbb{R}^n$. Another useful picture is to think of a manifold as a generalization of a surface embedded in $\mathbb{R}^{n+1}$. More formally, a *topological manifold* M of dimension n is a completely separable Hausdorff topological space that has the property that every point $x \in M$ has a neighbourhood U homeomorphic to an open subset of the n -dimensional Euclidean space. Denoting the homeomorphism by ϕ , the pair (U, ϕ) is called a *chart*. A point x is said to be in the chart if $x \in U$. An *atlas* is a family of charts $(U_i, \phi_i)_{i=1}^m$ such that every point of M is in at least one of the charts. An atlas is smooth if for every pair of charts (U_i, ϕ_i) and (U_j, ϕ_j) it holds that $\phi_i(U_i \cap U_j)$ and $\phi_j(U_i \cap U_j)$

are open sets in $\mathbb{R}^n$ and $\phi_j \circ \phi_i^{-1} : \phi_i(U_i \cap U_j) \rightarrow \phi_j(U_i \cap U_j)$ is smooth in the sense that each of its components is a smooth real-valued function. A smooth manifold is a topological manifold equipped with a smooth atlas that cannot be contained in a larger smooth atlas. Charts allow us to define *local coordinates*. If $x \in (U, \phi)$, then the set of local coordinates for x is $(x_1, x_2, \dots, x_n) = \phi(x)$. The concept of charts is illustrated in Figure 1.2.




Figure 1.2: Two charts on a smooth manifold M . Figure inspired by Lee [2012, Figure 1.6].

Having introduced manifolds, we now consider some important objects associated to each manifold.

1. The set of vectors v tangent to M at the point $x \in M$ is a vector space known as the *tangent space at x* and is denoted $T_x M$. Equivalently, $T_x M$ is the set of all *derivations* of smooth functions $C^\infty(M)$ at x , i.e., the linear maps $v : C^\infty(M) \rightarrow \mathbb{R}$ satisfying

$$v(fg) = f(x)v(g) + g(x)v(f)$$

for $f, g \in C^\infty(M)$.

2. The *tangent bundle* of the manifold, denoted by TM , is the disjoint union of the tangent spaces $T_x M$ at each $x \in M$. The tangent bundle is a $2n$ -dimensional manifold [Lee, 2012, Proposition 3.18]. There is a natural projection $p : TM \rightarrow M$ sending $v \in T_x M$ to x .
3. The *cotangent bundle* T^*M is the disjoint union of the duals $T_x^* M$ of the tangent spaces.

4. Let M and N be smooth manifolds. The space of isomorphisms of manifolds, or, the space of *diffeomorphisms* of M to N is denoted $\text{Diff}(M, N)$. If $M = N$, the notation $\text{Diff}(M)$ is used. Under composition, $\text{Diff}(M, N)$ is a group. Further, a diffeomorphism φ *acts* on a smooth function $f \in C^\infty(M)$ by $\varphi \cdot f = f \circ \varphi^{-1}$ [Lee, 2012, Proposition 2.15].

Example 1. The circle $\mathbb{S}^1 = \{x \in \mathbb{R}^2, \|x\|_{\mathbb{R}^2} = 1\}$ is a smooth manifold of dimension 1. It locally resembles $\mathbb{R}$. Its tangent space at x is $T_x\mathbb{S}^1 = \mathbb{R}$ and $T\mathbb{S}^1 \cong \mathbb{S}^1 \times \mathbb{R}$. Geometrically, $T\mathbb{S}^1$ is exactly a cylinder, see Figure 1.3a.

Whenever the tangent bundle is a copy of the manifold times the space it locally resembles, as is the case for $\mathbb{S}^1$, the tangent bundle is *trivial* and the manifold is said to be *parallelizable*.




Figure 1.3: Illustration of the tangent space and tangent bundle of the circle. In Figure 1.3a, note how the tangent space at each point $x \in \mathbb{S}^1$ is equal to the real line. By rotating the real lines attached to each point of the circle, one obtains a cylinder of infinite height.

A smooth vector field v on M is informally an assignment of a tangent vector to each point $x \in M$. This is made precise in the definition below.

Definition 1. A vector field is a section of TM , meaning that it is a continuous map $X : M \rightarrow TM$ with the property that for all $x \in M$,

$$p(X(x)) = x.$$

If the map from M to TM is smooth, then X is a smooth vector field. Equivalently, one can understand vector fields as the linear maps on the space of smooth functions $C^\infty(M)$ that are derivations, i.e.,

$$X(fg) = gX(f) + fX(g),$$

for all $f, g \in C^\infty(M)$. The space of all smooth vector fields on M is denoted $\mathfrak{X}(M)$. Two vector fields X and Y can be combined into a third by the *commutator* $[X, Y] \in \mathfrak{X}(M)$. For any $f \in C^\infty(M)$, it is given by

$$[X, Y](f) = X(Y(f)) - Y(X(f)).$$

When working with vector fields, we make use of two important concepts: the pushforward of a vector field by a smooth map, used to move a vector field from one manifold to another, and the integral curve of a vector field, used to move points on a manifold.

A smooth map $F : M \rightarrow N$ between two smooth manifolds M and N pushes a vector field $X \in \mathfrak{X}(M)$ forward by

$$F_*X(x) = dF_{F^{-1}(x)}(X(F^{-1}(x))) \in \mathfrak{X}(N),$$

where $x \in N$ and $dF_{F^{-1}(x)}$ is the differential of F at $F^{-1}(x) \in M$, i.e., the linear map $dF_{F^{-1}(x)} : T_{F^{-1}(x)}M \rightarrow T_xN$ defined by

$$dF_{F^{-1}(x)}(X)(f) = X(f \circ F) \tag{1.1}$$

for any $f \in C^\infty(N)$ and any $X \in T_{F^{-1}(x)}M$.

Equipped with the definition of differentials, we can define *submanifolds*. An m -dimensional *immersed submanifold* $\overline{M}$ of M is a manifold that is a subset of M , and for which the inclusion map $\iota : \overline{M} \rightarrow M$ is an immersion, meaning that the differential of ι is injective everywhere. Further, $\overline{M}$ is said to be an *embedded submanifold* if in addition ι is a topological embedding.

An *integral curve* $\gamma(t)$ of a vector field X is a differentiable curve on M satisfying

$$\frac{d}{dt}\gamma = \dot{\gamma} = X \circ \gamma$$

for all $t \in I := \text{Dom}(\gamma) \subset \mathbb{R}$.

Manifolds can be equipped with structures, of which we here consider two types. First, Riemannian manifolds are reviewed in some detail, after which we move on to symplectic manifolds.

