

Multiway Non-Rigid Point Cloud Registration via Learned Functional Map Synchronization

Jiahui Huang¹, Tolga Birdal², Zan Gojcic³,
Leonidas J. Guibas, *Senior Member, IEEE*, and Shi-Min Hu¹, *Senior Member, IEEE*

Abstract—We present SyNoRiM, a novel way to jointly register multiple non-rigid shapes by synchronizing the maps that relate learned functions defined on the point clouds. Even though the ability to process non-rigid shapes is critical in various applications ranging from computer animation to 3D digitization, the literature still lacks a robust and flexible framework to match and align a collection of real, noisy scans observed under occlusions. Given a set of such point clouds, our method first computes the pairwise correspondences parameterized via functional maps. We simultaneously learn potentially non-orthogonal basis functions to effectively regularize the deformations, while handling the occlusions in an elegant way. To maximally benefit from the multi-way information provided by the inferred pairwise deformation fields, we synchronize the pairwise functional maps into a *cycle-consistent* whole thanks to our novel and principled optimization formulation. We demonstrate via extensive experiments that our method achieves a state-of-the-art performance in registration accuracy, while being flexible and efficient as we handle both non-rigid and multi-body cases in a unified framework and avoid the costly optimization over point-wise permutations by the use of basis function maps.

Index Terms—3D point cloud, non-rigid registration, functional map synchronization

1 INTRODUCTION

THE prevalence of reliable 3D data capture fueled countless applications impacting from movie industry to robotics. In a wide variety of these applications, one needs to capture (in 3D) non-rigidly moving objects from multiple angles or over time [1], [2]. This leads to a dynamic, multi-scan alignment problem, further obstructed by the presence of occlusions, ambiguities, and noise.

Solving this challenging task fostered the development of a plethora of temporal, mesh-based dynamic non-rigid registration algorithms (e.g. [3]). However, these methods suffer from two main limitations: (1) the data provided by 3D sensors hardly come in mesh format, let alone the difficulty associated with preserving the mesh topology. Signed distance field based reconstruction methods like [4] can overcome some of these problems, but (2) they still assume a *streaming* depth map and fail to maintain correspondences, which are critical for defining non-rigid deformations. In addition to those nuisances, a drift-

free alignment almost surely demands a global optimization step, which exploits all possible loop closure constraints in the form of a graph optimization [5] or bundle adjustment [6].

In this paper, we set off in pursuit of alleviating both of these issues. In particular, we (1) leverage unstructured point cloud representations for maximal generality and flexibility; (2) assume a fully connected graph in lieu of the sequential order, incorporating drift reduction in the early stages. Additionally, compared to rigid motion [7], [8] or pure as-rigid-as-possible [9] deformation, our relaxed assumptions allow for more general object dynamics. In particular, we propose *SyNoRiM* framework, synchronization for non-rigid multi-way registration. Input to SyNoRiM, is a collection of point cloud scans containing a potentially non-rigidly deforming object. We then seek to recover a coherent and consistent 3D scene flow, originating from any source shape to a latent target shape that is to be discovered simultaneously. We start by a *motion coherence observation* [10] that pairwise correspondences between natural shapes are smooth and band-limited, i.e. nearby points in the source point cloud should map to nearby points in the target (except for topological changes and occlusions). This allows us to substitute the point-wise matches between pairs of scans by linear maps of *learned smooth functions* defined on the points. Such notion of a *functional map* was first introduced in the geometry processing community [11] for estimating correspondences between isometric meshes. Unfortunately, due to the lack of proper basis functions, an extension to point cloud analysis is not trivial [12]. In SyNoRiM we propose to directly learn the bases from a large-scale training corpus, either in a fully-supervised or unsupervised fashion. Such bases learned by neural networks are found to be robust to occlusions, partiality, and errors in the initialization because an end-to-end training informs the bases about the shape context while the

- Jiahui Huang and Shi-Min Hu are with the BNRist, Department of Computer Science and Technology, Tsinghua University, Beijing 100190, China. E-mail: huang-jh18@mails.tsinghua.edu.cn, shimin@tsinghua.edu.cn.
- Tolga Birdal and Leonidas J. Guibas are with Stanford University, Stanford, CA 94305 USA. E-mail: tbirdal@gmail.com, guibas@cs.stanford.edu.
- Zan Gojcic is with the Nvidia, Toronto, ON M5V 3M4, Canada. E-mail: zan.gojcic@gmail.com.

Manuscript received 10 November 2021; revised 14 March 2022; accepted 26 March 2022. Date of publication 5 April 2022; date of current version 6 January 2023.

This work was supported by the National Key R&D Program of China under Grant 2021ZD0112902, in part by Vannevar Bush faculty fellowship, ARL under Grant W911NF2120104, in part by NSF under Grant IIS-1763268, and in part by gift from the Autodesk Corporation.

(Corresponding author: Shi-Min Hu.)

Recommended for acceptance by V. Koltun.

