Machine learning and geometric data analysis

Nicolas Charon

High-dimensional data analysis group meeting

March 31, 2025

Geometric data science

The fields of **shape analysis** and more broadly of **geometric data science** can be viewed as the areas of applied mathematics concerned with building adequate statistical methods and machine learning tools for datasets of geometric objects, i.e. for data that does not naturally live in standard Euclidean or functional spaces.

Geometric data science

The fields of **shape analysis** and more broadly of **geometric data science** can be viewed as the areas of applied mathematics concerned with building adequate statistical methods and machine learning tools for datasets of geometric objects, i.e. for data that does not naturally live in standard Euclidean or functional spaces.

Such geometric data is ubiquitous in many applications to computer vision, biomedical imaging, computational anatomy, shape optimization. It also appears in various forms: landmarks, point clouds, curves, surfaces, graphs, tensors and tensor fields, measures...

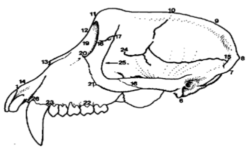
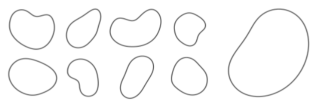


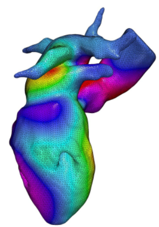
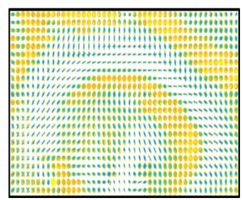

1. The Riemannian metric framework on shape spaces

2. Deep learning for shape analysis

3. Riemannian shape models in machine learning

WHY RIEMANNIAN METRICS?

Equipping shape spaces with *Riemannian metrics* is often a key step to extend statistical concepts and machine learning methods to geometric data. Indeed, the Riemannian setting allows for (relatively) natural generalizations of Fréchet means, regression, PCA, clustering, stochastic processes...

Geodesic path: interpolation/ extrapolation

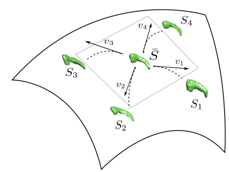

Fréchet mean: atlas estimation

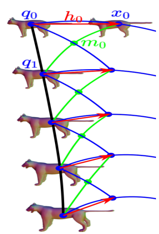

Parallel transport: motion transfer

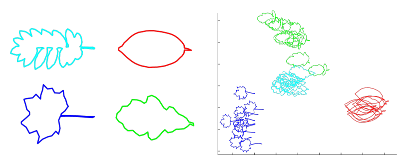

Metric-based clustering

The intrinsic and extrinsinc Riemannian frameworks

The difficulty in defining proper metrics on shape spaces lies in the non-canonical nature of those spaces that often carry the structure of infinite-dimensional quotient spaces: shapes are typically equivalence classes of objects modulo certain geometric transformations such as rigid motions, reparametrizations. . .

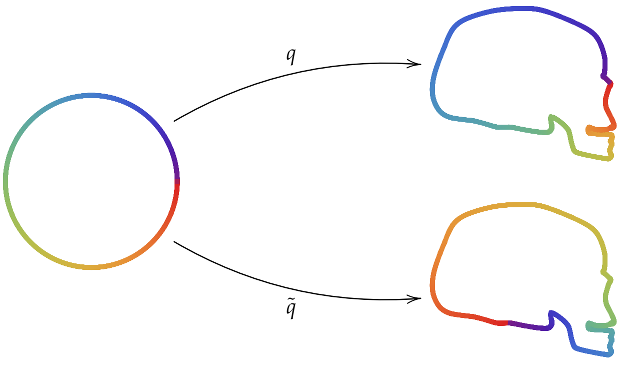

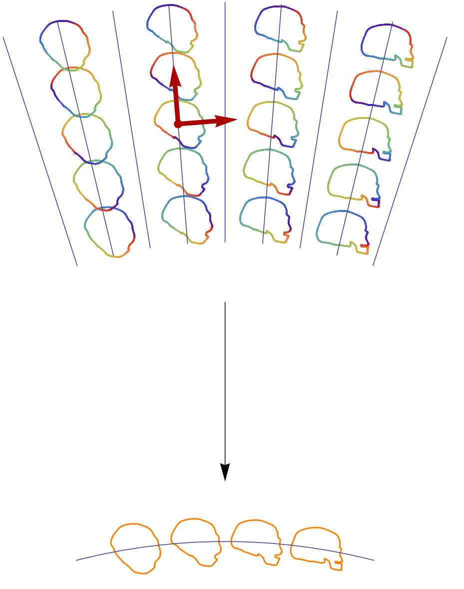

The intrinsic and extrinsinc Riemannian frameworks

Two mainstream approaches for building Riemannian metrics spaces of curves, surfaces or other types of shapes:

- Extrinsic model: build metrics from "fluid like" deformations of the whole ambiant space $\rightarrow$ derived from Grenander's shape space framework.
- Intrinsic model: build metrics on spaces of curves and surfaces from successive "elastic" transformations.

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

- A diffeomorphism of $\mathbb{R}^n$ is a smooth invertible map $\varphi : \mathbb{R}^n \to \mathbb{R}^n$.

[1] Beg, Miller, Trouvé, Younes. Computing Large Deformation Metric Mappings via Geodesic Flows of Diffeomorphisms. 2005.

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

- A diffeomorphism of $\mathbb{R}^n$ is a smooth invertible map $\varphi : \mathbb{R}^n \to \mathbb{R}^n$.
- The group of diffeomorphisms $\mathrm{Diff}(\mathbb{R}^n)$ (or any subgroup $\mathcal{G} \subset \mathrm{Diff}(\mathbb{R}^n)$) can "act" in a specific way on a given space of shapes embedded in $\mathbb{R}^n$. For example, a planar curve c can be transported by $\phi \in \mathrm{Diff}(\mathbb{R}^2)$ as $\varphi \circ c$. A surface $S \subset \mathbb{R}^3$ can be similarly transported by elements of $\mathrm{Diff}(\mathbb{R}^3)$. A function or image $I : \mathbb{R}^n \to \mathbb{R}$ is deformed by $\varphi \in \mathrm{Diff}(\mathbb{R}^n)$ via the action $I \circ \varphi^{-1}$. A measure μ on $\mathbb{R}^n$ can be transported by $\phi \in \mathrm{Diff}(\mathbb{R}^n)$ based on the pushforward action $\phi_{\#}\mu$.