A *Riemannian manifold* (M, g) is a manifold M equipped with a *Riemannian metric* g . This means that at each $x \in M$ the metric is an inner product on T_xM . The metric at $x \in M$ is denoted $g_x(\cdot, \cdot)$. All smooth manifolds admit a Riemannian metric [Lee, 2012, Proposition 13.3]. On a Riemannian manifold one can define for instance distances, angles, and curvature. The length of a vector $v \in T_xM$ is given by

$$|v|_g = \sqrt{g_x(v, v)}.$$

Further, if $\gamma : [a, b] \rightarrow M$ is a smooth curve, the length of γ is given by

$$L_g(\gamma) = \int_a^b |\dot{\gamma}(t)|_g dt.$$

The distance between two points $x, y \in M$ is the infimum of the length of all curves starting in x and ending in y . Finally, a geodesic is the curve minimizing the *energy functional* [Lee, 2018, Chapter 6]

$$E_g(\gamma) = \int_a^b |\dot{\gamma}(t)|_g^2 dt. \quad (1.2)$$

Example 2. The sphere $\mathbb{S}^2 = \{x \in \mathbb{R}^3, \|x\|_{\mathbb{R}^3} = 1\}$ is a smooth manifold. Note that $T_x \mathbb{S}^2 = \{v \in \mathbb{R}^3 : (v, x)_{\mathbb{R}^3} = 0\}$. This description of the tangent space relies on the embedding into $\mathbb{R}^3$. Equivalently, one can use local coordinates $(\theta, \phi) \in [0, 2\pi) \times [0, \pi]$. If tangent vectors are seen as derivations, $T_x \mathbb{S}^2$ is the span of $\frac{\partial}{\partial \theta}$ and $\frac{\partial}{\partial \phi}$. Therefore, one can specify a Riemannian metric on $\mathbb{S}^2$ by giving its values for the pairwise inner products of the basis vectors of the tangent space. On the sphere, the standard example is the *round metric*

$$[g] = \begin{bmatrix} 1 & 0 \\ 0 & \sin^2 \theta \end{bmatrix}.$$

Equipped with this metric, $\mathbb{S}^2$ is a Riemannian manifold. In fact, the round metric on $\mathbb{S}^2$ is the *induced metric* determined by the usual immersion of $\mathbb{S}^2$ into $\mathbb{R}^3$. See Figure 1.4 for an illustration.

In order to do calculus on manifolds we need *differential forms*. A differential k -form is informally an object that measures the (oriented) volume of k -dimensional parallelepiped. More formally, a k -form is a field of alternating multilinear forms, i.e., for each $x \in M$, there is a multilinear form $\alpha_x : (T_x M)^k \rightarrow \mathbb{R}$ that changes sign whenever two arguments are exchanged. The space of k -forms on M is denoted $\Omega^k(M)$.

Two differential forms can be combined with the *wedge product*, $\wedge : \Omega^k(M) \times \Omega^l(M) \rightarrow \Omega^{k+l}(M)$. The wedge product is associative and bilinear. It is commutative if kl is even, and anticommutative if kl is odd.

The *interior product* $\iota_X \alpha$ of a differential k -form α with a vector field X is the $(k-1)$ -form $\alpha(X, \cdot)$. The interior product is also called the *contraction* of α with X .

The exterior derivative $d\alpha$ of a k -form α is a $(k+1)$ -form. The exterior derivative is the $\mathbb{R}$ -linear mapping from $\Omega^k(M)$ to $\Omega^{k+1}(M)$ that satisfies $d^2 = 0$ and $d(\alpha \wedge \beta) = d\alpha \wedge \beta + (-1)^k \alpha \wedge d\beta$, where $\alpha \in \Omega^k(M)$ and $\beta \in \Omega^l(M)$. For 0-forms, which are functions, d coincides with the differential. A differential form $\alpha \in \Omega^k(M)$ is said to be *exact* if there is a $\beta \in \Omega^{k-1}(M)$ such that $\alpha = d\beta$. It is *closed* if $d\alpha = 0$.




Figure 1.4: The sphere, with the tangent plane $T_x \mathbb{S}^2$ at $x \in \mathbb{S}^2$. The geodesic between x and $y \in \mathbb{S}^2$ is depicted in blue. Note how the geodesic, as expected, follows the surface.

The *Lie derivative* of a differential form α with respect to a vector field X can be defined using *Cartan's magic formula*,

$$L_X \alpha := d(\iota_X \alpha) + \iota_X d\alpha.$$

The exterior derivative d allows us to define equivalents of many tools from vector calculus on Riemannian manifolds, for instance, *the gradient* of a function.

Definition 2. Let (M, g) be a Riemannian manifold. The gradient of a smooth function f is given by the vector field ∇f satisfying

$$df = g(\cdot, \nabla f).$$

And, or maybe the, key feature of differential forms is that they can be integrated. On $\mathbb{R}^k$, a k -form ω is given by $\omega = \phi(x) dx_1 \wedge \dots \wedge dx_k$, for some smooth function ϕ . The integral over a bounded convex polyhedron $D \subset \mathbb{R}^k$ is

$$\int_D \omega := \int_D \phi(x) dx_1 \dots dx_k,$$

in the usual sense, familiar from multivariate calculus.

In order to integrate $\omega \in \Omega^k(M)$, where M is n -dimensional, consider the k -cell $\sigma = (D, F, O)$, where D is a bounded convex polyhedron in $\mathbb{R}^k$, $F : D \rightarrow M$ is a differentiable map and O is an orientation of $\mathbb{R}^k$. The integral of ω over σ is

$$\int_\sigma \omega = \int_D F^* \omega,$$

where $F^*\omega(X_1, \dots, X_k) := \omega(F_*X_1, \dots, F_*X_k)$ is the *pullback* of ω . Given a sequence of $r \in \mathbb{N}$ k -cells $(\sigma_i)_{i=1}^r = (D_i, F_i, O_i)_{i=1}^r$ and a sequence of integers $(m_i)_{i=1}^r$, a k -chain c_k on M is the formal sum $c_k = \sum_{i=1}^r m_i \sigma_i$, and the integral of ω over c_k is

$$\int_{c_k} \omega = \sum_{i=1}^r m_i \int_{\sigma_i} \omega.$$

Integrating a k -form over a k -dimensional submanifold K of M is done in two steps. First, K is triangulated with bounded convex polyhedrons, whereafter one integrates over the k -chain defined by the triangulation.

