Correspondence-Free Non-Rigid Point Set Registration Using Unsupervised Clustering Analysis

Mingyang Zhao[1]   Jingen Jiang[2]   Lei Ma[3*]   Shiqing Xin[2]   Gaofeng Meng[1,4,5]   Dong-Ming Yan[4,5*]

[1]CAIR, HKISI, CAS [2]Shandong University [3]Peking University [4]MAIS, CASIA [5]UCAS

Figure 1. Non-rigid registration on 3D point sets. The blue and gray models represent the source and target point clouds, respectively, while the yellow models are our registration results. Our method achieves successful registrations even for shapes with challenging deformations.

Abstract

This paper presents a novel non-rigid point set registration method that is inspired by unsupervised clustering analysis. Unlike previous approaches that treat the source and target point sets as separate entities, we develop a holistic framework where they are formulated as clustering centroids and clustering members, separately. We then adopt Tikhonov regularization with an ℓ_1-induced Laplacian kernel instead of the commonly used Gaussian kernel to ensure smooth and more robust displacement fields. Our formulation delivers closed-form solutions, theoretical guarantees, independence from dimensions, and the ability to handle large deformations. Subsequently, we introduce a clustering-improved Nyström method to effectively reduce the computational complexity and storage of the Gram matrix to linear, while providing a rigorous bound for the low-rank approximation. Our method achieves high accuracy results across various scenarios and surpasses competitors by a significant margin, particularly on shapes with substantial deformations. Additionally, we demonstrate the versatility of our method in challenging tasks such as shape transfer and medical registration. [Code release]

1. Introduction

Non-rigid point set registration is to optimize a non-linear displacement field that accurately aligns one geometric shape with another. Due to its fundamental importance, non-rigid registration plays a dominant role in a wide range of applications, such as scene reconstruction [35, 52, 54], pose tracking [34], animation [38], deformable shape manipulation and editing [42], and so on.

However, given two point sets, one acting as the source and the other as the target, non-rigid registration presents a highly ill-posed and much more complex challenge compared to the rigid counterpart. This increased complexity is primarily attributed to the additional freedom of deformations allowed in non-rigid registration, especially when dealing with shapes that exhibit large deformations (Fig. 1).

To enhance the registration quality for shapes undergoing large deformations, numerous pioneering methods have been actively researched. Rather than directly optimizing the registration process, these methods usually employ a *two-step* approach [30, 36, 47, 48]. First, they perform *shape matching* by identifying corresponding points between the source and target shapes without considering geometry deformations. Then, they estimate the alignment transformation based on the established correspondences via off-the-shelf registration techniques. While there has

---

been significant attention and research dedicated to the initial shape matching stage, the exploration of *direct registration* methods for handling large deformations, without relying on shape matching, is comparatively limited and poses substantial challenges [20].

In this work, we address the problem of non-rigid point set registration without correspondences, with a specific emphasis on point sets exhibiting large deformations. To overcome this challenge, we present a fresh perspective and introduce a novel method. Our approach reformulates the non-rigid deformation process as an *unsupervised clustering problem* within the context of machine learning. Unlike previous approaches that treat the two point sets as *separate* entities, we consider them as *integral* parts of a whole.

Concretely, we designate the source point set as the clustering centroids, while the target one as the clustering samples. This *holistic* treatment enables us to leverage the interplay between these two sets. Then the dynamic optimization and update of the clustering centroids correspond to the underlying deformation of the source shape. We highlight the advantages of our novel registration function, which is built on clustering analysis, from both information theory and convex optimization perspectives. Furthermore, we provide *closed-form* solutions to our objective function during each iteration, which enables fast and efficient implementations. We introduce a sparsity-induced *Laplacian kernel* (ℓ_1-norm) in the Tikhonov regularization to ensure that the displacement field of clustering centroids remains as smooth as possible. This differs from the commonly used Gaussian kernel and exhibits higher robustness, as demonstrated by experimental results. Additionally, we leverage clustering analysis to adopt the improved Nyström low-rank approximation [51], which reduces the computational complexity and storage requirements of the Gram matrix to linear. Meanwhile, we give a rigorous proof of the approximation error bound associated with the Laplacian kernel.

Our method is independent of spatial dimensions, allowing us to evaluate and compare its performance in both 2D and 3D settings. The experimental results demonstrate the superiority of our method compared to baselines by a large margin. This is particularly evident in scenarios involving large deformations, such as shape transfer and medical data registration.

Our contributions can be summarized as follows:

- We propose a novel and correspondence-free method for non-rigid point set registration, utilizing unsupervised clustering analysis. The method achieves impressive results across various settings and mitigates the challenge without explicit correspondences.
- We incorporate the Laplacian kernel function for robust displacement regularization and provide a rigorous theoretical analysis to prove the approximation error bound of the Nyström low-rank method.

- Our method is dimension-independent, offering closed-form solutions during optimization, and significantly improves performance in handling large deformations.

2. Related Work

We review the work that is closely aligned with ours. Readers are directed to [14, 41] for comprehensive studies.

Non-rigid registration. Differing from shape matching that focuses on finding inlier correspondences, non-rigid registration aims to optimize the displacement field. Various pioneering algorithms employ an optimization paradigm that minimizes both the data and penalty terms simultaneously. Amberg *et al*. [1] extended the rigid iterative closest point algorithm [4] to non-rigid settings, while Yao *et al*. [47, 48] recently improved non-rigid ICP regarding both accuracy and efficiency through deformation graph optimization. *Coherent Point Drift* (CPD) [33] and GMM [23] developed probabilistic frameworks by minimizing the negative logarithm likelihood function to enhance the robustness for non-rigid point set registration. Ma *et al*. [30] further incorporated the shape context descriptor [2] to establish shape correspondences for better 2D registration. Hirose [20, 21] recently formulated CPD in a Bayesian setting, which effectively overcomes CPD's limitations and delivers impressive results.

With the advancement of deep learning, neural network-based methods have also been proposed for non-rigid point set registration [22, 25, 36]. Most of them utilize neural networks to extract features for point correspondences and then apply classical methods such as non-rigid ICP for registration. Instead of focusing on shape matching and heavily rely on data annotations, our method is unsupervised and reasons from a case-by-case geometric perspective. This allows us to achieve faithful registrations that are more generalizable to unknown categories.

Deformation representation. The representation of the deformation field is a key component in non-rigid registration. Several existing works are based on thin plate spline functions [12, 23, 37], which can be viewed as a regularization of the second-order derivatives of the transformations [33]. Another line of researches utilize kernel functions or a reproducing kernel Hilbert space to describe the deformation field [28, 29, 33, 43]. However, many of these methods are limited to the Gaussian kernel due to the reliance on fast Gauss transform [18]. Recently, the *Multi-Layer Perception* (MLP) network has been employed to represent the deformation field by mapping input coordinates to signal values [24, 26, 35] and the deformation degree is controlled by frequencies. These methods have shown promising results in dynamical reconstruction and scene flow estimation, which are typically considered less challenging tasks compared to dealing with large deformations.

