

Optimal Steps for Fast Diffeomorphic Shape Registration

Hadrien Bigo-Balland¹, Tom Boeken^{1,2}[0000–0003–2990–5180], and Jean Feydy¹[0000–0001–6049–563X]

¹ Inria, Université Paris Cité, Inserm, HeKA, F-75015 Paris, France

² Department of Vascular and Oncological Interventional Radiology, Hôpital Européen Georges Pompidou, AP-HP, 75015 Paris, France
hadrien.bigo-balland@inria.fr

Abstract. This paper presents a fast, learning-free registration method tailored to clinical constraints. While shape registration is fundamental to medical imaging, existing techniques often force a trade-off between computational scalability and strict topology preservation. Current diffeomorphic methods are either computationally expensive or demand extensive pre-processing, such as training neural networks on large datasets. This renders them impractical for clinical studies, which often rely on limited cohorts. Image-based methods also require large memory banks to process high-resolution data, while surface-based approaches have received less attention in the deep learning era. In this context, we propose a principled registration algorithm that alternates between a feature-aware matching scheme, a conjugate gradient solver and a diffeomorphic shooting step. Our pipeline can handle point clouds, curves and surface meshes. It is compatible with clinical workflows and can register a pair of shapes sampled with 10k points each in one second, without landmarks or supervision. We demonstrate consistent results on several anatomical structures and showcase the efficiency of our method by computing an atlas from hundreds of complex vertebral scans in a few minutes. Upon publication, our code will be released under the permissive MIT license at: scikit-shapes.github.io.

Keywords: Shape · Registration · Optimization · LDDMM.

1 Introduction and Related Works

Context. Non-rigid shape registration facilitates essential tasks in medical imaging, including statistical shape modeling [9] and the generation of patient-specific finite-element structures [31]. Within orthopedic imaging, it enables precise clinical evaluations of pathologies like idiopathic scoliosis [12]. As clinical cohorts increasingly scale to include hundreds or thousands of subjects, there is a need for methods that are both computationally efficient and fully automated.

Trade-offs. Traditional registration methods typically treat non-rigid alignment as a non-convex optimization problem [19], whereas the last decade has seen




Fig. 1. Our method can handle both point clouds and meshes. Using an affordable gaming GPU (Nvidia RTX 3090), we register vertebrae sampled with 10k points each in 1 s (left) and lung vascular trees sampled with 40k points each in 20 s (right).

a surge in learning-based alternatives [3]. These paradigms offer complementary advantages: while learning-based methods provide rapid inference and state-of-the-art accuracy, they necessitate computationally intensive training phases and may lack robustness when encountering atypical anatomical variations [14].

The selection of a shape encoding format is equally context-dependent. Although anatomical data is primarily acquired in volumetric formats, processing high-resolution 3D grids poses significant memory constraints. In contrast, point-based representations such as point clouds, curves, and surface meshes discard intensity information but are orders of magnitude more efficient, as they localize samples only on relevant boundary regions [15]. This facilitates precise morphometry on standard clinical workstations.

While preserving topological integrity is essential to prevent non-physical artifacts such as self-intersections or mesh folding, providing such a guarantee at an affordable computational cost remains challenging. Principled frameworks such as Large Deformation Diffeomorphic Metric Mapping (LDDMM) [4,34] and repulsive shells [27] offer rigorous guarantees but suffer from high computational complexity. Consequently, for large-scale clinical studies where high throughput is critical, researchers often pivot toward more efficient algorithms such as Non-rigid ICP [1] or functional maps [23]. However, these faster methods frequently introduce mesh tearing or triangle inversions, which can significantly undermine the reliability of downstream morphometric analyses.

Contributions. We propose a learning-free registration method that removes the need for extensive training data or human supervision, while preserving the topological guarantees of Large Deformation Diffeomorphic Metric Mapping (LDDMM) [4,34]. Compared to the canonical point-based LDDMM implementation Deformetrica [5], we achieve a $\sim 40\times$ speedup and better Chamfer and Hausdorff distances. This efficiency stems from three ingredients novel to point-based LDDMM: (i) an alternate minimization scheme inspired by the Majorize-Minimize framework and the Demons algorithm [33,36] to decouple matching and regularization, evaluated alongside a normalized kernel matrix (4); (ii) an analytic 3×3 Block-Jacobi preconditioner with closed-form eigenvectors; and (iii)