Digital Object Identifier no. 10.1109/TPAMI.2022.3164653

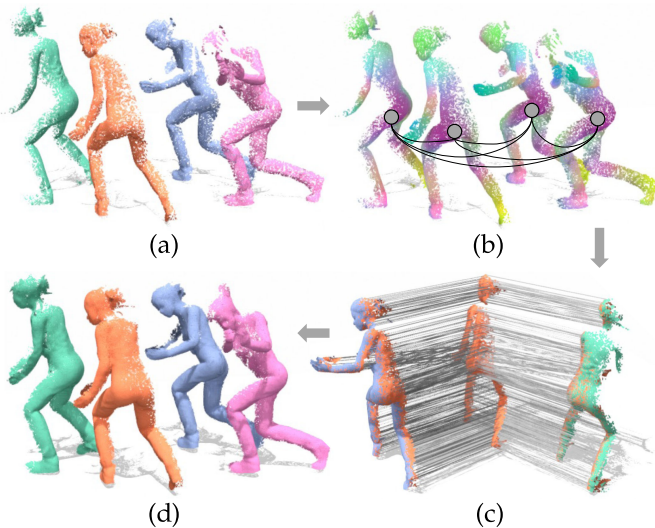


Fig. 1. *SyNoRiM* overview. (a) Input point clouds. (b) Synchronized canonical functions \mathbf{H} (visualized via Principal Component Analysis; see Section 5 for details). (c) Scene flow estimations for all pairs (here only two sets of flow, \rightarrow and \leftarrow , are shown for brevity). (d) Registered (accumulated) point clouds by gathering warped points from other inputs.

inherent smoothness enables meaningful extrapolations of the target coordinate functions. A subsequent *functional map synchronization* algorithm coerces the maps to be globally consistent, by making use of the information that individual observations are essentially samples from a common underlying surface. Based on this, we finally refine the pairwise deformations eliminating the ambiguities. An illustrative figure is shown in Fig. 1.

In summary, our contributions are:

- 1) We propose, to the best of our knowledge, the first end-to-end data-driven framework to learn consistent registration between multiple, possibly partial point cloud observations of non-rigidly moving bodies / objects.
- 2) Our deep network can learn potentially non-orthogonal functional bases on point sets without requiring to define Laplace-Beltrami operators (LBO). Functional maps between such bases can match non-isometric shapes.
- 3) We propose a novel *functional map synchronization* algorithm enforcing cycle-consistency among the pairwise deformation fields, estimated in isolation. This harmonizes the 3D flows into a coherent whole and thus enables multiway registration.

We demonstrate the efficacy of our algorithm through extensive evaluations on both rigid and non-rigid scenarios, showing superior performance on all datasets.

2 RELATED WORKS

Point-Set Registration. Finding reliable alignments between point clouds plays a fundamental role for many downstream tasks. Rigid registration estimates a single transformation matrix through either heuristic searches [13] or learned local/global descriptors [14], [15], [16], [17], [18], while common non-rigid registration techniques aim to best

align the clouds under the various data terms [19], [20], [21] and regularizations [4], [9] with different deformation representations [22], [23]. In this paper we consider the general non-rigid deformation scheme and also demonstrate results on hybrid ones such as multibody [24], [25], [26]. On the other hand, scene flow describes the transition between two point clouds using a three-dimensional vector field and is a low-level task agnostic of the deformation type. Existing techniques [27], [28], [29], [30] handle the task via accurate modeling of the point-wise features as well as its neighborhood context, and reach a good performance for cluttered scenes or driving scenarios, but is not robust under large deformations and ambiguities [31].

Function-Based Correspondence. In the field of geometry processing, the use of functional techniques is a recent trend for building reliable correspondences between 3D shapes (e.g., discretized manifold meshes). First introduced in [11], such methods compute the LBO eigenvectors as basis functions and infer a linear transformation of a subset of bases to indicate low-rank shape mapping. Many extensions have been vastly explored such as hierarchical matching [32], partial-to-full handling [33] or integration into deep learning frameworks [34], [35] with either supervised or unsupervised [36], [37], [38] methods. Contrarily in the domain of point set analysis, such methods are less popular despite a few [39], [40], [41]. Notably, [40] also propose to learn a linearly-invariant embedding as bases, yet their learning scheme is different from ours and they are not robust to partialities or occlusions, which are commonly observed in point cloud data. We further highlight that the idea to project natural signals to lower dimensions has also been partially explored in the vision community [42], [43].

Synchronization on 3D Geometry. Though initially established as a theory for clock systems, synchronization has now been widely used in many vision tasks such as structure from motion [5], [44], [45], semantic segmentation [46], correspondence refinement [47], [48] and multiway rigid/multibody registration [7], [49], [50]. By enforcing cycle consistencies within a system, the relative measurements are globally harmonized thanks to the averaging of local noise. In the field of 3D geometry analysis, analyzing and synchronizing spectral functional maps [51] from a mesh collection are widely applied to full shape matching, either with low-rank factorization [52], limit shapes extraction [53], coarse-to-fine strategy [54] or joint point-spectrum optimization [55]. In contrast, our novel synchronization formulation is tailored for the bases learned from raw point clouds that are free of the geometric impositions like (near-)isometries and directly takes correspondences into account, robustly.

3 OVERVIEW

Problem Formulation. The input to our method is a point cloud graph $\mathcal{G} := (\mathcal{V}, \mathcal{E})$, whose vertices \mathcal{V} represent the input set of K point clouds $\mathcal{V} := \{\mathbf{X}_k \in \mathbb{R}^{N_k \times 3}, k \in [1, K]\}$ and the edges $\mathcal{E} := \{(k, l), \mathbf{X}_k \in \mathcal{V}, \mathbf{X}_l \in \mathcal{V}\}$ represent the graph connectivity. By default we assume \mathcal{G} to be fully-connected. The output of our method contains all the pairwise per-point 3D flow vectors $\mathcal{F} := \{\mathbf{F}_{kl} \in \mathbb{R}^{N_k \times 3}, (k, l) \in \mathcal{E}\}$. The flow vectors naturally induce the non-rigid warp field from \mathbf{X}_k to \mathbf{X}_l as $\mathcal{W}_{kl}(\mathbf{X}_k) := \mathbf{X}_k + \mathbf{F}_{kl}$, optimally aligning the given