[1] Beg, Miller, Trouvé, Younes. Computing Large Deformation Metric Mappings via Geodesic Flows of Diffeomorphisms. 2005.

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

- A diffeomorphism of $\mathbb{R}^n$ is a smooth invertible map $\varphi : \mathbb{R}^n \to \mathbb{R}^n$.
- The group of diffeomorphisms $\mathrm{Diff}(\mathbb{R}^n)$ (or any subgroup $\mathcal{G} \subset \mathrm{Diff}(\mathbb{R}^n)$) can "act" in a specific way on a given space of shapes embedded in $\mathbb{R}^n$. For example, a planar curve c can be transported by $\phi \in \mathrm{Diff}(\mathbb{R}^2)$ as $\varphi \circ c$. A surface $S \subset \mathbb{R}^3$ can be similarly transported by elements of $\mathrm{Diff}(\mathbb{R}^3)$. A function or image $I : \mathbb{R}^n \to \mathbb{R}$ is deformed by $\varphi \in \mathrm{Diff}(\mathbb{R}^n)$ via the action $I \circ \varphi^{-1}$. A measure μ on $\mathbb{R}^n$ can be transported by $\phi \in \mathrm{Diff}(\mathbb{R}^n)$ based on the pushforward action $\phi_{\#}\mu$.
- The problem can be then reframed as building Riemannian metrics on $\mathrm{Diff}(\mathbb{R}^n)$ which can then be "projected" to the shape space via the particular group action.

[1] Beg, Miller, Trouvé, Younes. Computing Large Deformation Metric Mappings via Geodesic Flows of Diffeomorphisms. 2005.

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

- A diffeomorphism of $\mathbb{R}^n$ is a smooth invertible map $\varphi : \mathbb{R}^n \to \mathbb{R}^n$.
- The group of diffeomorphisms $\mathrm{Diff}(\mathbb{R}^n)$ (or any subgroup $\mathcal{G} \subset \mathrm{Diff}(\mathbb{R}^n)$) can "act" in a specific way on a given space of shapes embedded in $\mathbb{R}^n$. For example, a planar curve c can be transported by $\phi \in \mathrm{Diff}(\mathbb{R}^2)$ as $\varphi \circ c$. A surface $S \subset \mathbb{R}^3$ can be similarly transported by elements of $\mathrm{Diff}(\mathbb{R}^3)$. A function or image $I : \mathbb{R}^n \to \mathbb{R}$ is deformed by $\varphi \in \mathrm{Diff}(\mathbb{R}^n)$ via the action $I \circ \varphi^{-1}$. A measure μ on $\mathbb{R}^n$ can be transported by $\phi \in \mathrm{Diff}(\mathbb{R}^n)$ based on the pushforward action $\phi_{\#}\mu$.
- The problem can be then reframed as building Riemannian metrics on $\mathrm{Diff}(\mathbb{R}^n)$ which can then be "projected" to the shape space via the particular group action.
- The Large Deformation Diffeomorphic Metric Mapping (LDDMM) model[1] provides such a construction by considering deformation groups modeled as flows of non-stationary velocity fields. The *flow* of a time-dependent vector field $t \mapsto v(t, \cdot) \in C^1(\mathbb{R}^n, \mathbb{R}^n)$ is the map $(t, x) \to \varphi^v(t, x)$ defined by:

$$\begin{cases} \varphi^v(0, x) = x, \quad x \in \mathbb{R}^n \\ \partial_t \varphi^v(t, x) = v(t, \varphi^v(t, x)), \quad (t, x) \in [0, 1] \times \mathbb{R}^n \end{cases}$$

[1] Beg, Miller, Trouvé, Younes. Computing Large Deformation Metric Mappings via Geodesic Flows of Diffeomorphisms. 2005.