early termination of the inner conjugate gradient loop at tolerance $r/10$, where r is the current kernel radius. Combined, these reduce the inner CG loop from 20–30 iterations to 1–5. Our framework couples this fast quadratic step with a diffeomorphic shooting mechanism [11,17] to ensure topological correctness without the high overhead of traditional LDDMM. The resulting method is robust and versatile; with default parameters, it accurately registers varied anatomical geometries (Fig. 1). We demonstrate the scalability of our approach by generating a statistical shape model from 978 vertebrae in 17 minutes on a typical workstation. Finally, we provide a portable implementation of our pipeline within the Scikit-Shapes Python library ([scikit-shapes.github.io](https://github.com/scikit-shapes/scikit-shapes)) under the permissive MIT license.

2 Methods

Our method takes as input a source and a target shape identified with 3D point clouds $\mathcal{S} = (x_1, \dots, x_N)$ and $\mathcal{T} = (y_1, \dots, y_M)$, respectively. We adopt an iterative refinement strategy [1,22] comprising three primary phases:

1. **Matching.** Identification of point-wise correspondences ($x_i \leftrightarrow y_j$).
2. **Regularization.** Integration of correspondences into a globally smooth displacement field.
3. **Deformation.** Update of the source point positions x_i to reduce the residual distance to the target $\mathcal{T}$.

Matching. While advanced coupling algorithms based on optimal transport theory [30] or functional maps [23] are becoming increasingly popular, we found an affordable closest point matching scheme suitable for our purposes. For every source point x_i , we compute a feature vector:

$$f_i = [x_i, \gamma \text{FPFH}(x_i)] \in \mathbb{R}^{3+33}, \quad (1)$$

where the Fast Point Feature Histograms (FPFH) [25] denote local descriptors in dimension 33 which are invariant to rigid transformations, and $\gamma \geq 0$ is a scaling weight chosen to ensure that the collections of N coordinate vectors x_i and rescaled FPFH vectors have the same standard deviations. Likewise, we use the same scaling factor γ to compute a feature vector g_j in $\mathbb{R}^{36}$ for every target point y_j . Then, we compute the forward correspondence map $\sigma : \{1, \dots, N\} \rightarrow \{1, \dots, M\}$ and its backward counterpart τ with a nearest-neighbor search:

$$\sigma(i) = \arg \min_{j \in \{1, \dots, M\}} \|f_i - g_j\|^2, \quad \tau(j) = \arg \min_{i \in \{1, \dots, N\}} \|f_i - g_j\|^2. \quad (2)$$

For every source point x_i , consider the set $\mathcal{T}_i = \{y_j \in \mathcal{T} \mid \tau(j) = i\}$ of target points whose closest source neighbor is x_i . If it is empty, we simply define the effective target position z_i of the point x_i as its nearest target point $y_{\sigma(i)}$. But if $\mathcal{T}_i$ is not empty, we compute its barycenter t_i and use an effective target position $z_i = (1 - \kappa) y_{\sigma(i)} + \kappa t_i$, where $\kappa = 0.7$ is a weight parameter that we choose to promote a full coverage of the target shape $\mathcal{T}$.

Regularization. Since the closest point projections σ and τ may introduce topological artifacts, we must regularize the vector field of point displacements $u_i = z_i - x_i$. To this end, we solve the optimization problem:

$$v^* = \arg \min_{v \in \mathbb{R}^{N \times 3}} \underbrace{\lambda \cdot v^\top K^{-1} v}_{\text{Reg}(v)} + \underbrace{\sum_{i=1}^N (v_i - u_i)^\top L_i (v_i - u_i)}_{\text{Loss}(x+v)} , \quad (3)$$

where $\lambda = 1$ is a parameter that controls the goodness of fit, $L_1, \dots, L_N$ are symmetric, positive-definite 3×3 matrices that we define in the next paragraph and K is a symmetric, positive-definite $3N \times 3N$ matrix that is often interpreted as a covariance matrix or as a cometric in the shape space literature [18,11]. Standard choices for the covariance matrix K include the scale-invariant Thin Plate Spline kernel matrix [6], or the Gaussian kernel matrix used by the Coherent Point Drift and LDDMM methods [22,5]. Following recent works on diffusion operators [17], we use a *normalized* kernel matrix $K \in \mathbb{R}^{3N \times 3N}$ made up of 3×3 blocks:

$$K_{ij} = s_i s_j k(x_i, x_j) I_3 , \quad (4)$$

where I_3 denotes the 3×3 identity block, $k(x_i, x_j) = \exp(-\|x_i - x_j\|/r)$ is an exponential kernel of radius $r > 0$ and $s_i = 1/(\sum_{j=1}^N k(x_i, x_j))^{1/2}$ is a positive scaling factor.