3. Preliminaries on Clustering Analysis

As one of the representative unsupervised learning frameworks, clustering analysis plays a fundamental role in various scientific research domains [46]. The pioneering work [27, 53] explored clustering metrics for rigid point cloud registration. In contrast, we distinguish ourselves by addressing a more challenging non-rigid problem, which we have completely reformulated as a clustering process with a different objective function. We present a concise overview on two commonly used clustering approaches: *fuzzy clustering* and *Elkan k-means clustering* analysis.

3.1. Fuzzy Clustering Analysis

Given a dataset $\mathbf{X} = \{\boldsymbol{x}_i \in \mathbb{R}^n\}_{i=1}^M$, fuzzy clustering analysis solves the following problem:

$$\min_{\mathbf{U}, \mathbf{V}} \sum_{j=1}^C \sum_{i=1}^M (u_{ij})^r ||\boldsymbol{x}_i - \boldsymbol{v}_j||_2^2, s.t. \sum_{j=1}^C u_{ij} = 1, u_{ij} \in [0, 1], \quad (1)$$

where $\mathbf{U} = [u_{ij}]_{M \times C} \in \mathbb{R}^{M \times C}$ is the *fuzzy membership degree matrix*, $\mathbf{V} = \{\boldsymbol{v}_j \in \mathbb{R}^n\}_{j=1}^C$ is the set of *clustering centroids* consisting of $C \in \mathbb{Z}_+$ classes, and $r \in [1, +\infty)$ is the *fuzzy factor*, which controls the clustering fuzziness. To enhance the clustering performance on unbalanced datasets, Miyamoto *et al.* [32] proposed the inclusion of *cluster size controlling variables* $\boldsymbol{\alpha} = [\alpha_1, \cdots, \alpha_C] \in \mathbb{R}^C$ in Eq. (1), and thus classes with more samples may lead to higher fuzzy membership degree. Since Euclidean distance-based clustering algorithms are primarily suitable for spherical data, *Mahalanobis distance* is latter introduced to generalize the fuzzy clustering analysis to accommodate ellipsoidal structures [19]. Recently, [9] combined the merits of previous fuzzy clustering approaches and developed a novel clustering framework based on the $\ell_{2,p}$ norm, which achieves appealing results on a set of clustering analysis tasks:

$$\min_{\mathbf{U}, \mathbf{V}, \boldsymbol{\Sigma}, \boldsymbol{\alpha}} \sum_{j=1}^C \sum_{i=1}^M u_{ij} ||\boldsymbol{\Sigma}_j^{-\frac{1}{2}} (\boldsymbol{x}_i - \boldsymbol{v}_j)||_2^p + u_{ij} \log |\boldsymbol{\Sigma}_j| + \lambda u_{ij} \log \frac{u_{ij}}{\alpha_j},$$

$$s.t. |\boldsymbol{\Sigma}_j| = \theta_j, \sum_{j=1}^C u_{ij} = 1, \sum_{j=1}^C \alpha_j = 1, u_{ij}, \alpha_j \in [0, 1].$$

$$(2)$$

where $\lambda \in \mathbb{R}^+$ is a regularization parameter, and $\boldsymbol{\Sigma}_j \in \mathbb{S}_{++}^n \triangleq \{\mathbf{A} \in \mathbb{R}^{n \times n} | \boldsymbol{x}^T \mathbf{A} \boldsymbol{x} > 0, \forall \boldsymbol{x} \in \mathbb{R}^n\}$ denotes the covariance matrix of the j-th class, with the corresponding determinant equivalent to $|\boldsymbol{\Sigma}_j| \in \mathbb{R}$. We explore the application of this clustering analysis framework to non-rigid point set registration and demonstrate its superior performance over previous registration approaches.

3.2. Elkan k-Means Clustering

In contrast to fuzzy clustering analysis, the k-means algorithm [31] has emerged as one of the most widely used clustering methods due to its simplicity. Elkan k-means clustering further introduced the *triangle inequality* into the k-means framework to avoid unnecessary distance calculations, which dramatically speeds up the primary k-means clustering process. More details of Elkan k-means clustering can be found in [15].

4. Proposed Method

Problem formulation. Given two point sets $\mathbf{X} = \{\boldsymbol{x}_i \in \mathbb{R}^n\}_{i=1}^M$ and $\mathbf{Y} = \{\boldsymbol{y}_j \in \mathbb{R}^n\}_{j=1}^N$, where $\mathbf{X}$ and $\mathbf{Y}$ are named as the target and the source, separately, the objective of non-rigid point set registration is to find the optimal deformation map $\mathcal{T}$ that minimizes the shape deviation between $\mathcal{T}(\mathbf{Y}) \triangleq \mathbf{Y} + \nu(\mathbf{Y})$ and $\mathbf{X}$, where ν represents the displacement filed acting on each source point $\boldsymbol{y}_j$.

4.1. Clustering-Induced Non-Rigid Registration

Observations. We notice that during the clustering process, the spatial position of clustering centroids $\mathbf{V}$ are dynamically updated until the distance between the centroids and their members is minimized. This dynamic process bears resemblance to the iterative update of each source point $\mathcal{T}(\boldsymbol{y}_j)$. Inspired by this, we propose to formulate non-rigid registration as an unsupervised clustering process. We consider $\mathbf{Y}$ as the clustering centroids and $\mathbf{X}$ as the clustering members. We customize Eq. (2) to optimize the overall clustering loss by

$$\min F(\mathbf{U}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}, \nu) = \sum_{j=1}^C \sum_{i=1}^M u_{ij} ||\boldsymbol{\Sigma}_j^{-\frac{1}{2}} (\boldsymbol{x}_i - (\boldsymbol{y}_j + \nu(\boldsymbol{y}_j)))||_2^2$$

$$+ u_{ij} \log |\boldsymbol{\Sigma}_j| + \lambda u_{ij} \log \frac{u_{ij}}{\alpha_j},$$

$$s.t. |\boldsymbol{\Sigma}_j| = \theta_j, \sum_{j=1}^C u_{ij} = 1, \sum_{j=1}^C \alpha_j = 1, u_{ij}, \alpha_j \in [0, 1].$$

$$(3)$$

Here we set $p = 2$ to ease the computation, which also ensures closed-form solutions as derived in the following.

Regularization. As in [20, 33], we incorporate *Tikhonov regularization* [17] to promote smoothness in the displacement field of clustering centroids. Thus, our objective function is optimized to find the optimal locations of clustering centroids as follows:

$$\min F(\mathbf{U}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}, \nu) + \zeta \mathcal{R}(\nu), \quad (4)$$

where ζ is a trade-off parameter. $\mathcal{R}(\cdot)$ is an operator that penalizes the high-frequency component of ν if we consider it in the Fourier domain, *i.e.*,

$$\mathcal{R}(\nu) = \int_{\mathbf{R}^n} d\mathbf{s} \frac{||\tilde{\nu}(\mathbf{s})||_2^2}{\tilde{K}(\mathbf{s})}. \quad (5)$$

$K(\mathbf{s})$ is a kernel function regarding the frequency variable $\mathbf{s}$, and $\tilde{f}$ indicates the Fourier transform of the function f.