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

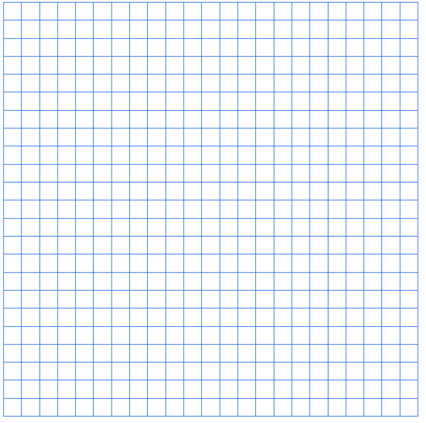

Initial grid

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

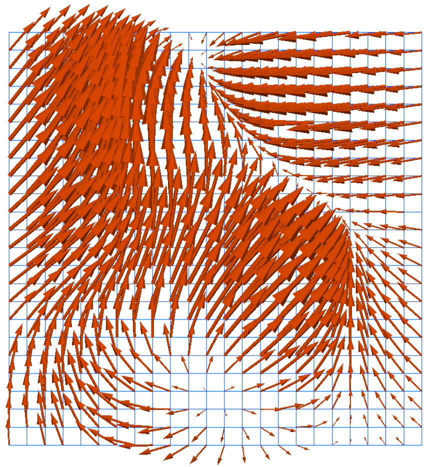

Deformation field $t = 0$

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

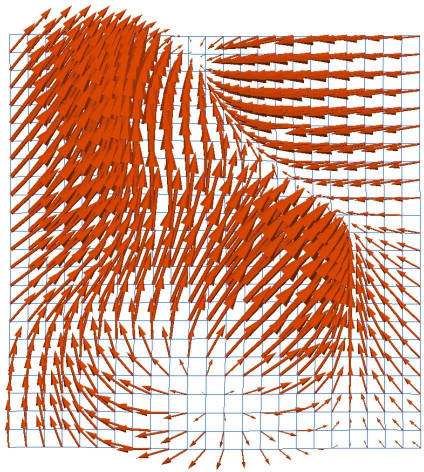

Deformation field $t = 0.1$

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

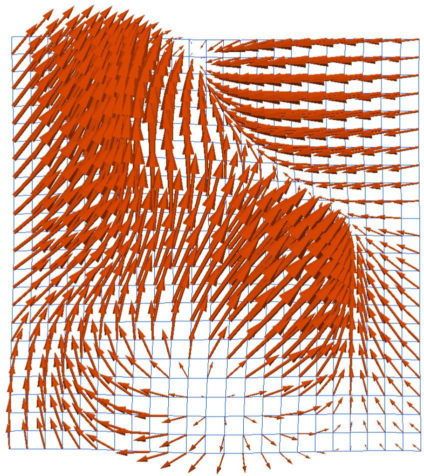

Deformation field $t = 0.2$

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

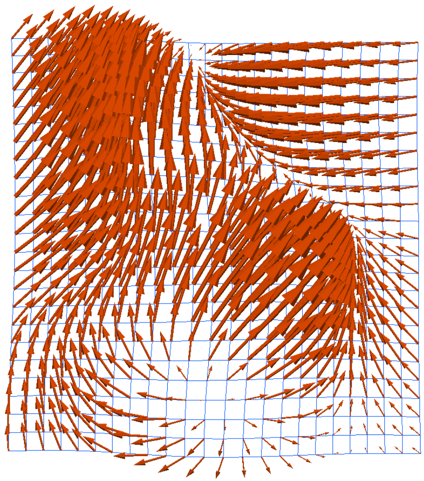

Deformation field $t = 0.3$

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

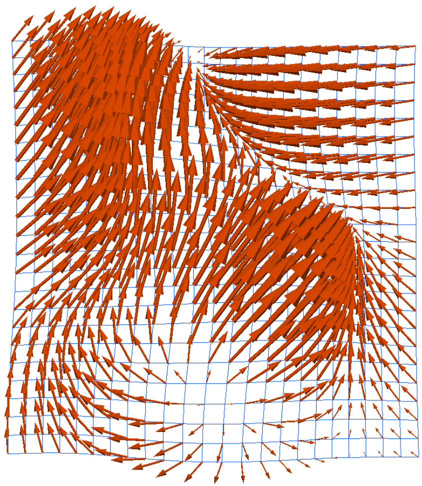

Deformation field $t = 0.4$

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

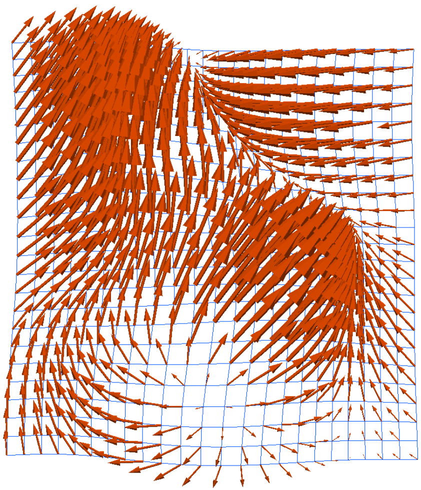

Deformation field $t = 0.5$

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

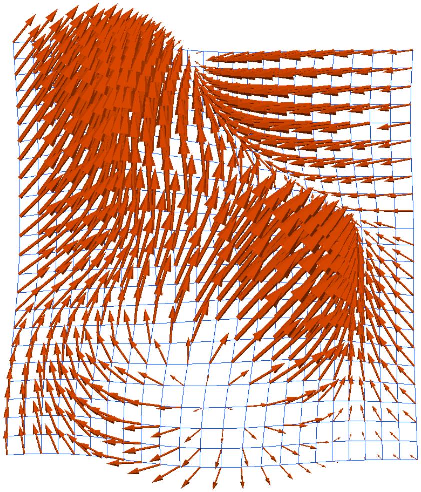

Deformation field $t = 0.6$

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

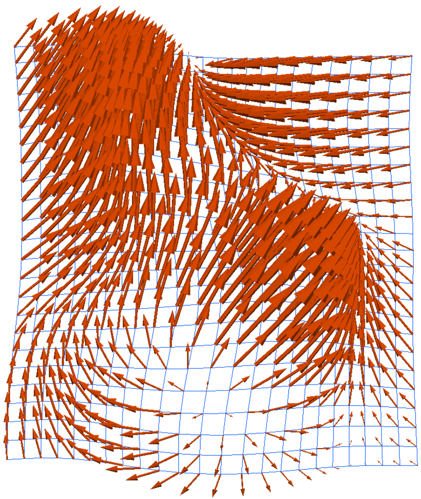

Deformation field $t = 0.7$

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

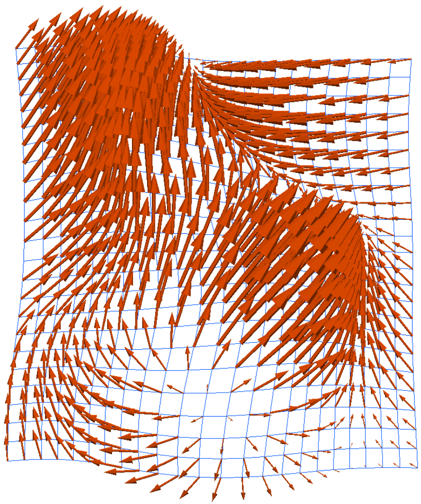

Deformation field $t = 0.8$

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

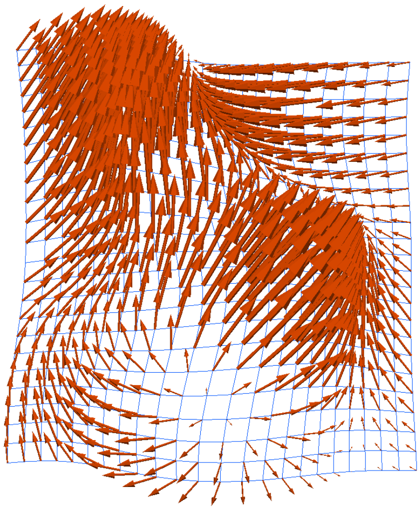

Deformation field $t = 0.9$

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

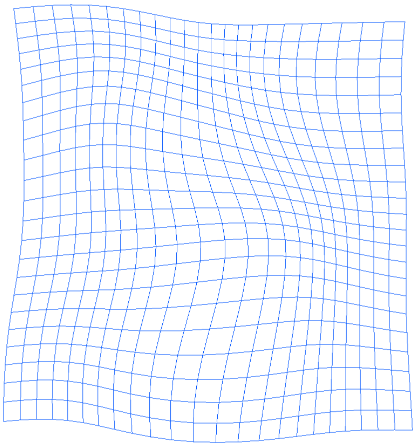

Final deformation grid

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

- Riemannian metrics on the set of resulting flow maps $\varphi^v(1,\cdot)$ can be then obtained by specifying a Hilbert space V of admissible vector fields with Hilbert norm $\|\cdot\|_V$ (e.g. a Sobolev norm). This then leads to a global energy for the deformation $\varphi^v(1,\cdot)$ given by $\int_0^1 \|v(t,\cdot)\|_V^2 dt$.
- Given two shapes S_0 and S_1 living in $\mathbb{R}^n$, the distance between them is computed by finding $v \in L^2([0,1],V)$ such that $\varphi^v(1,\cdot)$ maps S_0 to S_1 and has minimal total energy, which is a *registration problem*. The path of shapes $S(t) = \varphi^v(t, S_0)$ can be then viewed as a geodesic interpolating between S_0 and S_1.