Loss Functions. If the source and target shapes correspond to raw 3D point clouds, we can simply minimize a squared Euclidean error and use identity blocks $L_i = \frac{1}{N} I_3$. For surface meshes, however, this penalty could promote overfitting to the specific triangulation of the target surface. Assuming that we have access to unit point normals $\mathbf{n}_i$ and $\mathbf{m}_j$ in $\mathbb{R}^3$ for the points x_i and y_j , respectively, we thus prefer the plane-to-plane penalty [28] induced by the 3×3 blocks:

$$L_i = \frac{1}{N} [(\mathbf{n}_i \mathbf{n}_i^\top + \mathbf{m}_{\sigma(i)} \mathbf{m}_{\sigma(i)}^\top) + \alpha \cdot (2 I_3 - \mathbf{n}_i \mathbf{n}_i^\top - \mathbf{m}_{\sigma(i)} \mathbf{m}_{\sigma(i)}^\top)] \quad (5)$$

that use a scaling parameter $\alpha = 0.1$ to penalize normal displacements more strongly than sliding along the tangential planes of both surfaces. Likewise, we can use tangent vectors to define curve-to-curve penalties or rely on plane-to-point losses to match surface meshes onto raw point clouds.

Efficient Linear Solver. If L denotes the block-diagonal matrix obtained by concatenating the 3×3 blocks $L_1, \dots, L_N$, the solution v^* of Eq. (3) then reads:

$$v^* = K p^* , \quad \text{with} \quad (K + \lambda L^{-1}) p^* = u . \quad (6)$$

Since $(K + \lambda L^{-1})$ is symmetric positive-definite, we use a conjugate gradient solver to compute the optimal vector field $p^* \in \mathbb{R}^{N \times 3}$. In order to avoid wasteful computations, we make two crucial observations. First, since we are working with deformations at scale $r > 0$, we may stop the conjugate gradient loop as soon as the average norm of the residual vector field $[(K + \lambda L^{-1})p - u]$ is smaller than

a fixed tolerance $r/10$. Second, we can use a Block-Jacobi preconditioner P^{-1} to improve the condition number of the system. We propose to approximate the kernel matrix K by its diagonal blocks $K_{ii} = s_i^2 I_3$ and invert the 3×3 blocks L_i of L analytically. For the plane-to-plane penalty of Eq. (5), the inverse of each local preconditioner block $P_i = s_i^2 I_3 + \lambda L_i^{-1}$ then reads:

$$P_i^{-1} = \sum_{k=1}^3 \frac{1}{\mu_k} \frac{e_k e_k^\top}{\|e_k\|^2}, \quad (7)$$

with eigenvectors $e_1 = \mathbf{n}_i \times \mathbf{m}_{\sigma(i)}$, $e_2 = \mathbf{n}_i + \mathbf{m}_{\sigma(i)}$, $e_3 = \mathbf{n}_i - \mathbf{m}_{\sigma(i)}$ and corresponding eigenvalues $\mu_1 = s_i^2 + \frac{\lambda}{2\alpha}$, $\mu_2 = s_i^2 + \frac{\lambda}{1+\alpha+(1-\alpha)\mathbf{n}_i^\top \mathbf{m}_{\sigma(i)}}$, and $\mu_3 = s_i^2 + \frac{\lambda}{1+\alpha-(1-\alpha)\mathbf{n}_i^\top \mathbf{m}_{\sigma(i)}}$. In practice, these two remarks allow us to break the conjugate gradient loop after 1 to 5 iterations, instead of the 20 to 30 iterations required by an off-the-shelf solver to reach numerical precision. We rely on the KeOps library to implement products with the kernel matrix K on the GPU, with optimized run times and a linear memory footprint [7].