4.2. Virtues of the Newly-Defined Function

We provide a theoretical analysis of Eq. (3) from both information theory and optimization perspectives. This analysis allows us to highlight the virtues of our newly introduced loss function for non-rigid point set registration.

Information theory view. We re-write $F(\mathbf{U}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}, \nu)$ as

$$\sum_{j=1,i=1}^{C,M} u_{ij}||\boldsymbol{\Sigma}_j^{-\frac{1}{2}}(\boldsymbol{x}_i-(\boldsymbol{y}_j+\nu(\boldsymbol{y}_j)))||_2^2+u_{ij}\log(\frac{|\boldsymbol{\Sigma}_j|}{\alpha_j^2})-\lambda H(\mathbf{U}),$$

where $H(\mathbf{U}) = -\sum_{j=1}^{C}\sum_{i=1}^{M} u_{ij}\log(u_{ij})$ is the entropy of $\mathbf{U}$. From the perspective of information theory, this *entropy regularization* term serves to push $F(\mathbf{U}, \boldsymbol{\alpha}, \boldsymbol{\Sigma}, \nu)$ towards $\mathbf{U}$ with a uniform distribution that makes $H(\mathbf{U})$ the maximal and thus drags $F(\mathbf{U}, \alpha, \boldsymbol{\Sigma}, \nu)$ away from the sparse $\mathbf{U}$. This not only enhances the smoothness of the feasible set, but also improves the computational stability during optimization, *i.e.*, avoiding $\lim_{u_{ij}\to 0}\log(u_{ij}) = -\infty$ [13].

Optimization view. Alternatively, from an optimization point of view, $u_{ij}\log(u_{ij})$ is a convex function in terms of u_{ij} with λ controlling the degree of convexity. Moreover, $u_{ij}\log(u_{ij})$ acts as a *barrier function* that restricts u_{ij} to the range of $[0, 1]$ and prevents it from taking values outside this range [49].

4.3. Closed-Form Solutions

Our method enables closed-form solutions for each variable during the optimization step as derived in the following.

Update of $\mathbf{U}$. We fix $\boldsymbol{\alpha}, \boldsymbol{\Sigma}, \nu$ and update $\mathbf{U}$, which becomes a convex optimization problem. Utilizing the Lagrangian multiplier and ignoring parameters that are irrelevant to $\mathbf{U}$, we obtain

$$\mathcal{L}(\mathbf{U}, \boldsymbol{\beta}) = \sum_{j=1}^{C}\sum_{i=1}^{M} u_{ij}||\boldsymbol{\Sigma}_j^{-\frac{1}{2}}(\boldsymbol{x}_i - (\boldsymbol{y}_j + \nu(\boldsymbol{y}_j)))||_2^2$$

$$+u_{ij}\log|\boldsymbol{\Sigma}_j| + \lambda u_{ij}\log\frac{u_{ij}}{\alpha_j} + \sum_{i=1}^{M}\beta_i(\sum_{j=1}^{C} u_{ij} - 1),$$

where $\boldsymbol{\beta} = \{\beta_i \in \mathbb{R}\}_{i=1}^{M}$ are the set of Lagrangian multipliers. By equating $\frac{\partial\mathcal{L}}{\partial\mathbf{U}} = 0$, we have

$$\mathbf{U} = (\text{diag}(\mathbf{A}\mathbf{1}_C))^{-1}\mathbf{A} \qquad (6)$$

Here $\mathbf{A} = \exp(-\mathbf{D}/\lambda)\text{diag}(\boldsymbol{\alpha}\odot|\boldsymbol{\Sigma}|)$, $\mathbf{D} = [d_{ij}]_{M\times C} \in \mathbb{R}^{M\times C}$ is a squared Euclidean distance matrix with $d_{ij} = ||\boldsymbol{\Sigma}_j^{-\frac{1}{2}}(\boldsymbol{x}_i - (\boldsymbol{y}_j + \nu(\boldsymbol{y}_j)))||_2^2$, $\exp(\cdot)$ is the element-wise exponential operator of matrices, $\text{diag}(\mathbf{z})$ is an operator that creates a square diagonal matrix with the vector $\mathbf{z}$ on its main diagonal, and $|\boldsymbol{\Sigma}| = [|\boldsymbol{\Sigma}_1|, \cdots, |\boldsymbol{\Sigma}_C|]^T \in \mathbb{R}^C$. $\mathbf{1}_C$ is the C-dimensional vector of all ones, and $\odot$ represents the element-wise Hadamard product of two matrices or vectors.
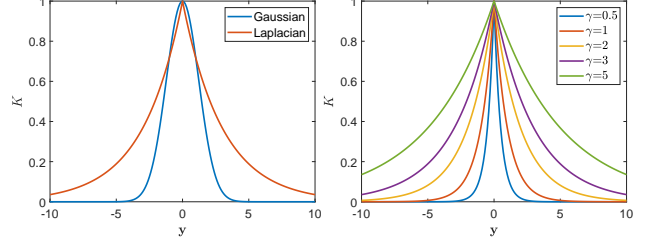

Figure 2. Left: Comparison between the Gaussian kernel and Laplacian kernel with the same bandwidth $\gamma = 3$, where the latter delivers considerably thicker lower and upper tails, indicating higher robustness. Right: Laplacian kernel with different γ.

Update of $\boldsymbol{\alpha}$. Likewise, the closed-form solution with respect to $\boldsymbol{\alpha}$ is $\boldsymbol{\alpha} = \frac{1}{M}\mathbf{U}^T\mathbf{1}_M$, which formally quantifies the clustering size for each class.

Update of $\boldsymbol{\Sigma}$. For simplicity, we relax each clustering centroid's covariance matrix to be isotropic, *i.e.*, $\boldsymbol{\Sigma}_j = \sigma^2\mathbf{I}$, where $\mathbf{I} \in \mathbb{R}^{n\times n}$ is the identity matrix. This ensures a closed-form solution to variance σ^2:

$$\sigma^2 = \frac{\text{tr}(\mathbf{X}^T\text{diag}(\mathbf{U}^T\mathbf{1}_M)\mathbf{X} - (2(\mathbf{U}\mathbf{X})^T + \mathbf{T}^T\text{diag}(\mathbf{U}\mathbf{1}_C))\mathbf{T})}{n \times M},$$

where $\text{tr}(\cdot)$ is the matrix trace operator.

Update of ν. By leveraging the *Riesz's representation theorem* [10], the closed-form solution to the regularization term ν can be expressed as

$$\nu(\boldsymbol{y}) = \sum_{j=1}^{C} c_j K(\boldsymbol{y}, \boldsymbol{y}_j) + \sum_{\eta=1}^{N} d_\eta\psi_\eta(\nu), \qquad (7)$$

where $\{c_j \in \mathbb{R}\}_{j=1}^{C}$ are the coefficient scalars, $K(\cdot, \cdot)$ is the kernel function defined in Eq. (5), and $\{\psi_\eta\}_{\eta=1}^{N}$ represent a set of basis in the N-dimensional null space of $\mathcal{R}(\nu)$, which is typically composed by a set of polynomials for most choices of the stabilizer $\mathcal{R}(\nu)$.