Example 3. As an example, let us integrate the 1-form $\omega = xdx - ydy$ on $\mathbb{R}^2$ over the submanifold K depicted in Figure 1.5. The manifold consists of two parts, so we will need two 1-cells in order to integrate over it. The polyhedron covering the circular arch is just the interval $D_1 = [0, \pi/2]$, with $\phi_1^{(t)} = (\cos(t), \sin(t))$. The line segment is covered by the interval $D_2 = [0, 1]$ with $\phi_2^{(t)} = (t, 0)$. Thus,

$$\begin{aligned} \int_K \omega &= \int_{\sigma_1 + \sigma_2} xdx - ydy = \int_{\sigma_1} xdx - ydy + \int_{\sigma_2} xdx - ydy \\ &= \int_0^{\pi/2} -\sin(2t) dt + \int_0^1 t dt = -\frac{1}{2}, \end{aligned}$$

as $dF = \omega$, with $F = x^2/2 - y^2/2$, ω is exact. Further, the boundary of K is $\partial K = [(0, 1)] - [(0, 0)]$ and

$$\int_{\partial K} F = F(-1, 0) - F(0, 0) = -1/2.$$

The integral of ω over K is equal to the integral of F over ∂K . In fact, this is an example of *Stoke's theorem* [Fortney, 2018, Theorem 11.1], which states that given a $(k-1)$ -form η and a k -dimensional manifold M ,

$$\int_M d\eta = \int_{\partial M} \eta.$$

A Riemannian metric defines a linear map between tangent and cotangent spaces, by means of the interior product. If this map is an isomorphism, the metric is said to be *strong*, but if it is only injective, the metric is *weak*.

Having introduced Riemannian manifolds, we move on to an essential tool in geometric mechanics, *symplectic manifolds*. They can be thought of as generalizations of phase spaces of mechanical systems. A smooth manifold M is said to be symplectic if it can be equipped with a closed, non-degenerate 2-form ω . This means that $d\omega = 0$ and that for all $x \in M$, if $\omega_x(u, v) = 0$ for all $v \in T_x M$, then $u = 0$. A symplectic manifold is even-dimensional [Lee, 2012, Proposition 22.7]. For instance $\mathbb{R}^3$ cannot be symplectic. However, this is in fact indicative of why




Figure 1.5: The submanifold K used in Example 3. It consists of a quarter circle of radius 1 starting in $(1,0)$ and a line segment from the origin to $(1,0)$. As the submanifold consists of two distinct parts, we will need two 1-cells to integrate over it.

symplectic manifolds are generalizations of phase spaces. In a phase space, one doubles the number of variables by including momentum variables—so the number of variables and thus the dimension is even.

While symplectic manifolds may appear similar to Riemannian manifolds, there are important differences. Firstly, not all smooth manifolds (of even dimension) admit a symplectic form. For instance, it is only in 2 dimensions that the sphere can be equipped with a symplectic form [Lee, 2012, Chapter 22]. More importantly, on a symplectic manifold (M, ω) there are coordinates such that ω coincides with the canonical symplectic form on $\mathbb{R}^n$. This means that, in contrast with Riemannian manifolds, there is no local structure; from the symplectic point of view, locally everything is just as $\mathbb{R}^n$ [Koszul and Zou, 2019, Theorem 2.3.7].

Example 4. The sphere, $\mathbb{S}^2$, is a symplectic manifold when equipped with the closed and non-degenerate form

$$\omega_x(u, v) = x \cdot (u \times v),$$

where $x \in \mathbb{S}^2$, $u, v \in T_x \mathbb{S}^2$, $\cdot$ is the standard Euclidean inner product and $\times$ is the cross product. In cylindrical local coordinates, $(\theta, z) \in [0, 2\pi] \times [-1, 1]$, we have that $\omega = d\theta \wedge dz$.

Example 5. If V is a real vector space of dimension n , then

$$T^*V = V \times V^*,$$

where V^* is the dual of V . Just as V is an n -dimensional manifold, T^*V is a manifold of dimension $2n$. Moreover, $V \times V^*$ can be equipped with a *canonical symplectic form*, given by

$$\omega((v_1, \alpha_1), (v_2, \alpha_2)) = \alpha_2(v_1) - \alpha_1(v_2), \quad (1.3)$$

where $v_1, v_2 \in V$ and $\alpha_1, \alpha_2 \in V^*$. In fact, the cotangent bundle of any smooth manifold can be equipped with a canonical symplectic form [Lee, 2012, Proposition 22.11].

A natural question to ask is if a manifold can be both Riemannian and symplectic simultaneously, in the sense that the symplectic form fits together with the Riemannian metric. We have seen such an example already: the sphere. The round metric in Example 2 and the symplectic form on $\mathbb{S}^2$ are *compatible*, meaning that there is an integrable *almost-complex structure* J connecting the Riemannian metric and the symplectic form. In other words, there is a smooth field J of automorphisms of $T\mathbb{S}^2$ such that $J^2 = -I$, where I denotes the identity, and

$$g(u, v) = \omega(Ju, v),$$

for all $u, v \in T\mathbb{S}^2$. A manifold of this type is called a *Kähler* manifold. In addition to $\mathbb{S}^2$, another example of a Kähler manifold is $\mathbb{C}^n$ [da Silva, 2008, Chapter 16.4].

We are now almost ready to proceed to geometric mechanics. However, an important concept appearing in mechanics is that of continuous transformations acting on systems, such as rotations and translations. Properties that are preserved under continuous transformations are called continuous symmetries. The mathematical tool describing these are *Lie groups*. A Lie group H is a smooth manifold that also is a group with the property that the map $(h_1, h_2) \mapsto h_1 h_2^{-1}$ is smooth for all $h_1, h_2 \in H$.

The *Lie algebra* $\mathfrak{h}$ of a Lie group H is defined as the tangent space of H at the identity $e \in H$. It is equipped with a *Lie bracket*, which is a map $[\cdot, \cdot]: \mathfrak{h} \times \mathfrak{h} \rightarrow \mathfrak{h}$.

A Lie bracket is bilinear over scalars, anticommutative and should satisfy the *Jacobi identity*

$$[h_1, [h_2, h_3]] + [h_2, [h_3, h_1]] + [h_3, [h_1, h_2]] = 0,$$

for all $h_1, h_2, h_3 \in \mathfrak{h}$. A *Lie subalgebra* is a linear subspace of $\mathfrak{h}$ that is closed under the Lie bracket. Note that different Lie groups can have the same algebra.

Example 6. A classical example of a Lie group is *the general linear group* $\mathrm{GL}(n)$. It consists of all invertible $n \times n$ matrices with real entries equipped with standard matrix multiplication. Its Lie algebra is denoted by $\mathfrak{gl}(n)$ and is the space of all $n \times n$ matrices. The Lie bracket on $\mathfrak{gl}(n)$ is the standard matrix commutator, $[A, B] = AB - BA$. The general linear group contains several subgroups. One example is $\mathrm{O}(n)$, containing orthogonal matrices of dimension n . Another example is $\mathrm{SO}(n)$, consisting of elements of $\mathrm{O}(n)$ with determinant 1. Note that while $\mathrm{SO}(n) \subset \mathrm{O}(n)$, $\mathfrak{so}(n) = \mathfrak{o}(n) = \{A \in \mathfrak{gl}(n), A^T = -A\}$.