Fig. 2. *Method.* During training, our method is supervised in a pairwise fashion. We first train ① φ_{desc} (i.e., \mathfrak{g}_{pd}) to establish putative correspondences between each point cloud pair. We then estimate a set of basis functions $\{\Phi_k\}$ for each point cloud using φ_{basis} to obtain ② the initial functional map \mathbf{C}_{kl}^0 before refining the 3D flow estimates with φ_{refine} . During test time with multiple inputs, we estimate the map set $\{\mathbf{C}_{kl}^0\}$ for all pairs. ③ The estimated maps are subsequently synchronized to optimize cycle consistency among the inputs. Finally, ④ 3D flows are estimated from the optimized functional maps $\{\mathbf{C}_{kl}^0\}$ as our final output. The registered point cloud is a fusion of all initial point clouds warped by the estimated flows.

point cloud pairs by deforming the source onto the target. We additionally encourage the cyclic consistency of the estimates \mathcal{F} , defined loosely as:

$$\mathcal{W}_{k_1 k_2} \circ \dots \circ \mathcal{W}_{k_{p-1} k_p} \circ \mathcal{W}_{k_p k_1} = \mathbf{I}, \forall (k_1, \dots, k_p) \in \mathcal{C}(\mathcal{G}), \quad (1)$$

where $\mathcal{C}(\mathcal{G})$ are the set of cycles in the \mathcal{G} and \mathbf{I} is the identity warping. The domain of the above composed warp map is the union region $(\mathbf{X}_{k_1}, \dots, \mathbf{X}_{k_p})$.

Method Summary. Our method begins by establishing sparse point-level putative correspondences using a correspondence generator $\mathfrak{g}(\cdot)$ (Section 4.2). In the meantime, a basis network φ_{basis} is applied independently to each input point cloud to generate a set of basis functions defined on the points (Section 4.3). Given the sparse correspondences and bases we can compute the initial pairwise functional map matrices. The technique to recover flows from the matrices is realized through another network φ_{refine} (Section 4.4). We then feed the initial map estimates into a synchronization module that jointly optimizes for all pairwise functional mappings, considering the cycle consistency (Section 5), before utilizing the φ_{refine} module once more to get the optimized deformations. The full pipeline is illustrated in Fig. 2.

4 PAIRWISE FUNCTIONAL REGISTRATION

4.1 Preliminaries on Functional Maps

Given two (abstract) shapes \mathcal{S}_k and \mathcal{S}_l in the form of smooth continuous manifolds, a functional map $T_{kl}: L_2(\mathcal{S}_l) \mapsto L_2(\mathcal{S}_k)$ maps from the space of square-integrable real-valued functions (L_2 -space) defined on \mathcal{S}_l to \mathcal{S}_k . Such an operator is proved to be linear [11], i.e., $T_{kl}(\alpha_1 f_1 + \alpha_2 f_2) = \alpha_1 T_{kl}(f_1) + \alpha_2 T_{kl}(f_2)$, where f_1 and f_2 are functions defined on \mathcal{S}_l and α_1 and α_2 are the coefficients.

Every function f_k on \mathcal{S}_k (or f_l on \mathcal{S}_l) can be represented as a linear combination of *basis functions* $\{\phi_{k,m}\}$ (or $\{\phi_{l,m}\}$): $f_k = \sum_m h_{k,m} \phi_{k,m}$, $f_l = \sum_m h_{l,m} \phi_{l,m}$, with $h_{k,m}$ and $h_{l,m}$ being the coefficients. Moreover, most natural functions defined on shapes can be approximated by linearly combining a *finite* set of M bases, if correctly chosen. A functional map matrix $\mathbf{C}_{kl} \in \mathbb{R}^{M \times M}$ can be then defined as a replacement for T_{kl} , and satisfies $T_{kl}(\phi_{l,m_l}) = \sum_{m_k=1}^M (\mathbf{C}_{kl})_{m_k, m_l} \phi_{k, m_k}$, where $(\mathbf{C}_{kl})_{m_k, m_l}$ denotes the element at the m_k -th row and the m_l -th column.

We can then re-write the relation $f_k = T_{kl}(f_l)$ in linear

algebra as $\mathbf{h}_k = \mathbf{C}_{kl} \mathbf{h}_l$, where $\mathbf{h}_k := [h_{k,1}, \dots, h_{k,M}]^\top$, $\mathbf{h}_l := [h_{l,1}, \dots, h_{l,M}]^\top$. This re-formulation admits various computation tools available for optimization. We refer readers to a full tutorial provided in [56] for more in-depth discussions.

Discretization on Point Clouds. As points are samples from the surfaces, we can define real-valued functions on each point cloud (e.g., \mathbf{X}_k) as a column vector in \mathbb{R}^{N_k} . Horizontally stacking the set of M basis functions $\{\phi_{k,m}\}_{m=1}^M$ gives a full-rank compact basis matrix $\Phi_k \in \mathbb{R}^{N_k \times M}$. The pairwise linear functional map matrix \mathbf{C}_{kl} , when left-multiplied with Φ_k , linearly re-combines the bases from \mathbf{X}_k and yields the transferred set of bases Φ_l from \mathbf{X}_l to \mathbf{X}_k , i.e., $\Phi_k \mathbf{C}_{kl} \approx \Pi_{kl} \Phi_l$, where Π_{kl} is the point-wise permutation matrix between the two point clouds. For other arbitrary functions $\mathbf{f} \in \mathbb{R}^{N_k}$, their coordinates under the bases are given by $\Phi_k^+ \mathbf{f}$, where $^+$ is the Moore-Penrose pseudo-inverse operator.