In contrast to previous approaches that commonly utilize a *Gaussian Radial Basis Function* (RBF) [30, 33], we adopt the sparsity-induced *Laplacian kernel* with the robust ℓ_1-norm to characterize the displacement field ν, *i.e.*,

$$K(\boldsymbol{y}_i, \boldsymbol{y}_j) = \exp(-\gamma\|\boldsymbol{y}_i - \boldsymbol{y}_j\|_1), \quad \gamma > 0 \qquad (8)$$

in which $\|\boldsymbol{y}_i - \boldsymbol{y}_j\|_1$ is the *Manhattan distance* between the two input vectors. Compared to the RBF kernel, the Laplacian kernel exhibits stronger robustness due to its considerably thicker tails, as illustrated in Fig. 2. We also validate this conclusion through subsequent experiments. Since the Laplacian kernel is positive definite, we obtain $\psi_\eta \equiv 0$ [17]. By evaluating $\nu(\boldsymbol{y})$ at $\mathbf{Y} = \{\boldsymbol{y}_j \in \mathbb{R}^n\}_{j=1}^{C}$, following [20, 33], the coefficient vector $\mathbf{c} = [c_1, c_2, \cdots, c_C]^T \in \mathbb{R}^C$ is recovered from the following linear system:

$$\mathbf{c} = (\mathbf{L} + \zeta\sigma^2\text{diag}(\mathbf{U}\mathbf{1}_C)^{-1})^{-1}(\text{diag}(\mathbf{U}\mathbf{1}_C)^{-1}\mathbf{U}\mathbf{X} - \mathbf{Y}), \quad (9)$$

where $\mathbf{L}$ is the Gram matrix with $l_{ij} = K(\boldsymbol{y}_i, \boldsymbol{y}_j)$. Therefore, the newly deformed shape $\mathbf{T}$ from the source point set $\mathbf{Y}$ becomes $\mathbf{T} = \mathcal{T}(\mathbf{Y}) = \mathbf{Y} + \mathbf{Lc}$.

4.4. Improved Nyström Low-Rank Approximation

The matrix inverse operation in Eq. (9) leads to a computational complexity of $O(C^3)$ and a memory requirement of $O(C^2)$. Previous approaches often employ the *fast Gauss transform* (FGT) [18] to reduce memory usage and accelerate computation. However, FGT is merely limited to the Gaussian kernel. To circumvent this issue, BCPD [20] combined the Nyström method [11] and the KD tree search [3] for acceleration. However, there are still two major issues that remain unresolved. (1) Due to the random sampling scheme used in BCPD, it is unclear how effective the Nyström approximation performs. (2) In order to address convergence issues when σ^2 becomes small, BCPD need to switch from Nyström approximation to KD tree search. This transition may affect the optimization trajectory.

To overcome these challenges, we opt to use *clustering analysis* instead of random sampling. Concretely, we first employ the fast Elkan k-means algorithm (Sec. 3.2) to partition $\mathbf{Y}$ into C' disjoint clusters $\mathbf{P}_i \subset \mathbf{Y}$, with the corresponding clustering centroids as $\{\boldsymbol{z}_i \in \mathbb{R}^n\}_{i=1}^{C'}$ ($C' \ll C$). Then, we adopt the improved Nyström approximation [51] for efficient and consistent optimization:

$$\mathbf{L} \approx \mathbf{E}\mathbf{W}^{-1}\mathbf{E}^T, \qquad (10)$$

where $\mathbf{E} = [e_{ij}] \in \mathbb{R}^{C \times C'}$ and $\mathbf{W} = [w_{ij}] \in \mathbb{R}^{C' \times C'}$ are the low-rank Laplacian kernel matrices, with elements $e_{ij} = K(\boldsymbol{y}_i, \boldsymbol{z}_j)$ and $w_{ij} = K(\boldsymbol{z}_i, \boldsymbol{z}_j)$.

By incorporating clustering analysis, we achieve two key benefits: (1) rigorously proving the error bound of the Nyström approximation for our utilized Laplacian kernel, and (2) as demonstrated through experiments, providing encouraging results for non-rigid point set registration without compromising the optimization trajectory.

Proposition 1. *The low-rank approximation error* $\epsilon = \|\mathbf{L} - \mathbf{E}\mathbf{W}^{-1}\mathbf{E}^T\|_F$ *in terms of the Laplacian kernel is bounded by*

$$\epsilon \le 4\sqrt{2}T^{3/2}\gamma\sqrt{C'q} + 2C'\gamma^2 Tq\|W^{-1}\|_F, \qquad (11)$$

where $\|\cdot\|_F$ *is the matrix Frobenious norm,* $T = \max_i |\mathbf{P}_i|$, $q = \sum_{j=1}^C \|\boldsymbol{y}_j - \boldsymbol{z}_{c'(j)}\|_2^2$ *is the clustering quantization error with* $c'(j) = \operatorname{argmin}_{i=1,\cdots,C'} \|\boldsymbol{y}_j - \boldsymbol{z}_i\|_2$, *and* γ *is the Laplacian kernel bandwidth defined in Eq. (8).*

Proof. Please see the *Supplementary Material.* □

5. Experimental Results

We perform extensive experiments to demonstrate the performance of the proposed method and compare it with state-of-the-art approaches from both 2D and 3D categories.

Implementation details. Given a pair of point sets, for better numerical stability, we first perform shape normalization to make them follow the standard normal distribution. However, the registration evaluation is still based on the original inputs through denormalization. The Laplacian kernel bandwidth γ is set to 2 by default, and the number of clustering centroids in Elkan k-means equals to $0.3C$ for better trade-off between registration accuracy and efficiency. During optimization, we fix the two weight coefficients $\lambda = 0.5$ and $\zeta = 0.1$, which deliver impressive performance across various scenes. Our algorithm is implemented in MATLAB, on a computer running AMD Core Ryzen 5 3600XT (3.8GHz). We leverage publicly available implementations of baseline approaches for assessment, with their parameters either fine-tuned by ourselves or fixed by the original authors to achieve their best results.

Evaluation criteria. As in [20], we adopt the *Root Mean Squared Error* (RMSE) to quantitatively assess the registration accuracy. For point sets with known ground-truth correspondences, we compute the squared distance between corresponding points directly. However, for point sets without annotated correspondences, such as distinct types of geometries, we identify the corresponding point pairs through the nearest neighbor search. Accordingly, the RMSE is defined as:

$$\mathrm{RMSE}(\mathcal{T}(\mathbf{Y}), \mathbf{X}) = \sqrt{\mathrm{Tr}\{(\mathcal{T}(\mathbf{Y}) - \mathbf{X})^T (\mathcal{T}(\mathbf{Y}) - \mathbf{X})\}/M},$$

where $\mathcal{T}(\mathbf{Y})$ and $\mathbf{X}$ are the deformed and the target point sets, respectively.