A Lie group H can act on itself. The element $h_1 \in H$ can for instance act on an element $h_2 \in H$ by left translation, $L_{h_1} h_2 = h_1 h_2$, right translation $R_{h_1} h_2 = h_2 h_1$, and by conjugation $C_{h_1} h_2 = L_{h_1} R_{h_1^{-1}} h_2 = h_1 h_2 h_1^{-1}$. The conjugation action gives rise to the *adjoint representation*.

Definition 3. The adjoint representation at $h \in H$, denoted by $\text{Ad}_h \in \text{End}(\mathfrak{h})$, is given by the differential of C_h at the identity. Here, $\text{End}(\mathfrak{h})$ refers to the space of endomorphisms of $\mathfrak{h}$. By varying h we obtain a map $\text{Ad} : H \rightarrow \text{End}(\mathfrak{h})$. Differentiating Ad yields the map $\text{ad} : \mathfrak{h} \rightarrow \text{End}(\mathfrak{h})$. Explicitly, $\mathfrak{h} \ni \xi \mapsto [\xi, \cdot] = \text{ad}_\xi$.

These definitions can be rather opaque, and an enlightening example is found in matrix Lie groups.

Example 7. Let us again consider $\text{GL}(n)$. Let $G, H \in \text{GL}(n)$. The conjugation action is given by $C_G H = G H G^{-1}$. In order to explicitly compute $\text{Ad}_G X$, where $X \in \mathfrak{gl}(n)$, we take the curve in $\text{GL}(n)$ given by $\gamma(t) = I + tX$, $t \in \mathbb{R}$. Note that at $t = 0$, $\gamma(0) = I$. Then,

$$d(C_G)_I(X) = \left. \frac{d}{dt} \right|_{t=0} C_G(I + tX) = \left. \frac{d}{dt} \right|_{t=0} G(I + tX)G^{-1} = GXG^{-1} = \text{Ad}_G X.$$

In order to calculate $\text{ad}_Y X$, with $Y \in \mathfrak{gl}(n)$, take the curve $I + tY$ and differentiate,

$$\text{ad}_Y X = \left. \frac{d}{dt} \right|_{t=0} (I + tY)X(I + tY)^{-1} = YX - XY = [Y, X].$$

As we shall see later, the dual of Ad , the *coadjoint representation*, is central to geometric mechanics and in particular to our approach to shape analysis.

1.3 Geometric mechanics

In this section, we introduce necessary concepts from mechanics, of which there are two main flairs, Lagrangian and Hamiltonian. In both cases, one starts with a manifold, or in the language of physics, a *configuration space*, often denoted by Q . Having introduced geometric mechanics in general, we then will move on to the most central type of system for our purposes, Euler–Arnold systems. These are, as we shall see, are very useful in shape analysis.

Lagrangian mechanics studies systems where a state is specified by a position and a velocity. Geometrically, this means that Lagrangian systems are defined on TQ . A system is described by its *Lagrangian* $\mathcal{L}$, which is a smooth function from TQ to the reals. Given a curve $\gamma : [a, b] \rightarrow Q$, the *action functional* is

$$A(\gamma) = \int_a^b \mathcal{L}(\gamma, \dot{\gamma}). \quad (1.4)$$

The curve γ is a *motion* of the system if it minimizes the action functional. This is known as *Hamilton's principle*. Note here the similarity between geodesics, that extremize the energy functional of Equation (1.2), and motions of a system that extremize the more general action functional. As we shall see later, an interesting class of mechanical systems can indeed be understood as geodesics on a suitable Riemannian manifold.

From the fundamental lemma of calculus of variations and Hamilton's principle, one can deduce that a motion of a Lagrangian system follows the *Euler–Lagrange equations*,

$$\frac{d}{dt} \frac{\partial \mathcal{L}}{\partial \dot{\gamma}} - \frac{\partial \mathcal{L}}{\partial \gamma} = 0. \quad (1.5)$$

Hence, Lagrangian mechanics builds on a *variational principle* [Marsden and Ratiu, 1999, Chapter 1].

Example 8. A basic example of a Lagrangian system is a particle of unit mass in $\mathbb{R}^n$. Its position is given by $x \in \mathbb{R}^n$ and its velocity by $\dot{x}$. The particle is affected by some potential $V : Q \rightarrow \mathbb{R}$. The Lagrangian is the difference of kinetic and potential energy,

$$\mathcal{L}(x, \dot{x}) = \frac{\dot{x} \cdot \dot{x}}{2} - V(x).$$

By Equation (1.5), the equation of motion is $\ddot{x} + \nabla V(x) = 0$, which exactly is Newton's second law.

Having briefly introduced Lagrangian mechanics, we can move on to Hamiltonian mechanics, where one works not with positions and velocities, but with positions and *momenta*. In terms of geometry, we work not in TQ , but in T^*Q . Recall from Example 5 that T^*Q is a symplectic manifold.

Given a smooth function $\mathcal{H}$ on a symplectic manifold, it's *Hamiltonian vector field* $X_{\mathcal{H}}$ is defined by

$$d\mathcal{H} = \omega(\cdot, X_{\mathcal{H}}), \quad (1.6)$$

the function $\mathcal{H}$ is called the *Hamiltonian*. Often, the Hamiltonian describes the total energy of the system, which in the Lagrangian case is given by $E = \frac{\partial \mathcal{L}}{\partial \dot{\gamma}} \cdot \dot{\gamma} - \mathcal{L}$. In fact, with the momenta $p = \frac{\partial \mathcal{L}}{\partial \dot{\gamma}}$, we have that the energy, and thus the Hamiltonian, is given by

$$\mathcal{H} = p \cdot \dot{\gamma} - \mathcal{L}.$$

This is called the *Legendre transformation* and describes the relation between Lagrangian and Hamiltonian mechanics.

The governing equations of a Hamiltonian system is given by

$$\dot{z} = X_{\mathcal{H}}(z). \quad (1.7)$$

where $z = (\gamma, p)$ [Marsden and Ratiu, 1999, Chapter 5.4]. Note that on a Kähler manifold, such as the sphere, the Hamiltonian vector field is given by $X_{\mathcal{H}} = J^{-1} \nabla \mathcal{H}$, i.e., by rotating the gradient 90 degrees.