4.2 Generating Putative Correspondences

We estimate deep features $\mathbf{D}_k \in \mathbb{R}^{N_k \times F}$ for each point cloud in \mathcal{V} using a sparse-convolution-based [57] feature descriptor network $\varphi_{\text{desc}}: \mathbf{X}_k \rightarrow \mathbf{D}_k$. Given a pair of descriptors \mathbf{D}_k and \mathbf{D}_l , we construct the soft permutation matrix $\hat{\Pi}_{kl}^d := \text{softmax}(\hat{\mathbf{\Pi}}_{kl}^d) \in \mathbb{R}^{N_k \times N_l}$ and compute the scene flow as follows:

$$\mathbf{F}_{kl}^d := \hat{\Pi}_{kl}^d \mathbf{X}_l - \mathbf{X}_k, \quad (\hat{\Pi}_{kl}^d)_{ij} := -\frac{1}{t^d} \|\mathbf{D}_k\|_i - \|\mathbf{D}_l\|_j, \quad (2)$$

where t^d is a trainable parameter with the initial value of 1.0 and a minimum value of 0.02. $\text{softmax}(\cdot)$ performs softmax normalizations over all the rows of $\hat{\Pi}_{kl}^d$ so that the output $\hat{\Pi}_{kl}^d$ becomes row-stochastic (i.e., each row sums to 1). The flow loss \mathcal{L}_f introduced in Section 4.5 will be used to supervise φ_{desc} until convergence. Remarkably, such a simple strategy, similar to [29], already produces the flow \mathbf{F}_{kl}^d . However, such an estimation is corrupted with noise, occlusions, and inconsistencies hence performing a lot worse than our full pipeline. This is verified through the ‘Ours (φ_{desc} only)’ baseline in the experiments, where we simply map each point to its nearest neighbor in the space of \mathbf{D}_k .

After φ_{desc} is trained, it is fixed and we define the correspondence generator as $\mathfrak{g}_{\text{pd}}: (\mathbf{X}_k, \Phi_k), (\mathbf{X}_l, \Phi_l) \rightarrow (\Phi_k^{(kl)}, \Phi_l^{(kl)})$.

which selects *corresponding* rows (i.e., matched points) from the input bases Φ_k and Φ_l (detailed below), and produces $\Phi_k^{(kl)} \in \mathbb{R}^{I_{kl} \times M}$ and $\Phi_l^{(kl)} \in \mathbb{R}^{I_{kl} \times M}$, where I_{kl} is the number of matches. The row indices are determined by a nearest neighbor search performed on the descriptors \mathbf{D}_k and \mathbf{D}_l with cross-check, and the rows are matched if the L2 distance of the descriptors is smaller than a *conservative* threshold of 0.3.

4.3 Computing Basis Functions

The basis functions $\{\Phi_k\}$ play a critical role on the power of learned representations. For triangulated shapes where connectivity information is provided, bases can be easily formed via the standard method of decomposing the graph Laplacian [51]. Similarly, for point clouds one can also build a mesh structure such as a k-NN graph or intrinsic Delaunay triangulation [58] before computing the bases. However, the latter can be problematic because the constructed graph would have no notion of semantics or geodesics, leading to redundant or erroneous function approximation. Instead, we propose to learn the basis functions directly from point clouds using a sparse-convolution-based [57] neural network $\varphi_{\text{basis}}: \mathbf{X}_k \rightarrow \Phi_k$. Basis functions learned in such manner are both (1) accurate, i.e. focus only on the flow properties we are interested in, and (2) compact with no redundancies, as will be demonstrated in Section 6, which allows achieving higher accuracy even with a small number of bases.

In order to train φ_{basis} , we compute the optimal functional maps as

$$\mathbf{C}_{kl}^0 := \underset{\mathbf{C}}{\operatorname{argmin}} E_{kl}(\mathbf{C}), \quad (3)$$

where

$$\begin{aligned} E_{kl}(\mathbf{C}) &:= \sum_{i=1}^{I_{kl}} \rho\left(\|\Phi_{l,i}^{(kl)} - \Phi_{k,i}^{(kl)}\mathbf{C}\|\right), (\Phi_k^{(kl)}, \Phi_l^{(kl)}) \\ &= \mathfrak{g}((\mathbf{X}_k, \Phi_k), (\mathbf{X}_l, \Phi_l)). \end{aligned} \quad (4)$$

Here, $\mathfrak{g}(\cdot)$ is a putative correspondence generator detailed earlier. The rows of $\Phi_k^{(kl)}$ and $\Phi_l^{(kl)}$ are indexed via $(\cdot)_{i:}$. During training we use the strategy similar to [59] by randomly instantiating \mathfrak{g} either with the ground-truth \mathfrak{g}_{gt} (if provided) or the predicted one \mathfrak{g}_{pd} with a probability of 50%. We empirically observe the positive impact of this on final performance. During test time, when ground-truth information is not available, we simply let $\mathfrak{g} = \mathfrak{g}_{\text{pd}}$. $\rho(\cdot)$ is the Huber robust function [60] used for outlier rejection that individually treats each summand i . The scale of $\rho(\cdot)$ is chosen to be 0.05 and empirically we do not observe any significant difference using more advanced adaptive scales such as the median absolute deviations [61]. After obtaining \mathbf{C}_{kl}^0 we feed it into the method in Section 4.4 to compute the flow vectors, and supervise the networks using the loss function defined later in Section 4.5.