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

- Riemannian metrics on the set of resulting flow maps $\varphi^v(1,\cdot)$ can be then obtained by specifying a Hilbert space V of admissible vector fields with Hilbert norm $\|\cdot\|_V$ (e.g. a Sobolev norm). This then leads to a global energy for the deformation $\varphi^v(1,\cdot)$ given by $\int_0^1 \|v(t,\cdot)\|_V^2 dt$.
- Given two shapes S_0 and S_1 living in $\mathbb{R}^n$, the distance between them is computed by finding $v \in L^2([0,1],V)$ such that $\varphi^v(1,\cdot)$ maps S_0 to S_1 and has minimal total energy, which is a *registration problem*. The path of shapes $S(t) = \varphi^v(t, S_0)$ can be then viewed as a geodesic interpolating between S_0 and S_1.

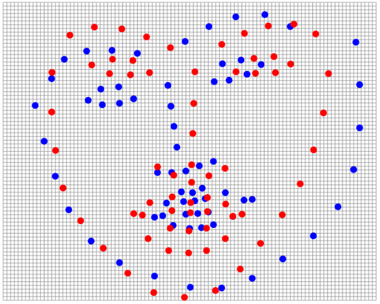

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

- Riemannian metrics on the set of resulting flow maps $\varphi^v(1,\cdot)$ can be then obtained by specifying a Hilbert space V of admissible vector fields with Hilbert norm $\|\cdot\|_V$ (e.g. a Sobolev norm). This then leads to a global energy for the deformation $\varphi^v(1,\cdot)$ given by $\int_0^1 \|v(t,\cdot)\|_V^2 dt$.
- Given two shapes S_0 and S_1 living in $\mathbb{R}^n$, the distance between them is computed by finding $v \in L^2([0,1],V)$ such that $\varphi^v(1,\cdot)$ maps S_0 to S_1 and has minimal total energy, which is a *registration problem*. The path of shapes $S(t) = \varphi^v(t, S_0)$ can be then viewed as a geodesic interpolating between S_0 and S_1.

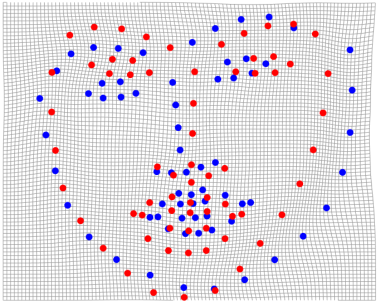

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

- Riemannian metrics on the set of resulting flow maps $\varphi^v(1,\cdot)$ can be then obtained by specifying a Hilbert space V of admissible vector fields with Hilbert norm $\|\cdot\|_V$ (e.g. a Sobolev norm). This then leads to a global energy for the deformation $\varphi^v(1,\cdot)$ given by $\int_0^1 \|v(t,\cdot)\|_V^2 dt$.
- Given two shapes S_0 and S_1 living in $\mathbb{R}^n$, the distance between them is computed by finding $v \in L^2([0,1],V)$ such that $\varphi^v(1,\cdot)$ maps S_0 to S_1 and has minimal total energy, which is a *registration problem*. The path of shapes $S(t) = \varphi^v(t, S_0)$ can be then viewed as a geodesic interpolating between S_0 and S_1.

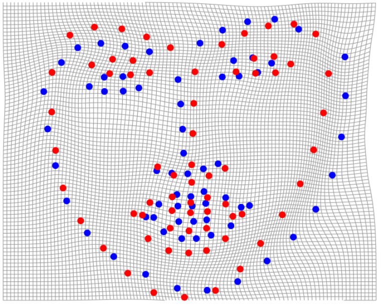

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

- Riemannian metrics on the set of resulting flow maps $\varphi^v(1,\cdot)$ can be then obtained by specifying a Hilbert space V of admissible vector fields with Hilbert norm $\|\cdot\|_V$ (e.g. a Sobolev norm). This then leads to a global energy for the deformation $\varphi^v(1,\cdot)$ given by $\int_0^1 \|v(t,\cdot)\|_V^2 dt$.
- Given two shapes S_0 and S_1 living in $\mathbb{R}^n$, the distance between them is computed by finding $v \in L^2([0,1],V)$ such that $\varphi^v(1,\cdot)$ maps S_0 to S_1 and has minimal total energy, which is a *registration problem*. The path of shapes $S(t) = \varphi^v(t,S_0)$ can be then viewed as a geodesic interpolating between S_0 and S_1.

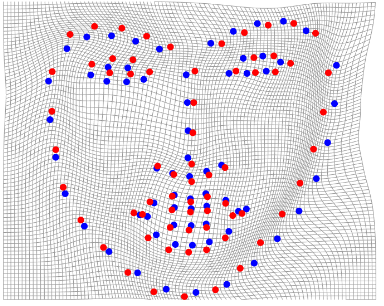

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

- Riemannian metrics on the set of resulting flow maps $\varphi^v(1,\cdot)$ can be then obtained by specifying a Hilbert space V of admissible vector fields with Hilbert norm $\|\cdot\|_V$ (e.g. a Sobolev norm). This then leads to a global energy for the deformation $\varphi^v(1,\cdot)$ given by $\int_0^1 \|v(t,\cdot)\|_V^2 dt$.
- Given two shapes S_0 and S_1 living in $\mathbb{R}^n$, the distance between them is computed by finding $v \in L^2([0,1],V)$ such that $\varphi^v(1,\cdot)$ maps S_0 to S_1 and has minimal total energy, which is a *registration problem*. The path of shapes $S(t) = \varphi^v(t, S_0)$ can be then viewed as a geodesic interpolating between S_0 and S_1.