Deformation. While the optimal vector field $v^* \in \mathbb{R}^{N \times 3}$ provides a smooth direction for deformation, updating source point positions via a single step $x_i \leftarrow x_i + v_i^*$ often leads to folding artifacts and self-intersections. To ensure topology preservation throughout the registration process, we employ an LDDMM-like shooting method by integrating the Hamiltonian flow:

$$\begin{cases} \dot{q}(t) = K_{q(t)} p(t) \\ \dot{p}(t) = -\frac{1}{2} \nabla_q (p(t)^\top K_{q(t)} p(t)) \end{cases} \quad (8)$$

on point clouds $q(t) = (q_1(t), \dots, q_N(t)) \in \mathbb{R}^{N \times 3}$ and momentum vector fields $p(t) = (p_1(t), \dots, p_N(t)) \in \mathbb{R}^{N \times 3}$. We initialize the system at $q(0) = (x_1, \dots, x_N)$ and set the initial momentum $p(0) = p^*$, corresponding to the solution of Eq. (6).

The final deformed positions are obtained by integrating the flow to time $t = 1$, such that $x_i \leftarrow q_i(1)$. Crucially, because the exponential kernel is globally Lipschitz, the resulting velocity field is uniformly Lipschitz; by the Cauchy-Lipschitz theorem, this guarantees unique ODE solutions and ensures the deformation is a topology-preserving homeomorphism. To mitigate the contraction-expansion artifacts inherent in standard LDDMM geodesics [20], we apply a normalization to the kernel matrix as detailed in [11, Fig. 5.15] or [17, Fig. 5] and expressed in Eq. (4). This bounds the spectrum of K in $[0, 1]$, strictly controlling the kinetic energy ($\|v\|^2 = p^\top K p \leq \|p\|^2$) and ensuring that deformation velocities do not explode in regions with high sampling density. Furthermore, to guarantee numerical stability and strict topological consistency, we employ an adaptive Euler integrator. We strictly constrain the step size dt so that point displacements $\|q_i(t+dt) - q_i(t)\|$ never exceed the local sampling length ℓ , ensuring each step is too small to invert a source triangle.

Affine Alignment and Multiscale Heuristics. Our iterative matching, regularization, and deformation process is designed to drive the source point cloud $\mathcal{S}$ toward the target shape $\mathcal{T}$. To ensure convergence to corresponding anatomical structures, we employ a two-stage initialization strategy: a robust affine pre-alignment followed by a multiscale annealing heuristic. Both stages are fully automatic and require no manual landmarks or template-specific tuning.

First, we perform a rigid alignment of $\mathcal{S}$ onto $\mathcal{T}$ using a RANSAC-ICP (Iterative Closest Point) algorithm informed by FPFH descriptors. To account for inter-subject variability, we subsequently apply anisotropic scaling along each canonical axis, ensuring that the bounding boxes of the source and target geometries coincide. This normalization procedure effectively mitigates initial pose and scale discrepancies.

To further refine the registration, we introduce a sequence of four logarithmically spaced kernel radii $r_1 > r_2 > r_3 > r_4$, bounded by $r_{\text{init}} = r_1$ (set to the source shape radius) and $r_{\text{final}} = r_4$ (set to the sampling length ℓ). Starting from the affine-registered output, we perform four iterations of the “match-regularize-deform” registration loop at each scale. During the first two scales (r_1, r_2), we leverage FPFH features to guide the matching. At the final two scales (r_3, r_4), we transition to a purely spatial matching strategy based on 3D point coordinates, which enhances precision. In total, our approach performs 16 iterations. The computational complexity of each iteration is governed by the $\mathcal{O}(N^2)$ runtime of the matrix-vector products associated with the kernel matrix K within the conjugate gradient solver and the Hamiltonian flow.

Background. We frame our method as an efficient, Gauss-Newton-like approach to Large Deformation Diffeomorphic Metric Mapping (LDDMM). While Gauss-Newton-Krylov schemes have been well-established for diffeomorphic registration on dense voxel grids [19], current state-of-the-art implementations for point-based representations typically rely on extensive quasi-Newton iterations to converge toward optimal deformations [5]. By conducting a rigorous analysis of the optimization landscape, we have developed a strategy that reduces the complexity of this fundamental processing step by an order of magnitude, maintaining high levels of robustness and accuracy while significantly accelerating the registration pipeline.