5.1. 2D Non-Rigid Point Set Registration

For 2D non-rigid point set registration, we utilize the benchmark IMM hand dataset [40] for evaluation. This dataset encompasses 40 real images, showing the left hands of four distinct subjects, and each contains 10 images. As illustrated in Fig. 3, the hand shape is described through 56 key points extracted from the contour of the hand. We employ the first pose from each group of hands as our target point set, while the remaining poses of the same subject serve as the source point sets. The quantitative comparison results with state-of-the-art 2D registration approaches including MR-RPM [30], BCPD [20], GMM [23], and ZAC [45] are reported in Tab. 1. We report the average RMSE for each subject along with the average registration timing of each method. As observed, our method consistently outperforms the comparative approaches with higher registration accuracy and efficiency across all subjects. Although without the need for constructing the initial point correspondences, like shape context [2] used in MR-RPM, our method still delivers RMSE that is orders of magnitude lower than that of most competitors, highlighting its compelling advantages. The qualitative comparison results regarding the inputs in the third row of Fig. 3 are presented in Fig. 4.
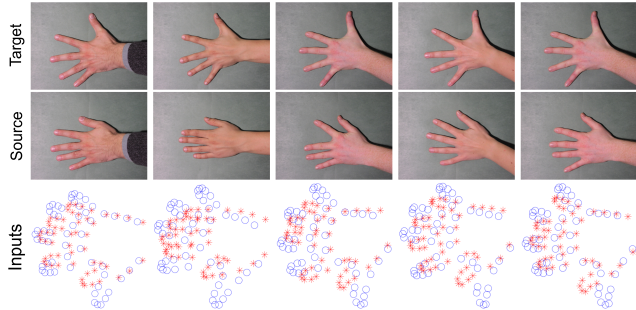
Figure 3. Test samples from the 2D hand pose dataset with different subjects. The red ∗ and the blue ◦ indicate the source and the target point sets, respectively.

Robustness. We further investigate the robustness of the designed method against external disturbances including noise and occlusion. We add a set of Gaussian noise with zero mean and varying standard deviations $\sigma \in [0.01, 0.06]$ to all of the source point sets defined in the above section. Additionally, we randomly erase several points, around $3\% \sim 20\%$ of the source, to construct a range of occlusion geometries. Fig. 5 summarizes the average RMSE values across all subjects. It can be observed that our method still achieves the highest or comparable registration accuracy on all settings, highlighting its stability and robustness. Qualitative comparison results are presented in the *Supplementary Material*.

Table 1. Quantitative comparisons on the 2D hand pose dataset. **Bold** values stand for the top performer.

Method	Subject 1	Subject 2	Subject 3	Subject 4	Time (s)
MR-RPM [30]	0.0940	0.0834	0.1028	0.1388	0.2382
BCPD [20]	0.1027	0.1055	0.1080	0.1579	0.6890
GMM [23]	0.0571	0.0547	0.0734	0.0917	0.1140
ZAC [45]	0.4886	0.4566	0.4879	0.4935	0.4254
Ours	**0.0383**	**0.0481**	**0.0537**	**0.0879**	**0.1074**

5.2. 3D Non-Rigid Point Cloud Registration

Since our method is dimension-independent, we further substantiate its efficacy on 3D point clouds and compare it with eight state-of-the-art 3D registration or deformation approaches, including BCPD [20], GBCPD [20], Fast_RNRR [50], AMM_NRR [48], Sinkhorn [16], as well as network-based ones Nerfies [35], NDP [26], and NSFP [24]. For efficiency, we downsample the point clouds from datasets FAUST [5] and TOSCA [6] using voxel grid filtering, with a point size of $3,000 \sim 4,000$. More experiments are presented in the *Supplementary Material*.

Registration for real human scans. Tab. 2 and Fig. 6 report the quantitative and qualitative comparison results on the FAUST human dataset, respectively. The evaluation is conducted using six sets of subjects in six different and challenging poses for each subject. We first perform *intra-class*
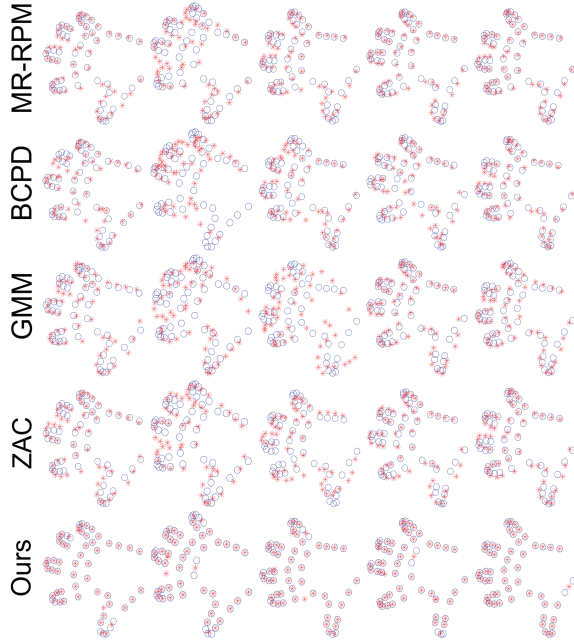

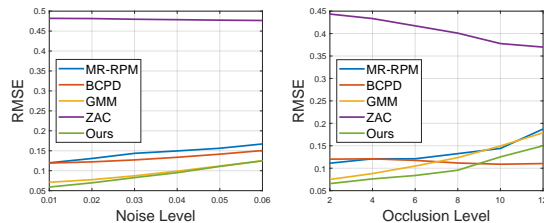Figure 4. Qualitative comparisons on the 2D hand pose dataset.

Figure 5. Robustness comparisons against external disturbances.

registration, *i.e.*, deforming the first human geometry to match the other poses for the same subject. Then, to validate the capability of the designed approach against large deformations, we further conduct an *inter-class* registration test by aligning the first human pose of the i-th subject to all the poses of the $(i+2)$-th subject ($i = 1, 2, 3, 4$). The statistical results summarized in Tab. 2 demonstrate that our method achieves the highest registration accuracy across all subjects and outperforms competitors by a significant margin, even several orders of magnitude higher. Notably, while achieving remarkable accuracy, our method also maintains efficiency comparable to most competitors, making it a highly practical and effective solution. The qualitative comparison results in Fig. 6 indicate that our method not only ensures higher-quality deformations but also recovers the geometric details as well as the topology of the target subject more accurately and faithfully.

Registration for larger deformations. We further verify whether the proposed method improves the registration performance for point cloud pairs with much larger deformations. We evaluate four classes of animals from the TOSCA dataset [6] and report the average RMSE for each class of them. As illustrated in Fig. 7, the source and tar-

Table 2. Quantitative comparisons on the real-world FAUST human scan dataset.