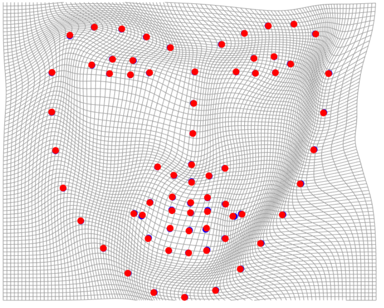

Extrinsic model: Large Deformation Diffeomorphic Metric Mapping (LDDMM)

The key features of the extrinsic/LDDMM approach are[2]:

- Shapes are compared via global diffeomorphic transformations of the whole ambiant space.
- It enforces a diffeomorphic mapping between shapes, and thus technically only applies to objects belonging to the same orbit under the action of the group (in particular these must have a consistent topology).
- It is very versatile: can be applied to shapes of various nature (landmarks, curves, surfaces, images...) as long as there is a well-defined action of diffeomorphisms on those shapes.

[2] Younes. Shapes and diffeomorphisms. 2019

Intrinsic model: elastic shape analysis

In contrast with LDDMM, the intrinsic approach aims at directly building Riemannian metrics on the space of smooth curves or surfaces. Formally, a tangent vectors to the shape space at a shape S can be viewed as smooth vector field $v : S \to \mathbb{R}^n$ defined on S itself, representing an infinitesimal variation of S.

A Riemannian metric should then assign a cost/energy $\|v\|_S$ to the infinitesimal deformation of S in the direction of v. Geodesics and geodesic distances are obtained by finding shape paths $t \mapsto S(t)$ with minimal total energy.

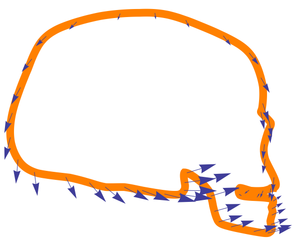

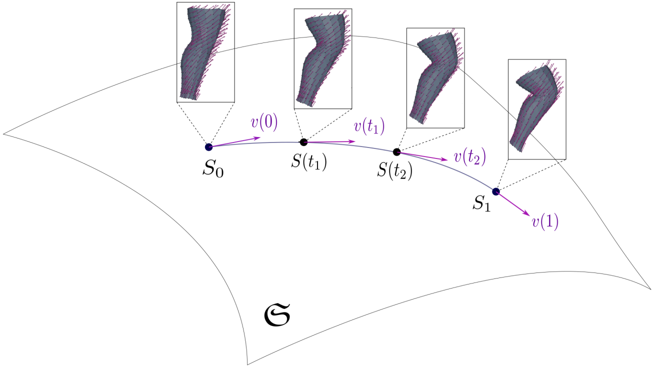

INTRINSIC RIEMANNIAN METRICS

The definition of the Riemannian energy $\|v\|_S$ must follow certain conditions to eventually yield a proper Riemannian distance on the shape space[3]. A natural idea is to define $\|v\|_S$ based on linear membrane elasticity theory, combining bending/stretching energies of the infinitesimal deformation v of S.

For mathematical reasons, it is often preferable to introduce higher than first order metrics and consider the class of second-order intrinsic Sobolev metrics taking the generic form:

$$\|v\|_S^2 = \int_S (a_0|v(x)|^2 + a_1|dv(x)|^2 + a_2|\Delta_S v|^2) d\mathrm{vol}_S(x)$$

with dv being the Jacobian of v along S, Δ_S the Laplace-Beltrami operator on S, and vol_S the arclength/area measure over the curve/surface S. $a_0, a_1, a_2 > 0$ are weighing constants.

[3] Michor & Mumford. An overview of the Riemannian metrics on spaces of curves using the Hamiltonian approach. 2007

INTRINSIC VS EXTRINSIC MODEL[4] [5]

In comparison to the extrinsic deformation approach, the key features of the intrinsic model are:

- This model is specific to shape spaces of curves, surfaces (and more generally submanifolds) but not easily extendable to other types of geometric data such as images, tensors...
- The deformation energy measures local regularity of a vector field along the shape itself rather than considering deformation fields over the whole ambiant space.
- Intrinsic metrics have a somewhat clearer “physical” interpretation and usually lead to less intensive numerical computations than LDDMM on discretized curves/surfaces.
- Intrinsic metrics not guarantee that geodesics remain diffeomorphically equivalent to the initial shape: in particular self-intersections may be created and removed along a geodesic path.

[4] Bauer, Charon & Younes. Metric registration of curves and surfaces using optimal control. 2019.

[5] Charon & Younes. Shape spaces: Shape spaces: From geometry to biological plausibility. 2022.

Riemannian shape analysis: successes and limitations

Both the intrinsic and extrinsic frameworks have witnessed, over the past two decades, considerable developments on the theoretical and numerical levels.

Multiple code libraries now exist which allow to numerically approximate distances and geodesics between shapes and perform some basic statistical tasks on shape spaces (e.g. Fréchet mean estimation, tangent PCA, regression...). For the LDDMM model, these include among others: **Deformetrica**, **FshapesTk**, **Keops**, **CLAIRE**. . .
For intrinsic metrics: **H2_SurfaceMatch**, **fdasrvf**, **GeomStats**. . .

Riemannian shape analysis: successes and limitations

Both the intrinsic and extrinsic frameworks have witnessed, over the past two decades, considerable developments on the theoretical and numerical levels.

Multiple code libraries now exist which allow to numerically approximate distances and geodesics between shapes and perform some basic statistical tasks on shape spaces (e.g. Fréchet mean estimation, tangent PCA, regression...). For the LDDMM model, these include among others: **Deformetrica**, **FshapesTk**, **Keops**, **CLAIRE**...
For intrinsic metrics: **H2_SurfaceMatch**, **fdasrvf**, **GeomStats**...

A clear downside of Riemannian methods for geometric data is the high computational cost involved. In both models, computing a single geodesic or distance involves solving a high-dimensional ($\sim$ number of vertices in the shape) and non-convex optimal control problem. This often prevents such methods to scale up to modern geometric datasets with large number of high resolution observations.

1. The Riemannian metric framework on shape spaces

2. Deep learning for shape analysis

3. Riemannian shape models in machine learning

Deep learning for functions and images

The use and training of neural networks have resulted in considerable improvement in the performance of algorithms in most tasks involving time series, functions or images.

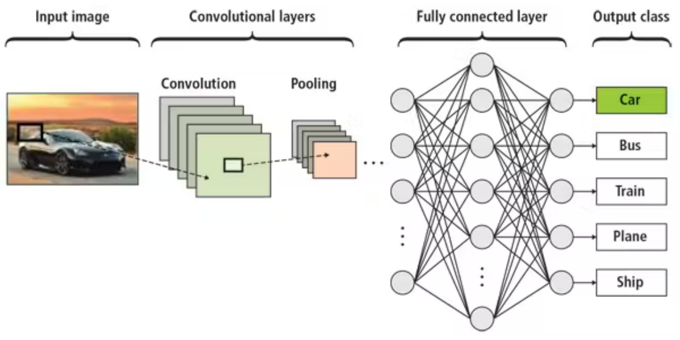

A typical deep learning pipeline for image classification.