3 Experiments

Datasets. To evaluate the robustness and versatility of our method, we conduct experiments on two distinct medical imaging datasets. First, we perform vertebral registration, a task characterized by high morphological variability across the cervical, thoracic, and lumbar regions. We use the VerSe dataset, a comprehensive collection of 374 CT scans from 355 patients, and adhere to the standardized training, validation and test partitions [29]. We process the expert-annotated Nifti masks through a dedicated pipeline: we apply anisotropic Gaussian smoothing to suppress voxel-based staircasing artifacts, followed by boundary surface

extraction using the PyVista library [32]. To optimize computational efficiency for point-based registration methods, we decimate the resulting surfaces from 50k to 10k points using quadric error metrics [13]. As source templates for these target vertebrae, we use standardized anatomical meshes from the BodyParts3D database [21], refined with Loop subdivision and uniform remeshing [35].

To demonstrate topological versatility, we also register 54 pairs of lung vascular trees from the Lung250M-4B dataset [10], sampled with around 40k points each. This experiment shows that our pipeline can handle complex branching structures without domain-specific heuristics.

General Setup & Baselines. We run all experiments on a Dell workstation equipped with 64 Gb of RAM, an Intel Xeon W-2255 CPU (3.70GHz) and an Nvidia RTX 3090 GPU. We benchmark against two established registration frameworks: the LDDMM-based **Deformetrica** and the deep-learning architecture **VoxelMorph**. We exclude the widely-used ANTs toolkit [2], as its SyN algorithm required over ten minutes per vertebral pair, proving computationally prohibitive for our cohort. To ensure a fair comparison, all baselines were initialized with our affine pre-alignment and rigorously tuned on the VerSe dataset to optimize the trade-off between registration accuracy and topological preservation. More precisely:

- For **our method**, we used kernel radii ranging from $r_{\text{init}} = 10$ mm to $r_{\text{final}} = 4$ mm, corresponding to the expected scale of the largest displacements and to the spatial resolution of the coarsest source volumes, respectively.
- **VoxelMorph** was implemented in PyTorch [24]. Due to architectural and memory constraints, each vertebra was resampled as a $96 \times 96 \times 96$ segmentation mask. To optimize boundary alignment, we tested both a Dice loss on binary masks and a Mean Squared Error loss on exact Euclidean distance fields. Both struggled with high-fidelity alignment; we report the distance-field results for their superior gradient stability. The model was trained for 100 epochs (Adam optimizer, learning rate $2 \cdot 10^{-4}$) with an ℓ_2 spatial gradient penalty to ensure smooth deformations.
- **Deformetrica** was driven by an L-BFGS optimizer and a varifold loss [8]. The deformation and matching kernel widths were both set to 5 mm, maintaining consistency with the final scale of our proposed method. The noise standard deviation was fixed at 0.7.

Evaluation. We evaluate performance using three surface metrics: **Chamfer Distance (CD)** for global alignment, **95th percentile Hausdorff Distance (HD95)** for worst-case local misalignment, and the **variance of the surface log-Jacobian** for deformation regularity. For the lung registration task, accuracy is measured using the 100 landmarks per tree provided in Lung250M-4B.

As shown in Table 1, our framework combines the speed of deep learning with the precision of classical LDDMM, achieving sub-millimeter surface alignment in 1.05 ± 0.10 s. While VoxelMorph provides the fastest inference (15.2 ms),




Fig. 2. Registration of a C7 vertebral template (a) onto a segmented sample (d). Unlike Deformetrica (b), which induces tangential compression manifesting as excessive stretching of the mesh triangulation, our method (c) achieves high geometric fidelity with minimal distortion. By enforcing strict topological constraints, our model preserves the structural integrity of the source, allowing for the accurate localization of the transverse foramina even in the presence of segmentation errors in the target.

Table 1. Quantitative comparison on 757 vertebrae from the VerSe dataset. Our method achieves the highest surface alignment accuracy and preserves topological structures, without requiring access to a training dataset and the related computational overhead.

Method	Chamfer (mm) ↓	HD95 (mm) ↓	Log-Jacobian Variance ↓	Training Time ↓	Inference Time ↓
VOXELMORPH	2.39 ± 1.20	7.68 ± 4.58	0.12 ± 0.10	46 min	15.2 ± 0.05 ms
DEFORMETRICA	0.24 ± 0.12	0.74 ± 1.61	0.38 ± 0.47	0 s	43.50 ± 11.6 s
OURS	0.22 ± 0.08	0.64 ± 0.40	0.16 ± 0.14	0 s	1.05 ± 0.10 s

its low log-Jacobian variance (0.12 ± 0.10) is a byproduct of under-deformation; the network struggles to align complex vertebral features (yielding a HD95 of 7.68 ± 4.58 mm) and predicts overly conservative displacements. In contrast, our method successfully captures these large deformations while maintaining excellent topological integrity (log-Jacobian variance of 0.16 ± 0.14). As illustrated in Fig. 2, whereas Deformetrica struggles with tangential compression artifacts, our approach achieves smooth, high-fidelity alignments, strictly outperforming Deformetrica while operating roughly $40\times$ faster. Crucially, while traditional CPU-bound frameworks like ANTs SyN would require days of continuous computation to register our test cohort, our approach requires only a single consumer-grade GPU (Nvidia RTX 3090) to process the same data in minutes.