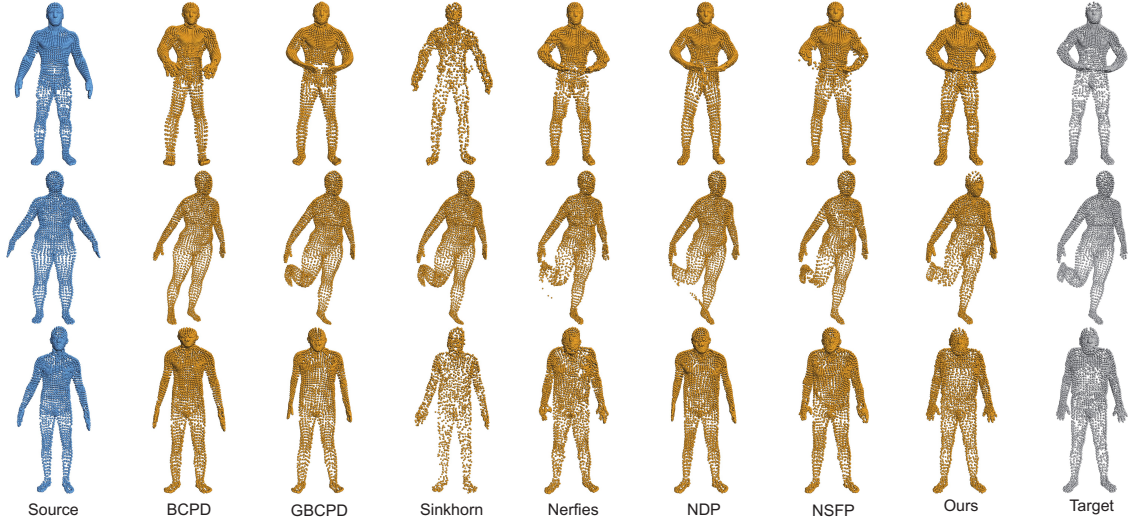| Method \ Settings	Intra-1	Intra-2	Intra-3	Intra-4	Intra-5	Intra-6	Inter-1	Inter-2	Inter-3	Inter-4	Average	Time (s)
BCPD [20]	0.0913	0.1011	0.0872	0.0577	0.1004	0.0746	0.1196	0.0705	0.0935	0.0923	0.0888	3.0359
GBCPD [21]	0.0285	0.0212	0.0211	0.0260	0.0244	0.0339	0.0359	0.0340	0.0212	0.0190	0.0265	1.9346
Fast_RNRR [47]	0.0430	0.0487	0.0397	0.0504	0.0429	0.0391	0.1358	0.0743	0.0477	0.0358	0.0557	**0.6324**
AMM_NRR [48]	0.0544	0.0486	0.0400	0.0539	0.0405	0.0393	0.0838	0.0686	0.0422	0.0399	0.0511	2.0438
Sinkhorn [16]	0.0654	0.0638	0.1372	0.1096	0.0749	0.0821	0.2467	0.0781	0.1400	0.1720	0.1170	2.0377
Nerfies [35]	0.0120	0.0107	0.0138	0.0129	0.0135	0.0118	0.0121	0.0144	0.0140	0.0140	0.0129	9.4287
NDP [26]	0.0183	0.0199	0.0192	0.0152	0.0170	0.0149	0.0181	0.0198	0.0164	0.0155	0.0174	1.7590
NSFP [24]	0.0126	0.0134	0.0132	0.0118	0.0137	0.0142	0.0167	0.0162	0.0148	0.0166	0.0143	2.4607
Ours	**0.0086**	**0.0089**	**0.0103**	**0.0096**	**0.0089**	**0.0081**	**0.0097**	**0.0099**	**0.0094**	**0.0081**	**0.0092**	2.3746

Figure 6. Qualitative comparisons on the FAUST dataset. The top and the bottom rows represent the test registrations of intra and inter-subject, respectively.

Table 3. Quantitative comparisons on the TOSCA dataset.

Method	Cat	Centaur	Dog	Gorilla	Average
BCPD [20]	3.9884	8.1017	7.2800	5.6253	5.9935
GBCPD [21]	1.5631	2.9480	1.5300	3.5751	2.6523
Nerfies [35]	3.2704	2.8826	1.3612	2.2809	2.3211
NDP [26]	4.3639	3.4373	3.1285	2.8312	3.2560
NSFP [24]	1.8774	2.6425	1.6734	2.2044	2.0710
Ours	**1.3496**	**1.8125**	**1.2088**	**1.6807**	**1.5247**

Figure 7. Qualitative comparisons on the TOSCA dataset, where much larger shape deformations exist.

get point sets exhibit significant pose differences, making the registration quite challenging. Tab. 3 summarizes the quantitative comparison results. We exclude Fast_RNRR, AMM_NRR, and Sinkhorn from our analysis because they exhibit significant deviations from the target poses, rendering the error metrics unreliable. Our method consistently outperforms all the baselines by a large margin. Fig. 7 demonstrates that our method delivers highly stable and accurate registration results for point clouds with large deformations, even without the point-wise correspondences.

5.3. Ablation Study

Effect of the improved Nyström method. Fig. 8 reports the registration error, running time, and the matrix approximation error (defined in Proposition 1), between the im-
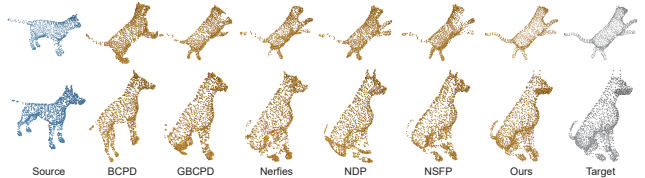
proved or clustered Nyström approximation method (Ours) and the random one on two randomly extracted FAUST models. We vary the approximation ratio $R \in [0.02, 0.4]$ with $\Delta R = 0.02$. It can be seen that our method obtains significant registration and timing performance boost and decrease in matrix approximation error by a large margin, especially on lower ratios.

Laplacian VS. Gaussian. Fig. 9 summarizes both the quantitative and qualitative comparison results between the kernel functions of Gaussian and Laplacian. We validate the merits and robustness of the Laplacian Kernel by aligning the source Bunny model [39] contaminated by a set of noise ($\sigma \in [0, 0.06]$) to a randomly deformed Bunny. The kernel bandwidth γ is varied in $[1, 3]$. We report the RMSE and
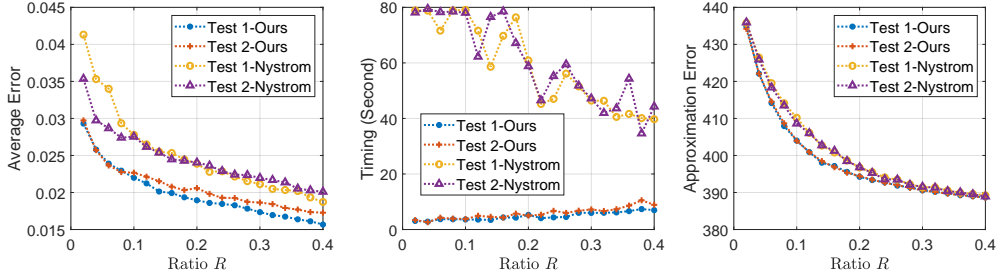
Figure 8. Comparisons between the Nyström low-rank approximation and our clustering-improved one for non-rigid registration.
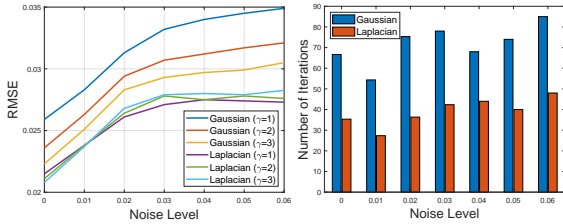