4.4 Recovering Flows From Functional Maps

Given a computed functional map matrix \mathbf{C}_{kl} , which could either be \mathbf{C}_{kl}^0 or the optimized one \mathbf{C}_{kl}^* from Section 5, one can again compute the soft permutation matrix $\hat{\mathbf{\Pi}}_{kl} := \operatorname{softmax}(\hat{\mathbf{\Pi}}_{kl})$ between the two clouds and extract the flow as:

$$\begin{aligned} \mathbf{F}_{kl}^n &:= \hat{\mathbf{\Pi}}_{kl} \mathbf{X}_l - \mathbf{X}_k, \\ (\hat{\mathbf{\Pi}}_{kl})_{ij} &:= -\frac{1}{t} \|\left(\Phi_k \mathbf{C}_{kl}\right)_i - (\Phi_l)_j\|. \end{aligned} \quad (5)$$

However, as pointed out before, $\hat{\mathbf{\Pi}}_{kl}$ is row-stochastic and the warped position $\hat{\mathbf{\Pi}}_{kl} \mathbf{X}_l$ hence never extends beyond the *convex hull* of \mathbf{X}_l , which is detrimental for computing flows for occluded points. An alternative way to compute the flow is to treat the target positions \mathbf{X}_l *themselves* as three functions and map their coordinates via \mathbf{C}_{kl} , evaluated as:

$$\mathbf{F}_{kl}^f := \Phi_k \mathbf{C}_{kl} \Phi_l^+ \mathbf{X}_l - \mathbf{X}_k. \quad (6)$$

Note that the ground-truth mapping already implies $\Phi_k \mathbf{C}_{kl} \approx \hat{\mathbf{\Pi}}_{kl} \Phi_l$ as noted in Section 4.1. However, the two methods for computing flows are fundamentally different, conceptually: \mathbf{F}_{kl}^f elegantly handles the issue of occlusions because the mapping allows us to map unknown positions to valid ranges given the continuity of the bases, leading to a learned extrapolation scheme via band-limited regularization. Nevertheless, the truncated basis may over-regularize the positions in geometrically-complicated parts and limit the representation power of the flows. In contrast, \mathbf{F}_{kl}^n can perform well on those regions thanks to notion of point-wise correspondences.

Hence, we combine the best of the two worlds by an additional refinement network φ_{refine} which digests the concatenation of the two, along with \mathbf{X}_k , and outputs the residual:

$$\mathbf{F}_{kl} := \mathbf{F}_{kl}^n + \varphi_{\text{refine}}([\mathbf{X}_k, \mathbf{F}_{kl}^f, \mathbf{F}_{kl}^n]). \quad (7)$$

The refinement network φ_{refine} is instantiated with a sparse-convolution-based network, which works in a coarse-to-fine fashion by first aligning the general cloud structure with \mathbf{F}_{kl}^f and then fixing small-scale detailed errors by considering the information from \mathbf{F}_{kl}^n . The smoothness property of such networks also helps prune predicted outliers as already demonstrated by previous works (e.g. [29]). A comparison between the different flows is shown in Fig. 6.

4.5 Loss and Training

We train our network in two steps, where: (1) We train the correspondence generator \mathfrak{g} described in Section 4.2 using a *flow loss* \mathcal{L}_f only, and (2) we jointly train φ_{basis} and φ_{refine} (Sections 4.3 and 4.4) end-to-end adding the *consistency loss* \mathcal{L}_c , i.e.:

$$\mathcal{L}_f + \lambda_c \mathcal{L}_c. \quad (8)$$

Notably, the argmin operator in (3) is differentiated by computing the solution with Iteratively Reweighted Least Squares (IRLS) and unrolling the re-weighting iterations. We provide the details of this algorithm in Appendix C, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2022.3164653>.

Flow Loss \mathcal{L}_f . If ground-truth annotations $\mathbf{F}_{kl}^{\text{gt}}$ are provided, we can directly use:

$$\mathcal{L}_{f,\text{sup}} := \|\mathbf{F}_{kl} - \mathbf{F}_{kl}^{\text{gt}}\|_{\text{F}}^2, \quad (9)$$

where $\|\cdot\|_{\text{F}}$ denotes Frobenius norm. Otherwise, we use the self-supervised loss inspired from [28]:




Fig. 3. *Learned bases visualization.* We visualize for two shapes (left) a chosen subset of its M bases. The basis range is normalized.

$$\mathcal{L}_{f,\text{unsup}} := \sum_{\text{type} \in \{\text{chamber}, \text{smooth}, \text{lap}\}} \lambda_{\text{type}} \mathcal{L}_{\text{type}}, \quad (10)$$

where for conciseness, the respective loss terms $\mathcal{L}_{\text{type}}$, the weights being λ_{type} , are defined in Appendix A. We evaluate both settings using either $\mathcal{L}_{f,\text{sup}}$ or $\mathcal{L}_{f,\text{unsup}}$ in our experiments (see Section 6.1 for details). The more general \mathcal{L}_f refers to both for convenience.

Consistency Loss \mathcal{L}_c We impose the pairwise consistency loss [37] as:

$$\mathcal{L}_c := \|\mathbf{C}_{kl}^0 \mathbf{C}_{lk}^0 - \mathbf{I}_M\|_F^2, \quad (11)$$

where \mathbf{I}_M is the identity matrix with size $M \times M$. This loss regularizes the compound mapping from functions on \mathbf{X}_k to functions on \mathbf{X}_l and its backward forms an identity, ensuring a local consistency of the learned bases. Additionally, we tried other regularization terms over the bases, such as Laplacian commutativity [11]¹ or descriptor preservation [62] but empirically no improvement is observed.