Example 9. Let us revisit Example 8. There, the particle was described by its position and velocity. The cotangent bundle of $\mathbb{R}^3$ is identified with $\mathbb{R}^3 \times \mathbb{R}^3$. The symplectic form is given by Equation (1.3). The momenta is given by

$$p = \frac{\partial \mathcal{L}}{\partial \dot{x}} = \dot{x}.$$

By the Legendre transform, $\mathcal{H} = p^2 - \mathcal{L} = p^2/2 + V(x)$ and by rotating the gradient of $\mathcal{H}$, one obtains that $X_{\mathcal{H}} = J^{-1}(\nabla V, p) = (p, -\nabla V)$. The Hamiltonian dynamics are then

$$\begin{aligned}\dot{x} &= p, \\ \dot{p} &= -\nabla V.\end{aligned}$$

The symplectic form gives rise to a *Poisson bracket* $\{\cdot, \cdot\} : C^\infty(M) \times C^\infty(M) \rightarrow C^\infty(M)$ by

$$\{H, F\} = \omega(X_F, X_H).$$

A Poisson bracket is a Lie bracket that also satisfies the *Leibniz rule*,

$$\{\omega_1 \omega_2, \omega_3\} = \{\omega_1, \omega_2\} \omega_3 + \omega_1 \{\omega_2, \omega_3\}$$

for all $\omega_1, \omega_2, \omega_3 \in C^\infty(M)$.

Hamilton's equations (1.7) can be formulated using the Poisson bracket,

$$\dot{z} = \{z, H\}.$$

This formulation of Hamiltonian mechanics allows us to easily define for instance *conserved quantities*. A smooth function f is a conserved quantity if it is in *involution* with H , meaning that $\{H, f\} = 0$ [Marsden and Ratiu, 1999, Chapter 5.5].

Before moving on to shape analysis, we describe how geodesics of left-invariant metrics on Lie groups are a useful tool for the study of a large class of systems. As a warm-up, let us consider the *rigid body*.

Example 10. A classical example of a mechanical system is the rotating rigid body. We assume that the rigid body consists of a continuous body B of uniform density. The configuration of the rigid body is described by its rotation, i.e., by an element in $Q = \text{SO}(3)$. The body is assumed not to be influenced by any external torques, so the Lagrangian is just the kinetic energy of the body. The rotation of the body is given by $A \in \text{SO}(3)$. Then, at $x \in B$, the velocity is given by $\frac{d}{dt}Ax = \dot{A}x$ and the infinitesimal kinetic energy is $\dot{A}x \, dx$, so the Lagrangian is

$$\mathcal{L}(A, \dot{A}) = \frac{1}{2} \int_B \|\dot{A}x\|_{\mathbb{R}^3}^2 \, dx. \quad (1.8)$$

Firstly, note that $\mathcal{L}$ is quadratic and positive definite, so it is a metric on $\mathrm{SO}(3)$. Further, if we let $X \in \mathrm{SO}(3)$ act on A by left translation, $\mathcal{L}$ remains unchanged as X is an orthogonal matrix. This means that $\mathcal{L}$ is an example of a *left-invariant metric*. In fact, this invariance is central to the understanding of the behaviour of the rigid body.

By using left-invariance and that $\mathrm{SO}(3)$ has a trivial tangent bundle, the problem can be moved to $T_I \mathrm{SO}(3) \cong \mathfrak{so}(3)$ and the analysis can be continued to obtain the classical equations of motions for a rigid body [Arnold, 1978, Appendix 2]. The important observation for us, however, is that the Lagrangian is a left-invariant Riemannian metric, so the motion of the rigid body *is* a geodesic. The rigid body is an example of an *Euler–Arnold* system, a type of system that can be used to study a large class of mechanical systems, including those that arise in shape analysis [Arnold, 1966, Bruveris and Holm, 2013].

In what follows, let G be a Lie group and denote by $\mathfrak{g}$ its Lie algebra. A Riemannian metric on G is left-invariant if its value is unchanged by left translations. A left-invariant metric can be wholly specified by an inner product on $\mathfrak{g}$, as left translations can be used to move the metric to the correct tangent space. This inner product can be specified with a positive-definite, symmetric linear operator $A : \mathfrak{g} \rightarrow \mathfrak{g}^*$, by $(u, v) = Av(u)$.

The key fact of the Euler–Arnold framework is the following theorem.

Theorem 1. *Let the Lagrangian $\mathcal{L}$ be given by a left-invariant Riemannian metric defined by the operator $A : \mathfrak{g} \rightarrow \mathfrak{g}^*$. Then, the curve $\gamma : [0, 1] \rightarrow G$ extremizing the functional*

$$\int_0^1 \mathcal{L}(\gamma, \dot{\gamma}) dt$$

satisfies the Euler–Arnold equation

$$\begin{aligned} \dot{m} - \mathrm{ad}_v^* m &= 0, \\ m &= Av \end{aligned} \tag{1.9}$$

where $v \in \mathfrak{g}$. The curve γ is a geodesic on G .

For a proof, see Modin [2019].

There is a Hamiltonian interpretation of the Euler–Arnold framework. In order to do Hamiltonian mechanics, we shall specify an appropriate symplectic form.

Definition 4. Consider the coadjoint representation $\mathrm{Ad}^* : G \rightarrow \mathrm{End}(\mathfrak{g}^*)$. The *coadjoint orbit* through $\xi \in \mathfrak{g}^*$ is given by

$$\mathcal{O}_\xi = \{\mathrm{Ad}_g^* \xi, g \in G\}.$$

It is possible to show that $T_\eta \mathcal{O}_\xi = \{\text{ad}_g^* \eta, g \in \mathfrak{g}\}$. The coadjoint orbits can be endowed with a natural symplectic structure. If $\eta \in \mathcal{O}_\xi$, the *Kirillov form* Ξ_η is given by

$$\Xi_\eta(\text{ad}_{g_1}^* \eta, \text{ad}_{g_2}^* \eta) = \eta([g_1, g_2]).$$

Two important facts are that Euler–Arnold systems evolve on the coadjoint orbits and that they are Hamiltonian systems. The Hamiltonian is given by $\mathcal{H}(\eta) = \frac{1}{2}\eta(A^{-1}\eta)$, where $\eta \in \mathcal{O}_\xi$. Further, the coadjoint orbits “slices up” the dual algebra. More precisely, the set of coadjoint orbits *foliate* $\mathfrak{g}^*$ [Kirillov, 2004]. Again, the rigid body provides an illustrating example.