Figure 9. Comparisons between the Gaussian and the Laplacian kernel functions. Left: Results with various kernel bandwidth $\gamma \in [1, 3]$. Right: The number of iterations when algorithms converge.
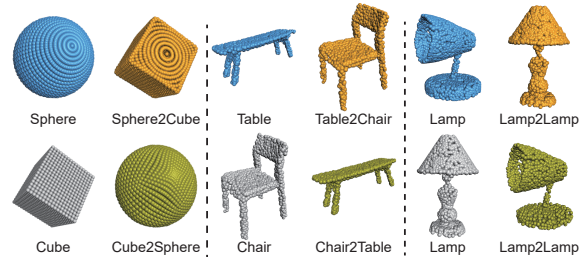
the average number of iterations when algorithms converge in Fig. 9. We find that the Laplacian kernel consistently outperforms the Gaussian kernel across all settings, suggesting the merits of the sparsity-induced ℓ_1 norm. Moreover, the Laplacian kernel delivers faster convergence and more accurate registration results (see the *Supplementary Material*).

5.4. Applications

Shape transfer. As depicted in Fig. 10, we apply the proposed method to transfer shapes belonging to *different categories* that require substantial deformations. We first transfer two geometries with the identical topology (sphere and cube) and then proceed to transfer CAD models from ShapeNet [8], which presents a more challenging task. Results indicate the effectiveness of our method in achieving accurate shape deformation while faithfully preserving the geometric details of the source shapes. Notably, our method consistently produces high-quality deformation results even when the shapes possess significantly distinct topology. More results on shape transfer are presented in the *Supplementary Material*.

Medical registration. Deforming a standard medical template to match those captured from individual patients is a crucial step in the field of medical data analysis. In Fig. 11, we demonstrate the efficacy of our method by aligning a 3D inhale lung volume to two exhale lungs [7] and two brain vessels [44], extracted from real-world CT and MRA images. Despite the presence of complex structures, large deformations, and mutual interference, our method consistently achieves impressive results in accurately deforming the template models to align the target shapes.
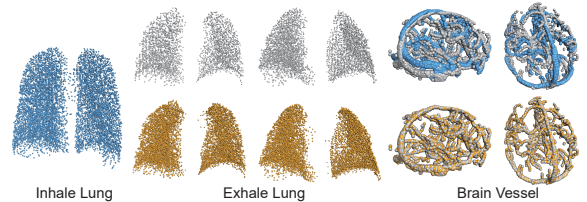

Figure 10. Application of the proposed method to shape transfer.

Figure 11. Application of our method to medical data registration.

6. Conclusions

We proposed an algorithm for solving non-rigid point set registration without prescribed correspondences. The key contribution of our method lies in reformulating non-rigid registration as an unsupervised clustering process that enables holistic optimization, dimension-independent, closed-form solutions, and handling large deformations simultaneously. Moreover, we introduce the ℓ_1-induced Laplacian kernel to achieve a more robust solution than the Gaussian kernel and provide a rigorous approximation bound for the Nyström method. Our method achieves higher-quality results than traditional methods and recent network models, particularly on geometries that exhibit significant deformations. We also showcase its applicability in challenging tasks such as shape transfer and medical registration.

References

[1] Brian Amberg, Sami Romdhani, and Thomas Vetter. Optimal step nonrigid icp algorithms for surface registration. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1–8, 2007. 2

[2] Serge Belongie, Jitendra Malik, and Jan Puzicha. Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(4):509–522, 2002. 2, 5

[3] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517, 1975. 5

[4] PJ Besl and HD McKay. A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(2), 1992. 2

[5] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, Piscataway, NJ, USA, 2014. 6

[6] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical geometry of non-rigid shapes.* 2008. 6

[7] Richard Castillo, Edward Castillo, David Fuentes, Moiz Ahmad, Abbie M Wood, Michelle S Ludwig, and Thomas Guerrero. A reference dataset for deformable image registration spatial accuracy evaluation using the copdgene study archive. *Phys. Med. Biol.*, 58(9):2861, 2013. 8

[8] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 8

[9] Qiang Chen, Feiping Nie, Weizhong Yu, and Xuelong Li. $l_{2,p}$-norm and mahalanobis distance based robust fuzzy c-means. *IEEE Trans Fuzzy Syst*, 2023. 3

[10] Zhe Chen and Simon Haykin. On different facets of regularization theory. *Neural Comput.*, 14(12):2791–2846, 2002. 4

[11] KI Christopher, Matthias Seeger, et al. Using the nyström method to speed up kernel machines. *Proc. Int. Conf. Neural Inf. Process. Syst.*, 13:682–688, 2001. 5

[12] Haili Chui and Anand Rangarajan. A new point matching algorithm for non-rigid registration. *Comput. Vis. Image. Und.*, 89(2-3):114–141, 2003. 2

[13] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Proc. Int. Conf. Neural Inf. Process. Syst.*, 26, 2013. 4

[14] Bailin Deng, Yuxin Yao, Roberto M Dyke, and Juyong Zhang. A survey of non-rigid 3d registration. In *Comput. Graph. Forum*, pages 559–589, 2022. 2

[15] Charles Elkan. Using the triangle inequality to accelerate k-means. In *Proc. Int. Conf. Mach. Learn.*, pages 147–153, 2003. 3

[16] Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-ichi Amari, Alain Trouvé, and Gabriel Peyré. Interpolating between optimal transport and mmd using sinkhorn divergences. In *Int. Conf. Artif. Intell. Stat.*, pages 2681–2690. PMLR, 2019. 6, 7

[17] Federico Girosi, Michael Jones, and Tomaso Poggio. Regularization theory and neural networks architectures. *Neural Comput.*, 7(2):219–269, 1995. 3, 4

[18] Leslie Greengard and John Strain. The fast gauss transform. *SIAM J. Sci. Stat. Comput.*, 12(1):79–94, 1991. 2, 5

[19] Donald E Gustafson and William C Kessel. Fuzzy clustering with a fuzzy covariance matrix. In *1978 IEEE conf. Decis. Control Including the 17th Sym. Adaptive Processes*, pages 761–766, 1979. 3

[20] Osamu Hirose. A bayesian formulation of coherent point drift. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020. 2, 3, 4, 5, 6, 7