4.6 Discussions

Further Notice of Φ and \mathbf{C} . Examples of the learned basis functions are visualized in Fig. 3. Remarkably, the learned bases Φ and the functional map matrix \mathbf{C} are not constrained to a particular geometric structure, e.g., orthonormality $\Phi^T \Phi = \mathbf{I}$ or orthogonality $\mathbf{C}\mathbf{C}^T = \mathbf{I}$, which hold only under isometric deformations with full geometry for meshes. The coarse-to-fine structure of Φ is not enforced either: the network finds its own way of representing high-frequency bases via learning small differences. Empirically we find that enforcing any of these regularizations harms the performance. On the other hand, unlike most existing works (e.g., [34], [40]) that formulates (4) by projecting *probe functions* $\mathbf{P}_k, \mathbf{P}_l$ via $\mathbf{C}_{kl}^0 := \arg\min_{\mathbf{C}} \|\mathbf{C}\Phi_l^+ \mathbf{P}_l - \Phi_k^+ \mathbf{P}_k\|$, our method is more robust due to the per-match outlier filtering and gets rid of the trade-off between the flexibility of probe functions and the compactness of the bases for flow estimation.

Special Case: Affinity Bases. For large-scale scenes with multiple moving *rigid* objects as often observed in the case of autonomous driving, the resulting flow can be effectively regularized by constraining the basis structure. Specifically, if the segmentation of S rigid bodies (each indexed by s) is known beforehand (e.g., acquired via [8] or semantic




Fig. 4. An illustration of affinity bases.

segmentations), instead of learned ones, we can manually define the basis functions as:

$$\forall s, (\Phi_k)_{:,4s:4s+4} = \text{diag}(\mathbf{G}_k = s)[\mathbf{X}_k \quad \mathbf{1}], \quad (12)$$

where $\text{diag}(\cdot)$ is the diagonal matrix of an indicator vector, $\mathbf{G}_k \in \mathbb{R}^{N_k}$ are the rigid body indices corresponding to \mathbf{X}_k and the final bases $\Phi_k \in \mathbb{R}^{N_k \times 4S}$ are split into S blocks (as visualized in the inset as a 2D example). Note that the resulting map $\mathbf{C}_{kl} \in \mathbb{R}^{4S \times 4S}$ not only shows the affine transformation matrices in its 4×4 sub-blocks, but the sparsity of the sub-blocks also reveals rigid-body-level permutations. As a demonstration of the effectiveness of such bases, we provide evaluations on driving scenes in Section 6.6.

5 FUNCTIONAL SYNCHRONIZATION

For a geometrically coherent outcome, we further ask the loops in \mathcal{G} to be closed. Note that this step happens only during test time after the pairwise networks are trained. In our setting of functional mappings, this amounts to enforcing the cycle-consistency [48]:

$$\mathbf{C}_{k_1 k_2} \dots \mathbf{C}_{k_{p-1} k_p} \mathbf{C}_{k_p k_1} = \mathbf{I}, \quad \forall (k_1, \dots, k_p) \in \mathcal{C}(\mathcal{G}). \quad (13)$$

It can be easily shown by construction that the above formulation is equivalent to defining a set of V *canonical* functions $\mathbf{H}_k = [\mathbf{h}_k^1, \dots, \mathbf{h}_k^V] \in \mathbb{R}^{M \times V}$ on each point cloud \mathbf{X}_k , expressed as the coordinates under Φ_k . By enforcing the canonical functions to evaluate the same for the corresponding regions on different shapes, we define the following cycle consistency energy:

$$E_{\text{cycle}}(\{\mathbf{H}_k\}, \{\mathbf{C}_{kl}\}) := \sum_{(k,l) \in \mathcal{E}} \|\mathbf{H}_k - \mathbf{C}_{kl} \mathbf{H}_l\|_F^2. \quad (14)$$

Note that a naive solution will collapse to $\mathbf{H}_k = \mathbf{0} \forall k$, so an additional orthonormality constraint is added as $\mathbf{H}^T \mathbf{H} = \mathbf{I}_V$, where $\mathbf{H} := [\mathbf{H}_1^T, \dots, \mathbf{H}_K^T]^T \in \mathbb{R}^{KM \times V}$. This constraint eliminates the trivial solution by *fixing the gauge* and encourages a larger functional space spanned by the canonical functions.

Additionally, we define the data term as the summation of all pairwise energies:

$$E_{\text{data}}(\{\mathbf{C}_{kl}\}) := \sum_{(k,l) \in \mathcal{E}} E_{kl}(\mathbf{C}_{kl}), \quad (15)$$

with E_{kl} defined in Eq. (4). Recall that during test time we set \mathfrak{g} (needed in E_{kl}) to \mathfrak{g}_{pd} without any reliance on ground-truth information. The solution to the problem

$$\arg\min_{\{\mathbf{H}_k\}, \{\mathbf{C}_{kl}\}} E_{\text{cycle}}(\{\mathbf{H}_k\}, \{\mathbf{C}_{kl}\}) + E_{\text{data}}(\{\mathbf{C}_{kl}\}), \quad (16)$$

gives us the synchronized functional maps $\{\mathbf{C}_{kl}^*\}$ as well as one possible set of canonical functions. After the

1. This means to project the Laplacian operator \mathbf{L} to the current bases via $\Phi^+ \mathbf{L} \Phi$.

optimization, we re-use the technique described in Section 4.4 to recover \mathcal{F} as our final output, by feeding in the optimized result $\{\mathbf{C}_{kl}^*\}$ instead of $\{\mathbf{C}_{kl}^0\}$.

Analysis. In our formulation, each summand E_{kl} in E_{data} of Eq. (15) now contains the robust function $\rho(\cdot)$ whose residual not only reflects the alignment with the current pairwise map, but also the consistency among all the maps in \mathcal{G} , related by $\{\mathbf{H}_k\}$. Remarkably, compared to existing point-based synchronization techniques (e.g. [7]), the use of functional maps is *lightweight* reducing the number of variables from the order of $O(N^2)$ to $O(M^2)$ where $M \ll N$. Besides, its additional degree-of-freedom endow us with more *flexibility* in representing non-rigid transformations compared to, e.g. [50].