A convenient aspect of such data is that the input/output of neural networks are typically well-structured, e.g. they are samples from fixed regular grids of pixels or voxels, which is a key feature to enable the use of CNN layers for instance.

The challenges of deep learning on geometric data

When dealing with (unstructured) geometric data, the design of adequate neural networks architectures is an active area of research[6]. This is due to several specific challenges:

- Shapes such as curves or surfaces are often obtained via segmentation of images or signals $\rightarrow$ data usually presents various types of irregularities such as noise, inconsistent sampling or topology...

[6] Bronstein, Bruna, Cohen, Velickovic. Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges. 2021

The challenges of deep learning on geometric data

When dealing with (unstructured) geometric data, the design of adequate neural networks architectures is an active area of research[6]. This is due to several specific challenges:

- Shapes such as curves or surfaces are often obtained via segmentation of images or signals $\rightarrow$ data usually presents various types of irregularities such as noise, inconsistent sampling or topology…
- Geometric data is fundamentally non-Euclidean as objects are equivalence classes modulo the action of certain transformations $\rightarrow$ neural network models and loss functions should embed these fundamental invariances in their design.

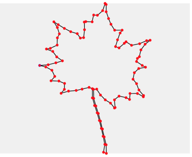
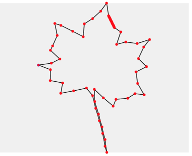
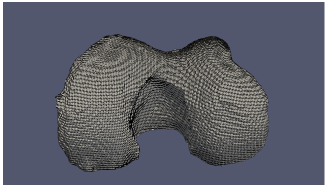
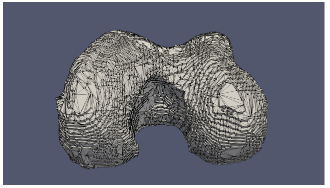

[6] Bronstein, Bruna, Cohen, Velickovic. Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges. 2021

The challenges of deep learning on geometric data

When dealing with (unstructured) geometric data, the design of adequate neural networks architectures is an active area of research[6]. This is due to several specific challenges:

- Shapes such as curves or surfaces are often obtained via segmentation of images or signals $\rightarrow$ data usually presents various types of irregularities such as noise, inconsistent sampling or topology...
- Geometric data is fundamentally non-Euclidean as objects are equivalence classes modulo the action of certain transformations $\rightarrow$ neural network models and loss functions should embed these fundamental invariances in their design.

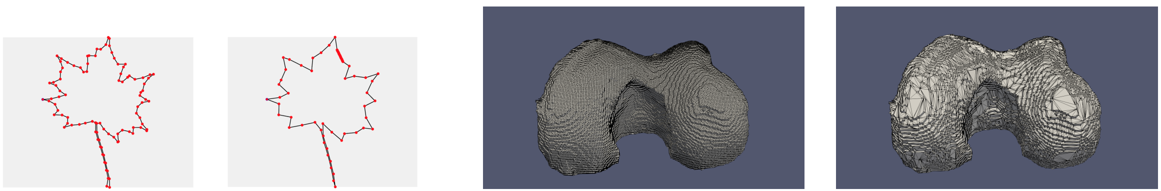

- Unlike with images, no access to massive training sets $\rightarrow$ geometric deep learning architectures cannot have too many trainable parameters to prevent overfitting.

[6] Bronstein, Bruna, Cohen, Velickovic. Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges. 2021

Naive strategies: normalization, data augmentation

For simple enough geometric objects (e.g. single curves), one possible approach is to rely on standard architectures by first applying some prior normalization step to the input data such as: rigid registration to a template, resampling to arclength parametrization with fixed number of points... This can be further combined, during training, with data augmentation strategies.

[7] Hartman, Sukurdeep, Charon, Klassen, Bauer. Supervised deep learning of elastic SRV distances on the shape space of curves. 2021.

NAIVE STRATEGIES: NORMALIZATION, DATA AUGMENTATION

For simple enough geometric objects (e.g. single curves), one possible approach is to rely on standard architectures by first applying some prior normalization step to the input data such as: rigid registration to a template, resampling to arclength parametrization with fixed number of points... This can be further combined, during training, with data augmentation strategies.

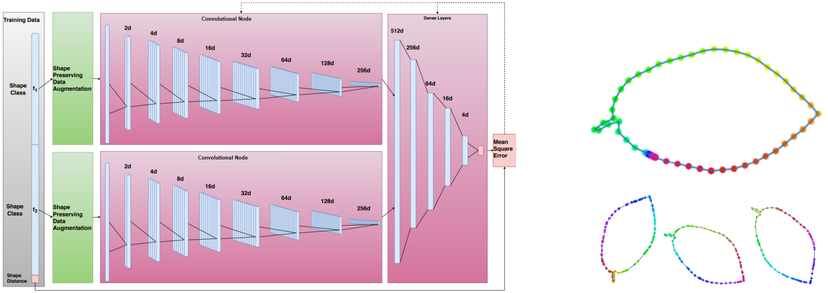

Figure. Supervised deep learning for prediction of elastic distance between closed curves, using pre-normalization to a fixed number of 100 points combined with data augmentation[7].

[7] Hartman, Sukurdeep, Charon, Klassen, Bauer. Supervised deep learning of elastic SRV distances on the shape space of curves. 2021.

Naive strategies: normalization, data augmentation

For simple enough geometric objects (e.g. single curves), one possible approach is to rely on standard architectures by first applying some prior normalization step to the input data such as: rigid registration to a template, resampling to arclength parametrization with fixed number of points... This can be further combined, during training, with data augmentation strategies.

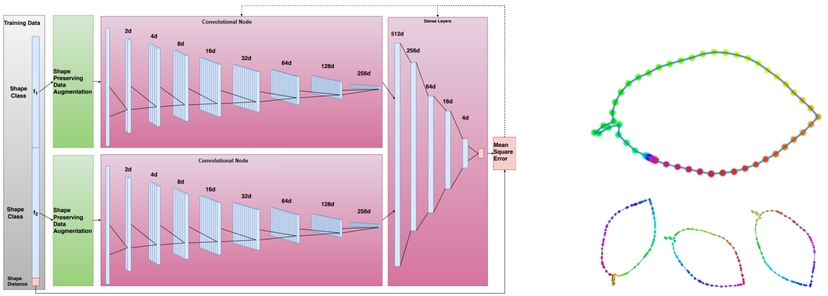

Figure. Supervised deep learning for prediction of elastic distance between closed curves, using pre-normalization to a fixed number of 100 points combined with data augmentation[7].