Going beyond surface meshes, and as a proof-of-concept, our pipeline successfully registered the complex vascular trees of the Lung250M-4B dataset sampled at 40k points in just 20 seconds per pair, achieving a mean landmark distance of 3.08 mm without any specialized pre-processing or extensive parameter tuning.

Atlas Estimation. We demonstrate our framework’s scalability by computing a Statistical Shape Model (SSM) from 978 VerSe test vertebrae in just 17 minutes

on a single workstation. Using the refined BodyParts3D templates introduced above, we compute mean anatomical vertebrae and their modes of variation. Because our method strictly regularizes local expansions and compressions, the resulting atlas preserves high-frequency structural details such as transverse and spinous processes.

Limitations and Future Works. While our method is orders of magnitude faster than other topology-preserving baselines, our run times remain quadratic due to our brute-force implementation of kernel matrix-vector products with KeOps [7]. Future work will explore approximate schemes like the Fast Gauss Transform, giga-scale GPU matrix-vector multiplication [16] and other similar techniques [26] to mitigate quadratic complexity. Additionally, extending this framework to accommodate partial overlaps, quantifying atlas template bias in heterogeneous cohorts, and broadening the anatomical scope of our evaluations remain critical objectives.

Acknowledgments. This study was funded by the France 2030 program managed by the Agence Nationale de la Recherche (ANR-23-IACL-0008, PR[AI]RIE-PSAI).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Amberg, B., Romdhani, S., Vetter, T.: Optimal step nonrigid ICP algorithms for surface registration. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1–8 (2007)
2. Avants, B.B., Tustison, N., Song, G., et al.: Advanced normalization tools (ANTS). *Insight j* **2**(365), 1–35 (2009)
3. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: VoxelMorph: a learning framework for deformable medical image registration. *IEEE Transactions on Medical Imaging* **38**(8), 1788–1800 (2019)
4. Beg, M.F., Miller, M.I., Trounev, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision* **61**(2), 139–157 (2005)
5. Bône, A., Louis, M., Martin, B., Durrleman, S.: Deformetrica 4: an open-source software for statistical shape analysis. In: International workshop on shape in medical imaging. pp. 3–13. Springer (2018)
6. Bookstein, F.L.: Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on pattern analysis and machine intelligence* **11**(6), 567–585 (2002)
7. Charlier, B., Feydy, J., Glaunes, J.A., Collin, F.D., Durif, G.: Kernel operations on the GPU, with autodiff, without memory overflows. *Journal of Machine Learning Research* **22**(74), 1–6 (2021)
8. Charon, N., Trounev, A.: The varifold representation of nonoriented shapes for diffeomorphic registration. *SIAM journal on Imaging Sciences* **6**(4), 2547–2580 (2013)