Numerical Optimization. Inspired by [46], we use an alternating method to solve Eq. (16). Specifically, we iterate over optimizing $\{\mathbf{H}_k\}$ with fixed $\{\mathbf{C}_{kl}\}$ and optimizing $\{\mathbf{C}_{kl}\}$ with fixed $\{\mathbf{H}_k\}$ until convergence. Details of the solver are deferred to Appendix B.

Basis Preconditioning. While we do not regularize the columns of Φ_k to be orthogonal during the training of φ_{basis} , optimizations related to it can be highly ill-posed due to the emergence of near-parallel columns. This is due to the use of pseudo-inverse that implicitly factors out the high-frequency bases from the network output, which by nature tends to produce over-smoothed results. Such a phenomenon, though does not affect network training without synchronization (i.e. Section 4), jeopardizes the convergence and stability when jointly optimizing Eq. (16). In order to mitigate this issue, during test time, we explicitly leverage orthogonality as a means of preconditioning by performing the singular value decomposition (SVD) as $\Phi_k = \mathbf{U}\Sigma\mathbf{V}^T$ and replace Φ_k with \mathbf{U} as the new set of basis functions. Note that Eqs. (5) and (6) are calibrated during training to engulf the original basis scale from the raw network output, hence the multiplier $\Sigma\mathbf{V}^T$ should be applied back after synchronization to recover the final 3D flow vectors.

6 EXPERIMENTS

6.1 Dataset and Settings

Datasets. State-of-the-art methods for non-rigid registration or matching usually evaluate on different datasets with different settings. In order to provide a comprehensive and fair evaluation, we experiment with a full spectrum of possible input data, ranging from partial to full point clouds, from non-rigid human/clothes/animals to multi-body articulated objects, and from synthetic to real-world scenes. To this end, we extend the following datasets [63], [64], [65], [66], [67] to the task of multi-way non-rigid registration (hence the prefix **Multi-Point Cloud** used below). For all the datasets we keep only the points from the foreground deforming objects.

- *MPC-CAPE* is sampled from CAPE [63], [64] dataset, which contains scanned sequences of clothed human using a commercial 3dMD scanner. Raw scans are cropped using a fixed-height horizontal plane to remove the static ground points. We randomly sample non-consecutive K -frame snippets from the sequences and obtain the ground-truth flow from the fitted template mesh skinned via a Linear Blend Skinning [68] scheme.
- *MPC-DT4D* is from the DeformingThings4D [65] dataset, with hundreds of dynamic sequences of

TABLE 1
Dataset Statistics

Dataset	# Train / Val.	# Test	Partial	Non-rigid	Real
MPC-CAPE [63], [64]	3015 / 798	209	-	✓	✓ [†]
MPC-DT4D [65]	3907 / 1707	1299	✓	✓	-
MPC-DD [66]	1754 / 200	267	✓	✓	✓
MPC-SAPIEN [67]	530 / 88	266	-	-	-

We report the number of training, validation and test samples used in the quantitative evaluations.

[†]: For train/val we use the points sampled from the parametric meshes while for test we use the raw-scanned point clouds.

humanoids and animals whose skeletons and motions are designed by experts. We extract the partial point clouds by synthetically scanning the object under different motions. We ensure that the (*instance, action*) tuples never overlap when we create the splits.

- *MPC-DD* [66] contains partial views captured with a real RGB-D camera. The scene flows are provided by the original dataset via non-rigidly tracking the temporally densely-sampled frames. The provided object masks are used to remove the static background points. We use the original validation set as our test set and split the original training set for training and validation, without instance overlap between the splits.
- *MPC-SAPIEN* [67] consists of simulated, realistic, articulated object-level models that are commonly observed in daily life. We follow the data generation and sampling strategies implemented in [7], which ensure that there is no instance overlap between the splits.

Statistics of the dataset splits are shown in Table 1.

Metrics. The main metric for evaluating our flow quality is the (1) L2 Error, also termed as Mean Absolute Error (MAE), or End-Point Error (EPE), i.e., the average norm of the flow error vectors over all points. We also include metrics from the scene flow [69] community but adapt their thresholds to our object-level setting. Specifically, we add: (2) 3D Accuracy Strict (AccS), the percentage of points whose relative error $< 5\%$ or $< 2\text{cm}$, (3) 3D Accuracy Relaxed (AccR), the percentage of points whose relative error $< 10\%$ or $< 5\text{cm}$, and (4) Outlier Ratio, the percentage of points whose relative error $> 30\%$.

In datasets with occlusions, we additionally compute all the metrics both for non-occluded points and the full point clouds. We subsequently collect the statistics among all the $K(K-1)$ pairs of the metrics and report the mean value and standard deviation (\pm), the latter of which shows the consistency of the estimates among all individual pairs, as desired in the multiway setting.