Trained networks tend to not recover invariances well and have limited generalizability far from training set. Normalization can be significantly more challenging with more complex data such as mesh surfaces, graphs...

[7] Hartman, Sukurdeep, Charon, Klassen, Bauer. Supervised deep learning of elastic SRV distances on the shape space of curves. 2021.

Designing invariant neural network models

An arguably sounder strategy is to instead incorporate all the desired invariances to e.g. parametrization/sampling, rigid motion, scaling... into the network architecture itself.

This could be achieved in different ways:

1. By building adequate extensions of convolutional layers for more general mesh objects.
2. By embedding geometric data into invariant feature representation spaces on which to perform learning and training.

Designing invariant neural network models

An arguably sounder strategy is to instead incorporate all the desired invariances to e.g. parametrization/sampling, rigid motion, scaling. . . into the network architecture itself.

This could be achieved in different ways:

1. By building adequate extensions of convolutional layers for more general mesh objects.
2. By embedding geometric data into invariant feature representation spaces on which to perform learning and training.

Designing invariant neural network models

An arguably sounder strategy is to instead incorporate all the desired invariances to e.g. parametrization/sampling, rigid motion, scaling. . . into the network architecture itself.

This could be achieved in different ways:

1. By building adequate extensions of convolutional layers for more general mesh objects.
2. By embedding geometric data into invariant feature representation spaces on which to perform learning and training.

The field **geometric measure theory** (GMT) has paved the idea of viewing geometric shapes as measures, which often provides an effective representation to frame and tackle variational problems, and to perform geometric processing tasks.

Point clouds as measures of $\mathbb{R}^n$

A point cloud in $\mathbb{R}^n$ is an unordered set of points $m = \{x_i\}_{i=1,\dots,p}$ with $p \in \mathbb{N}$ and $x_i \in \mathbb{R}^n$ for all $i = 1, \dots, p$. In terms of shape space, point clouds are elements of the reunion of quotient spaces $\mathcal{S} = \bigcup_{p\in\mathbb{N}} ((\mathbb{R}^n)^p/\mathrm{Sym}(p))$. However, $\mathcal{M}$ has a natural embedding into the space $\mathcal{M}^+(\mathbb{R}^n)$ of positive Radon measures of $\mathbb{R}^n$: $\{x_i\} \in \mathcal{S} \mapsto \sum_{i=1}^p \delta_{x_i} \in \mathcal{M}^+(\mathbb{R}^n)$.

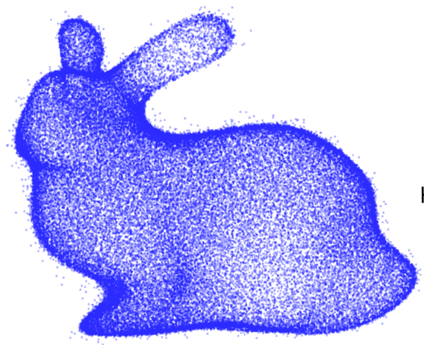

$$\longmapsto \mu = \sum_{k=1}^{p} \delta_{x_i}$$

This embedding provides a natural way to build distances and loss functions between point clouds based on metrics on $\mathcal{M}^+(\mathbb{R}^n)$, which include for instance Wasserstein metrics or kernel maximum mean discrepancies.

CURVES, SURFACES AS VARIFOLDS[8]

One natural generalization of such embeddings to geometric structures like curves and surfaces is provided by the concept of **varifold**. For $n = 2$ or $n = 3$, a varifold is a measure on the product space $\mathbb{R}^n \times \mathbb{S}^{n-1}$. To any shape (continuous or discrete) can be associated a varifold that represents the distribution of point positions with oriented directions.

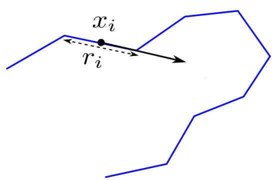

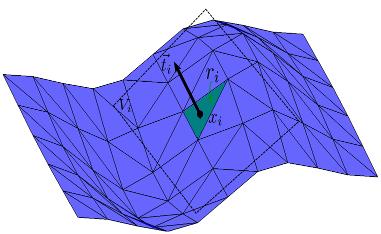

Varifold embedding: $S \mapsto \mu_S = \sum_{i=1}^{p} r_i \delta_{(x_i, \vec{t}_i)} \in \mathcal{M}^+(\mathbb{R}^n \times \mathbb{S}^{n-1})$ where $x_i, r_i, \vec{t}_i$ are the barycenter, length/area, unit tangent/normal vector of cell i.

[8]Kaltenmark, Charon, Charlier. A general framework for curve and surface comparison and registration with oriented varifolds. 2017.

Curves, surfaces as varifolds[8]

- A key property of the varifold representation is that it is robust to changes in parametrization (or discrete sampling for meshes). As such, any metric on the space of varifolds will automatically induce a well-defined distance between shapes independent of their parametrization.
- *Reproducing Kernel Hilbert Space* (RKHS) metrics in particular, defined via a positive definite kernel function K on $\mathbb{R}^n \times \mathbb{S}^{n-1}$, have been proposed and can be computed explicitly from evaluations of K. For $\mu_S = \sum_{i=1}^p r_i \delta_{(x_i, \vec{t}_i)}$ and $\mu_{S'} = \sum_{j=1}^q r'_j \delta_{(x'_j, \vec{t}'_j)}$, it writes:

$$\begin{aligned}
&\|\mu_S - \mu_{S'}\|^2_{\mathrm{Var}} = \langle \mu_S, \mu_S \rangle_{\mathrm{Var}} - 2\langle \mu_S, \mu_{S'} \rangle_{\mathrm{Var}} + \langle \mu_{S'}, \mu_{S'} \rangle_{\mathrm{Var}} \\
&= \sum_{i=1}^p \sum_{j=1}^p K((x_i, \vec{t}_i), (x_j, \vec{t}_j)) r_i r_j - 2 \sum_{i=1}^p \sum_{j=1}^q K((x_i, \vec{t}_i), (x'_j, \vec{t}'_j)) r_i r'_j + \sum_{i=1}^q \sum_{j=1}^q K((x'_i, \vec{t}'_i), (x'_j, \vec{t}'_j)) r'_i r'_j
\end{aligned}$$

- These are especially well-suited to build differentiable loss functions to quantify shape discrepancy.