Example 11. The configuration space of the rotating rigid body is $G = \text{SO}(3)$. Thus $\mathfrak{g} = \mathfrak{so}(3) = \mathfrak{g}^*$. The identification of $\mathfrak{g}$ with its dual is performed by equipping $\mathfrak{so}(3)$ with the Frobenius inner product

$$\langle \omega_1, \omega_2 \rangle_F = \text{Tr}(\omega_1 \omega_2^T).$$

With $A \in \text{SO}(3)$ and $\xi, \eta \in \mathfrak{so}(3)$ arbitrary, the coadjoint action on $\mathfrak{so}(3)$ is found by

$$\langle \text{Ad}_A^* \xi, \eta \rangle_F = \langle \xi, \text{Ad}_A \eta \rangle_F = \langle \xi, A \eta A^{-1} \rangle_F = \langle A^{-1} \xi A, \eta \rangle_F.$$

Thus, the orbits are of the form

$$\mathcal{O}_\xi = \{A^{-1} \xi A, A \in \text{SO}(3)\}.$$

It is possible to write the matrix ξ as a vector $\hat{\xi}$ in $\mathbb{R}^3$, and it holds that $\widehat{A^{-1} \xi A} = A^{-1} \hat{\xi}$. The coadjoint action of A is thus the ordinary action of A^{-1} on vectors. As rotation matrices preserve distance, the vector $\hat{\xi}$ can only be transported to vectors of the same length. The coadjoint orbit $\mathcal{O}_\xi$ is therefore the sphere with radius $\|\hat{\xi}\|_{\mathbb{R}^3}$. This gives rise to a foliation of the space by concentric spherical shells. In order to obtain an intuitive understanding, it is constructive to picture this as resembling an onion, which also consists of concentric layers, see Figure 1.6.

¹Image source: Bell, Darwin, 2007. One a day food item #26 - January 2006 [Online]. Available from: https://upload.wikimedia.org/wikipedia/commons/3/35/Red_onions_%28cross-sections%29.jpg [Accessed 21 September 2022].



Figure 1.6: Some sliced onions. Notice how onions consist (approximately) of concentric spherical shells, thus being similar to the coadjoint orbits of the rigid body.¹

Chapter 2: Shape Analysis

In this chapter we briefly introduce shape analysis. We first describe the matching problem, and then describe the geometric structure of the problem, in particular how geometric mechanics is used to obtain an understanding and reduction of the matching problem to a differential equation, using the Euler–Arnold framework described in Section 1.3. For a more detailed description on shape matching, see Bruveris and Holm [2013], Joshi and Miller [2000], Younes [2010].

2.1 Diffeomorphic shape matching

In this section, we let (M, g) be an orientable compact Riemannian manifold.

We consider a class of weak right-invariant Riemannian metrics on $\text{Diff}(M)$ given by

$$\langle \dot{\gamma}, \dot{\gamma} \rangle_{\gamma} = \int_M v \cdot Lv dx, \quad (2.1)$$

where $\gamma(t)$ is a smooth curve taking values in $\text{Diff}(M)$, $L: \mathfrak{X}(M) \rightarrow \mathfrak{X}^*(M)$ is an invertible elliptic differential operator acting on vector fields and $v = \dot{\gamma} \circ \gamma^{-1}$. One example of such an operator $L = (1 - \Delta)^k$ for some $k \in \mathbb{N} \cup \{0\}$.

In the following, let us model shapes as smooth functions. The *shape matching problem* in the large deformation diffeomorphic metric mapping (LDDMM) setting matches two shapes f_0 and f_1 , which are represented by smooth functions, i.e., as elements of $C^\infty(M)$. The time-dependent vector field $v: [0, 1] \rightarrow \mathfrak{X}(M)$ resulting in the best matching is the one minimizing the energy functional

$$E(v) = \|f_0 \circ \gamma(1)^{-1} - f_1\|^2 + \frac{1}{2\sigma} \int_0^1 \int_M v(t) \cdot Lv(t) dx dt, \quad (2.2)$$

where γ is obtained by the differential equation $\dot{\gamma}(t) = v(t) \circ \gamma(t)$, $\gamma(0) = \text{Id}$.

Let us dissect the right-hand side in Equation (2.2). The first term is a *matching term*, measuring how close the target image and the warped source are. The second term is a regularization term, penalizing deformations that are too extreme. The parameter σ , which must be positive, is a parameter that tunes the influence of

the regularization parameter. The second term can be thought of as either the energy of the deformation, or as a measure of how different γ is from the identity.

One approach to spatially discretize the matching problems is to use *landmarks*. The f_0 and f_1 are approximated by points $x_1, \dots, x_m \in M$ that discretizes f_0 , and $z_1, \dots, z_m \in M$ that discretizes f_1 . These points are known as landmarks. Since M is a Riemannian manifold, the metric induces a distance function $d_M: M \times M \rightarrow \mathbb{R}$. The landmark matching problem is to find the time-dependent vector field v that solves

$$\min_v \sum_{i=1}^m d_M^2(y_i(1), z_i) + \frac{1}{2\sigma} \int_0^1 \int_M v \cdot Lv dx dt, \quad (2.3)$$

$$\text{s.t. } \dot{y}_i(t) = v(t, y_i(t)), \quad t \in [0, 1], \quad y(0) = x_i. \quad (2.4)$$

Just as for Equation (2.2), the first term in Equation (2.3) is a matching term, and the second is a regularization term.

Minimization in the landmark matching problem is over all smooth vector field. The optimal vector field should warp ambient space as little as possible as this results in lower energy. Therefore, the optimal momentum should only have support on the landmark paths $y_1(t), y_2(t), \dots, y_m(t)$, meaning that

$$Lv_t = \sum_{i=1}^n p_i \delta_{y_i}(t),$$

for some variables $p_1, \dots, p_n$, and the optimization problem can be reduced to a finite-dimensional Hamiltonian system, which is why moving to landmarks results in a spatial discretization of the vector fields [Joshi and Miller, 2000].

Remark 1. Note that in both Equation (2.2) and Equation (2.3), the matching terms are explicitly written down. As we shall see later, the geometric structure of the matching problem, which leads to its reduction to a differential equation, is determined only by the regularization term. Therefore, this treatment holds also for more general matching terms.

2.2 The geometry of shape matching

In this section we will describe the geometric structure of the shape matching problem. Firstly, note that $\text{Diff}(M)$ can be given the structure of a Fréchet manifold. It is an infinite-dimensional manifold, which in contrast with finite-dimensional manifolds resembles not a Euclidean space but a *Fréchet space*—a particularly well-behaved topological vector space. We omit further discussion, but remark that the group multiplication (composition) in $\text{Diff}(M)$ is smooth with respect to the Fréchet structure. Thus, $\text{Diff}(M)$ is a Lie group whose Lie algebra is given

by the smooth vector fields. As we shall see, the matching problem will give rise to an Euler–Arnold equation on $\mathfrak{X}^*(M)$. In other words, matching problems have a dynamic formulation. The initial image is warped to match the target via a geodesic path of diffeomorphisms.