9. Durrleman, S., Pennec, X., Trounev, A., Ayache, N.: Statistical models of sets of curves and surfaces based on currents. *Medical Image Analysis* **13**(5), 793–808 (2009)
10. Falta, F., Großbröhmer, C., Hering, A., Bigalke, A., Heinrich, M.: Lung250M-4B: a combined 3D dataset for CT-and point cloud-based intra-patient lung registration. *Advances in Neural Information Processing Systems* **36**, 54819–54832 (2023)
11. Feydy, J.: Geometric data analysis, beyond convolutions. Ph.D. thesis, Université Paris-Saclay (2020)
12. Forsberg, D., Lundström, C., Andersson, M., Knutsson, H.: Model-based registration for assessment of spinal deformities in idiopathic scoliosis. *Physics in Medicine and Biology* **59**(2), 311–326 (2014)
13. Garland, M., Heckbert, P.S.: Surface simplification using quadric error metrics. In: *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. pp. 209–216 (1997)
14. Hansen, L., Heinrich, M.P.: Revisiting iterative highly efficient optimisation schemes in medical image registration. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 203–212. Springer (2021)
15. Heimann, T., Meinzer, H.P.: Statistical shape models for 3D medical image segmentation: A review. *Medical Image Analysis* **13**(4), 543–563 (2009)
16. Hu, R., Chau, S.L., Sejdinovic, D., Glaunès, J.: Giga-scale kernel matrix-vector multiplication on gpu. *Advances in Neural Information Processing Systems* **35**, 9045–9057 (2022)
17. Kessler, N., Magnet, R., Feydy, J.: Normalizing diffusion kernels with optimal transport (2026)
18. Lüthi, M., Gerig, T., Jud, C., Vetter, T.: Gaussian process morphable models. *IEEE transactions on pattern analysis and machine intelligence* **40**(8), 1860–1873 (2017)
19. Mang, A., Gholami, A., Davatzikos, C., Biros, G.: CLAIRE: A distributed-memory solver for constrained large deformation diffeomorphic image registration. *SIAM Journal on Scientific Computing* **41**(5), C548–C584 (2019)
20. Micheli, M., Michor, P.W., Mumford, D.: Sectional curvature in terms of the cometric, with applications to the Riemannian manifolds of landmarks. *SIAM Journal on Imaging Sciences* **5**(1), 394–433 (2012)
21. Mitsuhashi, N., Fujieda, K., Tamura, T., Kawamoto, S., Takagi, T., Okubo, K.: BodyParts3D: 3D structure database for anatomical concepts. *Nucleic acids research* **37**(suppl_1), D782–D785 (2009)
22. Myronenko, A., Song, X.: Point set registration: Coherent point drift. *IEEE transactions on pattern analysis and machine intelligence* **32**(12), 2262–2275 (2010)
23. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)* **31**(4), 1–11 (2012)
24. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in PyTorch (2017)
25. Rusu, R.B., Blodow, N., Beetz, M.: Fast point feature histograms (FPFH) for 3D registration. In: *2009 IEEE International Conference on Robotics and Automation*. pp. 3212–3217 (2009)
26. Rux, N., Hertrich, J., Neumayer, S.: Numerical methods for kernel slicing. *arXiv preprint arXiv:2510.11478* (2025)
27. Sassen, J., Schumacher, H., Rumpf, M., Crane, K.: Repulsive shells. *ACM Transactions on Graphics* **43**(4), 1–22 (2024)

28. Segal, A., Haehnel, D., Thrun, S.: Generalized-ICP. In: *Robotics: Science and Systems*. vol. 2, p. 435 (2009)
29. Sekuboyina, A., Hussein, M.E., Bayat, A., Löffler, M., Liebl, H., Li, H., Tetteh, G., Kukačka, J., Payer, C., Štern, D., et al.: VerSe: a vertebrae labelling and segmentation benchmark for multi-detector CT images. *Medical image analysis* **73**, 102166 (2021)
30. Shen, Z., Feydy, J., Liu, P., Curiale, A.H., San Jose Estepar, R., San Jose Estepar, R., Niethammer, M.: Accurate point cloud registration with robust optimal transport. *Advances in Neural Information Processing Systems* **34**, 5373–5389 (2021)
31. Sigal, I.A., Hardisty, M.R., Whyne, C.M.: Mesh-morphing algorithms for specimen-specific finite element modeling. *Journal of Biomechanics* **41**(7), 1381–1389 (2008)
32. Sullivan, C., Kaszynski, A.: PyVista: 3D plotting and mesh analysis through a streamlined interface for the Visualization Toolkit (VTK). *Journal of Open Source Software* **4**(37), 1450 (2019)
33. Thirion, J.P.: Image matching as a diffusion process: an analogy with maxwell’s demons. *Medical image analysis* **2**(3), 243–260 (1998)
34. Vaillant, M., Glaunès, J.: Surface matching via currents. In: *Information Processing in Medical Imaging (IPMI)*. pp. 381–392. Springer (2005)
35. Valette, S., Chassery, J.M.: Approximated centroidal voronoi diagrams for uniform polygonal mesh coarsening. In: *Computer graphics forum*. vol. 23, pp. 381–389. Wiley Online Library (2004)
36. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Symmetric log-domain diffeomorphic registration: A demons-based approach. In: *International conference on medical image computing and computer-assisted intervention*. pp. 754–761. Springer (2008)