[8] Kaltenmark, Charon, Charlier. A general framework for curve and surface comparison and registration with oriented varifolds. 2017.

Leveraging measure representations for geometric deep learning

Beyond the construction of loss functions, can varifolds also be used as invariant feature representations for the design of geometric neural network models? Some preliminary work (c.f. Emmanuel's presentation) suggest two promising approaches:

1. Exploit the duality of varifolds with test functions $\omega \in C_c(\mathbb{R}^n \times \mathbb{S}^{n-1})$ and consider feature representations of the form $(\mu_S(\omega_1), \ldots, \mu_S(\omega_K)) \in \mathbb{R}^K$ for some (trainable) set of test functions $(\omega_1, \ldots, \omega_K)$.
2. Fix a template shape T and represent any shape S as the gradient field of the varifold distance from μ_S to μ_T with respect to the points of S. (**VariGrad**)

Leveraging measure representations for geometric deep learning

Beyond the construction of loss functions, can varifolds also be used as invariant feature representations for the design of geometric neural network models? Some preliminary work (c.f. Emmanuel's presentation) suggest two promising approaches:

1. Exploit the duality of varifolds with test functions $\omega \in C_c(\mathbb{R}^n \times \mathbb{S}^{n-1})$ and consider feature representations of the form $(\mu_S(\omega_1), \dots, \mu_S(\omega_K)) \in \mathbb{R}^K$ for some (trainable) set of test functions $(\omega_1, \dots, \omega_K)$.
2. Fix a template shape T and represent any shape S as the gradient field of the varifold distance from μ_S to μ_T with respect to the points of S. (**VariGrad**)

Some open questions:

- Can one obtain injectivity and/or stability properties for those feature representations?
- Varifolds provide invariance to parametrization but not to rigid motions or scaling. Is it possible to build useful feature maps on the space of shapes modulo rigid motion, and modulo rescaling?

1. The Riemannian metric framework on shape spaces

2. Deep learning for shape analysis

3. Riemannian shape models in machine learning

DEFORMATION MODELS FOR PDE OPERATOR LEARNING[9]

Operator learning is a supervised machine learning framework which attempts to estimate the solution of a PDE given initial/boundary conditions. This is traditionally done over a domain with a fixed geometry. By leveraging the representation of diffeomorphic transformations in the LDDMM setting, one can extend the approach to the prediction of PDE solutions over varying domains.

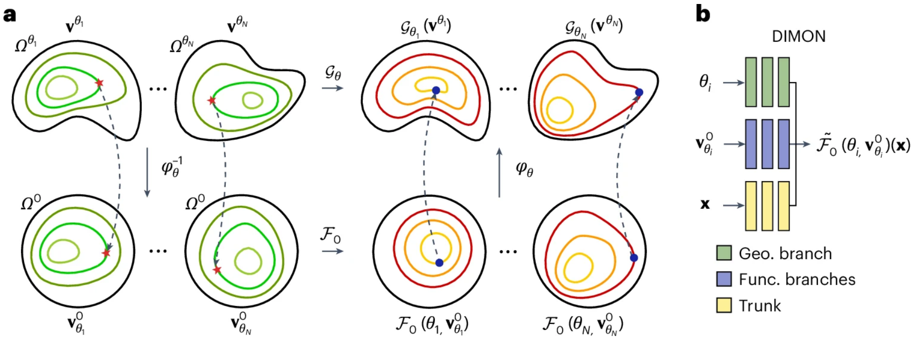

Overview of the DIMON approach

[9] Yin, Charon, Brody, Lu, Trayanova, Maggioni. A scalable framework for learning the geometry-dependent solution operators of partial differential equations. 2024.

DEFORMATION MODELS FOR PDE OPERATOR LEARNING

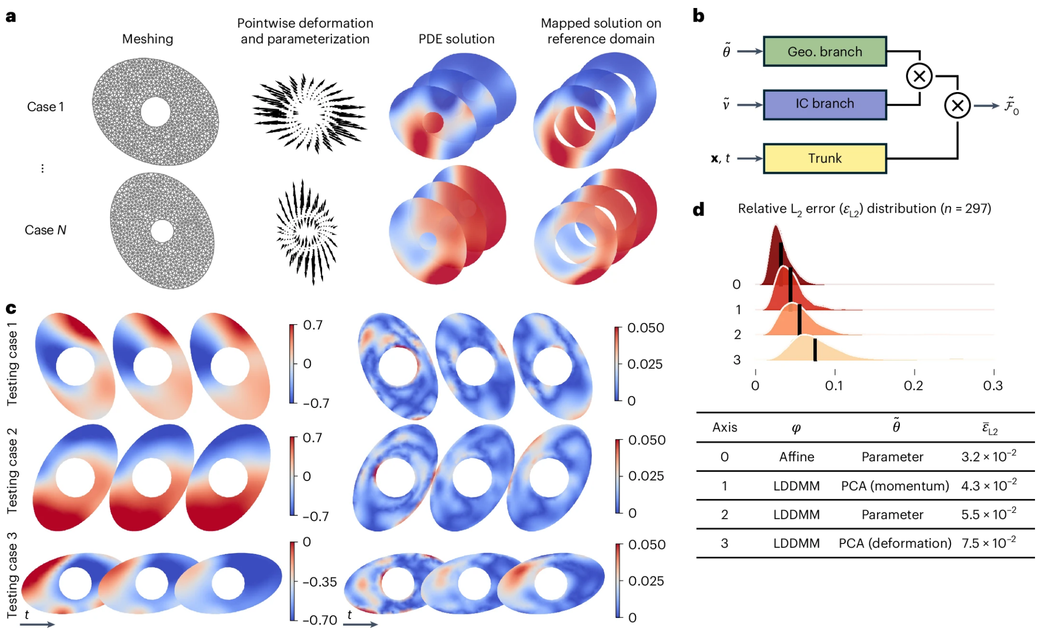

Axis	φ	$\tilde{\theta}$	$\bar{\varepsilon}_{L2}$
0	Affine	Parameter	3.2×10^{-2}
1	LDDMM	PCA (momentum)	4.3×10^{-2}
2	LDDMM	Parameter	5.5×10^{-2}
3	LDDMM	PCA (deformation)	7.5×10^{-2}

Application of DIMON to predict solutions of a reaction-diffusion equation on annuli domains