The matching term in Equation (2.2) only depends on $\gamma(1)$. A curve being a motion to the Lagrangian system governed by Equation (2.2) must therefore also be a motion in the Lagrangian system determined only by the regularization term,

$$\mathcal{L}(\gamma, \dot{\gamma}) = \frac{1}{2} \int_0^1 \int_M v(t) \cdot Lv(t) \, dx dt. \quad (2.5)$$

Equation (2.5) is right-invariant by construction. Therefore, the Euler–Arnold framework applies, and we conclude that the momentum $m = Lv$ evolves according to the equation

$$\begin{aligned} \dot{m} &= \text{ad}_v^* m, \\ Lv &= m. \end{aligned}$$

The group $G = \text{Diff}(M)$ is known, so we can compute ad^* explicitly. To this end, let $w \in \mathfrak{X}(M)$ be arbitrary and note that

$$\text{ad}_v^* m(w) = m(\text{ad}_v w).$$

By Definition 3, $\text{ad}_v w = [v, w]$. By the properties of the so-called *Levi-Cevita connection*, we have that $[v, w] = \nabla_v w - \nabla_w v$ [Lee, 2018, Chapter 5]. Finally, a calculation in local coordinates shows that

$$\text{ad}_v^* m(w) = m(\nabla_v w - \nabla_w v) = (\nabla_v^T m - \nabla_m v + \text{div}(v)m)(w).$$

The Euler–Arnold equations on $\text{Diff}(M)$ equipped with the metric given by the integrand in Equation (2.5) are

$$\begin{aligned} \dot{m} &= \nabla_v^T m - \nabla_m v + \text{div}(v)m, \\ Lv &= m. \end{aligned} \quad (2.6)$$

Equation (2.6) is known as the *EPDiff equation*.

We thus have a reduction of the matching problem to a set of differential equations. This reduction also suggests some numerical algorithms for the computing of the matching. For more details on computational aspects of matching problems, see for instance Beg et al. [2005], Joshi and Miller [2000].

The geometric structure of LDDMM is illustrated in Figure 2.1. The curve of diffeomorphisms $\gamma: [0, 1] \rightarrow \text{Diff}(M)$ is a geodesic in $\text{Diff}(M)$, which descends to a flow on the *orbit* of f_0 , given by

$$\text{Orb}(f_0) := \{f_0 \circ \psi^{-1} : \psi \in \text{Diff}(M)\} \subset C^\infty(M).$$




Figure 2.1: Illustration of the geometric structure of LDDMM. The red line in $\text{Diff}(M)$ is a geodesic, and moves the initial shape $f_0 \in C^\infty(M)$ along its orbit to get f_0 as close as possible to the target shape f_1 .

The goal of the minimization is to find the element in the orbit as close to f_1 as possible. Typically, $f_1 \notin \text{Orb}(f_0)$, so the matching will not be exact.

Furthermore, consider the set

$$\text{Diff}_{f_0}(M) := \{\psi \in \text{Diff}(M) : f_0 \circ \psi^{-1} = f_0\} \subset \text{Diff}(M).$$

In other words, it consists of the diffeomorphisms that leaves f_0 unchanged. Note that $\text{Diff}_{f_0}(M)$ is a subgroup of $\text{Diff}(M)$. Indeed, if $\psi, \varphi \in \text{Diff}_{f_0}(M)$, then $f_0 \circ (\psi \circ \varphi)^{-1} = f_0 \circ \varphi^{-1} \circ \psi^{-1} = f_0$, and $f_0 \circ \psi = f_0 \circ \psi^{-1} \circ \psi = f_0$, thus, $\text{Diff}_{f_0}(M)$ is closed under composition and inversion and is a subgroup. The orbit can be represented by the quotient set $\text{Diff}(M)/\text{Diff}_{f_0}(M)$, so points of the orbit can be thought of as fibers in $\text{Diff}(M)$.

Chapter 3: Summary of Included Papers

3.1 Paper 1

In paper 1, we consider landmark matching modified in two ways. The first modification is that the vector field warping the initial landmarks is parametrized in the sense that it is determined by a set of parameters $\mathcal{U}$. The second modification, inspired by Öktem et al. [2017], is that the matching term includes the case where the target landmarks $z_1, \dots, z_m$ are in a metric space N that may be different from the original manifold M . The main reason for introducing these modifications is that they allow us to connect deep learning and landmark matching. This is possible by viewing residual neural networks as temporal discretizations of time-continuous control problems, as in Celledoni et al. [2021], Li et al. [2017], and by describing how the control problems can be interpreted as high dimensional landmark matching problems.

The main result of the paper lies in describing the geometric structure of the modified problem. The parametrization of vector fields determines a subset of $\mathfrak{X}(M)$ which is not a Lie subalgebra of $\mathfrak{X}(M)$. The matching problem is considered as a nonlinear control system where the space of controls are given by $\mathcal{U}$. Inspired by Younes et al. [2020], we refer to this version of landmark matching as *sub-Riemannian landmark matching*.

It is showed that a dynamic formulation of sub-Riemannian landmark matching can be derived by means of an Euler–Arnold argument, just as in the non-modified case. In more detail, an equation governing the evolution of the time-dependent control variable u is derived. We describe how the dynamic formulation allows for matching to be performed by means of a shooting algorithm, meaning that it is only the initial state of the control u that needs to be determined. Finally, some numerical results illustrate the algorithms.

3.2 Paper 2

In paper 2, we consider the EPDiff equation

$$\begin{aligned}\dot{m} &= \nabla_v^T m - \nabla_m v + \operatorname{div}(v)m, \\ m &= (1 - \Delta)v,\end{aligned}\tag{3.1}$$

where $v \in \mathfrak{X}(M)$ and ∇ denotes the covariant derivative. It can be derived as the geodesic equation on $\operatorname{Diff}(\mathbb{S}^2)$ equipped with the H^1 metric.

The main result of the paper is an extension of the spatial discretization of Euler's equations introduced by Zeitlin [1991, 2004], often called Zeitlin's model, to the EPDiff setting. This approach builds on quantization theory, [Hoppe, 1989, Bordemann et al., 1991, 1994, Le Floch, 2018] and results in a discretization that preserves the underlying Lie–Poisson structure of the equations.