[21] Osamu Hirose. Geodesic-based bayesian coherent point drift. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(5):5816–5832, 2022. 2, 7

[22] Jiahui Huang, Tolga Birdal, Zan Gojcic, Leonidas J Guibas, and Shi-Min Hu. Multiway non-rigid point cloud registration via learned functional map synchronization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(2):2038–2053, 2022. 2

[23] Bing Jian and Baba C Vemuri. Robust point set registration using gaussian mixture models. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(8):1633–1645, 2010. 2, 5, 6

[24] Xueqian Li, Jhony Kaesemodel Pontes, and Simon Lucey. Neural scene flow prior. *Proc. Int. Conf. Neural Inf. Process. Syst.*, 34:7838–7851, 2021. 2, 6, 7

[25] Yang Li and Tatsuya Harada. Lepard: Learning partial point cloud matching in rigid and deformable scenes. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5554–5564, 2022. 2

[26] Yang Li and Tatsuya Harada. Non-rigid point cloud registration with neural deformation pyramid. *Proc. Int. Conf. Neural Inf. Process. Syst.*, 35:27757–27768, 2022. 2, 6, 7

[27] Qianfang Liao, Da Sun, and Henrik Andreasson. Point set registration for 3d range scans using fuzzy cluster-based metric and efficient global optimization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(9):3229–3246, 2020. 3

[28] Jiayi Ma, Weichao Qiu, Ji Zhao, Yong Ma, Alan L Yuille, and Zhuowen Tu. Robust $l_{2}e$ estimation of transformation for non-rigid registration. *IEEE Trans. Signal Process.*, 63(5):1115–1129, 2015. 2

[29] Jiayi Ma, Ji Zhao, and Alan L Yuille. Non-rigid point set registration by preserving global and local structures. *IEEE Trans. Image Process.*, 25(1):53–64, 2015. 2

[30] Jiayi Ma, Jia Wu, Ji Zhao, Junjun Jiang, Huabing Zhou, and Quan Z Sheng. Nonrigid point set registration with robust transformation learning under manifold regularization. *IEEE Trans. Neural Netw. Learn. Syst.*, 30(12):3584–3597, 2018. 1, 2, 4, 5, 6

[31] James MacQueen et al. Some methods for classification and analysis of multivariate observations. In *Proc. Berkeley Symp. Math. Stats. Pbty*, pages 281–297, 1967. 3

[32] Sadaaki Miyamoto, Hodetomo Ichihashi, Katsuhiro Honda, and Hidetomo Ichihashi. *Algorithms for fuzzy clustering.* 2008. 3

[33] Andriy Myronenko and Xubo Song. Point set registration: Coherent point drift. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(12):2262–2275, 2010. 2, 3, 4

[34] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 343–352, 2015. 1

[35] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 5865–5874, 2021. 1, 2, 6, 7

[36] Zheng Qin, Hao Yu, Changjian Wang, Yuxing Peng, and Kai Xu. Deep graph-based spatial consistency for robust non-rigid point cloud registration. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5394–5403, 2023. 1, 2

[37] Zsolt Sánta and Zoltan Kato. Correspondence-less non-rigid registration of triangular surface meshes. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2275–2282, 2013. 2

[38] Aliaksandr Siarohin, Willi Menapace, Ivan Skorokhodov, Kyle Olszewski, Jian Ren, Hsin-Ying Lee, Menglei Chai, and Sergey Tulyakov. Unsupervised volumetric animation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4658–4669, 2023. 1

[39] Stanford. The Stanford 3D Scanning Repository. http://graphics.stanford.edu/data/3Dscanrep/. 7

[40] Mikkel B Stegmann and David Delgado Gomez. A brief introduction to statistical shape analysis. *Informatics and mathematical modelling, Technical University of Denmark, DTU*, 15(11), 2002. 5

[41] Gary KL Tam, Zhi-Quan Cheng, Yu-Kun Lai, Frank C Langbein, Yonghuai Liu, David Marshall, Ralph R Martin, Xian-Fang Sun, and Paul L Rosin. Registration of 3d point clouds and meshes: A survey from rigid to nonrigid. *IEEE Trans. Vis. Comput. Graph.*, 19(7):1199–1217, 2012. 2

[42] Jiapeng Tang, Lev Markhasin, Bi Wang, Justus Thies, and Matthias Nießner. Neural shape deformation priors. *Proc. Int. Conf. Neural Inf. Process. Syst.*, 35:17117–17132, 2022. 1

[43] Te Tang, Changhao Wang, and Masayoshi Tomizuka. A framework for manipulating deformable linear objects by coherent point drift. *IEEE Robot. Autom. Lett.*, 3(4):3426–3433, 2018. 2

[44] TubeTK. The TubeTK Testing Data. https://public.kitware.com/Wiki/TubeTK/Data. 8

[45] Fudong Wang, Nan Xue, Jin-Gang Yu, and Gui-Song Xia. Zero-assignment constraint for graph matching with outliers. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3033–3042, 2020. 5, 6

[46] Rui Xu and Donald Wunsch. Survey of clustering algorithms. *IEEE Trans. Neural Netw. Learn. Syst.*, 16(3):645–678, 2005. 3

[47] Yuxin Yao, Bailin Deng, Weiwei Xu, and Juyong Zhang. Quasi-newton solver for robust non-rigid registration. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 1, 2, 7

[48] Yuxin Yao, Bailin Deng, Weiwei Xu, and Juyong Zhang. Fast and robust non-rigid registration using accelerated majorization-minimization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2023. 1, 2, 6, 7

[49] Alan L Yuille and JJ Kosowsky. Statistical physics algorithms that converge. *Neural Comput.*, 6(3):341–356, 1994. 4

[50] Juyong Zhang, Yuxin Yao, and Bailin Deng. Fast and robust iterative closest point. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2021. 6

[51] Kai Zhang, Ivor W Tsang, and James T Kwok. Improved Nyström low-rank approximation and error analysis. In *Proc. Int. Conf. Mach. Learn.*, pages 1232–1239, 2008. 2, 5

[52] Mingyang Zhao, Lei Ma, Xiaohong Jia, Dong-Ming Yan, and Tiejun Huang. Graphreg: dynamical point cloud registration with geometry-aware graph signal processing. *IEEE Trans. Image Process.*, 31:7449–7464, 2022. 1

[53] Mingyang Zhao, Xiaoshui Huang, Jingen Jiang, Luntian Mou, Dong-Ming Yan, and Lei Ma. Accurate registration of cross-modality geometry via consistent clustering. *IEEE Trans. Vis. Comput. Graph.*, 2023. 3

[54] Michael Zollhöfer, Matthias Nießner, Shahram Izadi, Christoph Rehmann, Christopher Zach, Matthew Fisher, Chenglei Wu, Andrew Fitzgibbon, Charles Loop, Christian Theobalt, et al. Real-time non-rigid reconstruction using an rgb-d camera. *ACM Trans. Graph.*, 33(4):1–12, 2014. 1