The main result of the paper is that we derive a discretization of Equation (3.1), obtaining the flow

$$\begin{aligned}\dot{W} &= [P, W] \\ P &= (1 - \Delta_N)^{-1} \Delta_N^{-1} W,\end{aligned}\tag{3.2}$$

where N is an integer, $P, W \in \mathfrak{gl}(N, \mathbb{C})$ and Δ_N is the so-called Hoppe–Yau Laplacian, see Hoppe and Yau [1998]. Formally, Equation (3.2) discretizes the vorticity equation associated to Equation (3.1), but as the velocity field v can be reconstructed from the vorticity, this is indeed a discretization of the EPDiff equation. The main point of interest of this discretization lies in its applications in for instance the simulation of compressible flows, optimal transport and shape matching.

Bibliography

- Arnold, V. (1966). Sur la géométrie différentielle des groupes de Lie de dimension infinie et ses applications à l’hydrodynamique des fluides parfaits. *Annales de l’institut Fourier*, 16(1):319–361.
- Arnold, V. I. (1978). *Mathematical Methods of Classical Mechanics*. Springer New York.
- Arnold, V. I. and Khesin, B. A. (1998). *Topological Methods in Hydrodynamics*. Springer New York.
- Azencot, O., Vantzios, O., and Ben-Chen, M. (2018). An explicit structure-preserving numerical scheme for EPDiff. *Computer Graphics Forum*, 37(5):107–119.
- Bauer, M., Joshi, S., and Modin, K. (2015). Diffeomorphic density matching by optimal information transport. *SIAM Journal on Imaging Sciences*, 8(3):1718–1751.
- Beg, F., Miller, M., Trouvé, A., and Younes, L. (2005). Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision*, 61(2):139–157.
- Bistoquet, A., Oshinski, J., and Skrinjar, O. (2008). Myocardial deformation recovery from cine mri using a nearly incompressible biventricular model. *Medical Image Analysis*, 12(1):69–85.
- Bordemann, M., Hoppe, J., Schaller, P., and Schlichenmaier, M. (1991). $\mathfrak{gl}(\infty)$ and geometric quantization. *Communications in Mathematical Physics*, 138(2):209–244.
- Bordemann, M., Meinrenken, E., and Schlichenmaier, M. (1994). Toeplitz quantization of Kähler manifolds and $\mathfrak{gl}(N), N \rightarrow \infty$ limits. *Communications in Mathematical Physics*, 165(2):281–296.
- Bruveris, M. and Holm, D. D. (2013). Geometry of Image Registration: The Diffeomorphism Group and Momentum Maps.

- Celledoni, E., Ehrhardt, M. J., Etmann, C., McLachlan, R. I., Owren, B., Schönlieb, C.-B., and Sherry, F. (2021). Structure-preserving deep learning. *European Journal of Applied Mathematics*, 32(5):888–936.
- Ceritoglu, C., Tang, X., Chow, M., Hadjiabadi, D., Shah, D., Brown, T., Burhanullah, M. H., Trinh, H., Hsu, J. T., Ament, K. A., Crocetti, D., Mori, S., Mostofsky, S. H., Yantis, S., Miller, M. I., and Ratnanather, J. T. (2013). Computational analysis of lddmm for brain mapping. *Frontiers in Neuroscience*, 7.
- da Silva, A. C. (2008). *Lectures on Symplectic Geometry*. Springer Berlin Heidelberg.
- Feydy, J., Charlier, B., Vialard, F.-X., and Peyré, G. (2017). Optimal transport for diffeomorphic registration. In *Medical Image Computing and Computer Assisted Intervention - MICCAI 2017*, pages 291–299. Springer International Publishing.
- Fortney, J. P. (2018). *A Visual Introduction to Differential Forms and Calculus on Manifolds*. Springer International Publishing.
- Hoppe, J. (1989). Diffeomorphism groups, quantization, and $SU(\infty)$. *International Journal of Modern Physics A*, 04(19):5235–5248.
- Hoppe, J. and Yau, S.-T. (1998). Some properties of matrix harmonics on S^2 . *Communications in Mathematical Physics*, 195(1):67–77.
- Joshi, S. and Miller, M. (2000). Landmark matching via large deformation diffeomorphisms. *IEEE Transactions on Image Processing*, 9(8):1357–1370.
- Kirillov, A. (2004). *Lectures on the Orbit Method*. American Mathematical Society.
- Koszul, J.-L. and Zou, Y. M. (2019). *Introduction to Symplectic Geometry*. Springer Singapore.
- Larsson, S., Matsuo, T., Modin, K., and Molteni, M. (2016). Discrete variational derivative methods for the epdiff equation.
- Le Floch, Y. (2018). *A Brief Introduction to Berezin–Toeplitz Operators on Compact Kähler Manifolds*. Springer International Publishing.
- Lee, J. M. (2012). *Introduction to Smooth Manifolds*. Springer.
- Lee, J. M. (2018). *Introduction to Riemannian Manifolds*. Springer International Publishing.
- Li, Q., Chen, L., Cheng, T., and E, W. (2017). Maximum Principle Based Algorithms for Deep Learning. *Journal of Machine Learning Research*, 18(1):5998–6026.

- Marsden, J. E. and Ratiu, T. (1999). *Introduction to Mechanics and Symmetry*. Springer.
- Modin, K. (2019). Geometric hydrodynamics: from Euler, to Poincaré, to Arnold.
- Mumford, D. and Michor, P. W. (2013). On Euler’s equation and ‘EPDiff’. *Journal of Geometric Mechanics*, 5(3):319–344.
- Qiu, A., Fennema-Notestine, C., Dale, A. M., and Miller, M. I. (2009). Regional shape abnormalities in mild cognitive impairment and Alzheimer’s disease. *Neuroimage*, 45(3):656–661.
- Risser, L., Vialard, F.-X., Baluwala, H. Y., and Schnabel, J. A. (2013). Piecewise-diffeomorphic image registration: Application to the motion estimation between 3D CT lung images with sliding conditions. *Medical Image Analysis*, 17(2):182–193.
- Vialard, F.-X., Kwitt, R., Wei, S., and Niethammer, M. (2020). A shooting formulation of deep learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 11828–11838. Curran Associates, Inc.
- Younes, L. (2010). *Shapes and Diffeomorphisms*. Springer.
- Younes, L., Gris, B., and Trounev, A. (2020). *Sub-Riemannian Methods in Shape Analysis*, pages 463–495. Springer.
- Zeitlin, V. (1991). Finite-mode analogs of 2D ideal hydrodynamics: Coadjoint orbits and local canonical structure. *Physica D: Nonlinear Phenomena*, 49(3).
- Zeitlin, V. (2004). Self-consistent finite-mode approximations for the hydrodynamics of an incompressible fluid on nonrotating and rotating spheres. *Physical Review Letters*, 93(26).
- Öktem, O., Chen, C., Onur Domanic, N., Ravikumar, P., and Bajaj, C. (2017). Shape-based image reconstruction using linearized deformations. *Inverse Problems*, 33(3):035004.