Baselines. We compare the performance of SyNoRiM against a number of baselines grouped by the learning strategy. Hereinafter, we use the underlined names for brevity:

- *Non-learned*: Coherent Point Drift method (CPD) [20], Group-wise unnormalized information potential with Rényi's entropy (GUIP) [70], and Non-Parametric Part (NPP) [71] using up to 25 RJ-MCMC sample steps;
- *Supervised*: Consistent ZoomOut (CZO) [54] where we replace the original Laplacian-Beltrami basis functions with the non-manifold one in [58] (the

TABLE 2
Quantitative Comparison on MPC-CAPE Dataset Using Fully-Supervised and Self-Supervised Schemes

	Supervision	L2 Error ↓	AccS ↑	AccR ↑	Outlier ↓	Run Time (s)
CPD [20]	-	7.06 ± 2.46	2.5 ± 2.1	42.0 ± 20.3	90.8 ± 5.5	1.03
GUIP [70]	-	5.55 ± 2.16	35.8 ± 11.1	69.3 ± 13.4	80.9 ± 6.5	(28.2)
NPP [71]	-	19.02 ± 5.63	1.3 ± 1.7	10.6 ± 8.7	99.3 ± 0.9	(70.0)
CZO [54]	Full	2.04 ± 0.27	64.2 ± 7.5	94.2 ± 1.9	51.1 ± 10.7	(24.2)
3D-CODED [72]	Full + Mesh	2.25 ± 0.58	61.6 ± 9.0	95.4 ± 3.6	53.8 ± 11.2	160 [‡]
PTF [73]	Full + Mesh	2.52 ± 0.54	47.9 ± 12.8	91.4 ± 5.7	57.8 ± 9.6	100
PWC [28]	Full	1.37 ± 0.34	82.2 ± 8.4	98.3 ± 1.6	39.1 ± 10.4	0.99
FLOT [29]	Full	1.79 ± 0.34	69.9 ± 9.3	96.9 ± 2.0	48.4 ± 10.7	1.45
EmbAlign [40]	Full	2.15 ± 0.49	61.5 ± 10.6	95.0 ± 3.5	53.3 ± 10.0	0.60
GeomFMaps [34]	Full	2.52 ± 0.25	54.3 ± 6.2	88.1 ± 2.0	54.4 ± 10.7	1.40
MBS [7]	Full	2.15 ± 0.22	54.0 ± 5.4	97.2 ± 1.8	55.9 ± 11.4	1.12
Ours (φ_{desc} only)	Full	1.39 ± 0.18	82.0 ± 5.3	99.1 ± 0.6	42.4 ± 10.9	0.12
Ours	Full	1.08 ± 0.17	90.1 ± 4.3	99.5 ± 0.5	35.3 ± 10.1	1.22
PWC-U [28]	Self	3.52 ± 1.68	50.0 ± 13.6	82.7 ± 11.7	67.4 ± 8.7	0.99
RMA [74]	Self [†]	5.47 ± 2.29	29.5 ± <u>15.0</u>	67.9 ± 15.5	83.4 ± 7.0	1.60
DeepShells [36]	Self	4.55 ± 3.61	57.1 ± <u>21.6</u>	81.2 ± 20.5	60.6 ± 16.4	58.0
Ours (φ_{desc} only)	Self	4.20 ± 1.44	34.4 ± 9.5	79.1 ± 10.0	71.5 ± 8.5	0.12
Ours (w/o sync.)	Self	3.07 ± 1.37	53.9 ± 15.3	88.1 ± 8.7	60.3 ± <u>11.0</u>	0.32
Ours	Self	2.95 ± 1.29	55.5 ± <u>15.0</u>	89.1 ± 8.0	59.1 ± <u>11.0</u>	1.22

↓ / ↑ : Lower/higher is better (for the standard deviation being lower is always better). Supervision with ‘Mesh’ requires an additional template (or parametric) human mesh as a canonical prior. ‘w/o sync.’ is inferred from the initial functional map matrices without synchronization. Boldfaced numbers highlight the best and underlined numbers mark the tied first. Run time in (parentheses) denotes no GPU acceleration.

[†]: Fails to converge on our dataset if trained from scratch so we finetune from the pretrained model provided by the official repository.

[‡]: Includes the time for best rotation selection and test-time Chamfer loss minimization.

number of bases is set to 80), scene flow estimation methods including FLOT [29] and PointPWC-Net (PWC) [28], MultiBodySync (MBS) [7] where for non-rigid case we remove the motion segmentation module and apply only the weighted permutation synchronization with downsampled 512 points, and EmbAlign [40] which also estimates point cloud bases.

- *Self-supervised*: PointPWC-Net (PWC-U) [28] with the unsupervised loss, and RMA-Net (RMA)[74] with the proposed differentiable rendering loss.

For MPC-CAPE where full human bodies are captured, we additionally compare to template-based methods 3D-CODED [72] and PTF [73], as well as deep-functional-map-based methods GeomFMaps [34] (supervised) and DeepShells [36] (self-supervised), where bases are replaced with [58]. For baselines that do not have a principled way of handling multiple inputs, we treat each pair of the input point clouds independently. Readers are referred to the appendix for more descriptions of the baseline settings.

Parameter Settings. The networks φ_{desc} , φ_{basis} , and φ_{refine} are all trained using AdamW [75] optimizer with the initial learning rate of 10^{-3} and a 30% decay every 50k samples. φ_{desc} and φ_{basis} are both 4-layer U-Net [76] with 32, 96, 64, 192 channels and 32, 64, 128, 256 channels respectively, while φ_{refine} is a 2-layer small network with 32 and 64 channels. All networks use InstanceNorm [77] as normalization. Further details about the network architectures are illustrated in Appendix D, available in the online supplemental material. The initial t in Eq. (5) is chosen to be 0.1, and the weight for \mathcal{L}_c is $\lambda_c = 5 \times 10^{-3}$. For our self-supervised training scheme, we balance the Chamfer/Laplacian/smoothness losses with the weights 1.0/5.0/1.0, respectively. For benchmarking we choose $K = 4$ and use $M = 24$ for all of our experiments

except for the one in Section 6.6. V is set to be $M - 4$ for MPC-CAPE and $M - 2$ for all the other datasets.

6.2 Evaluations on MPC-CAPE

MPC-CAPE is mainly used to measure the fitting tightness and the deformation quality of the warped point clouds. As shown in Table 2, our method achieves a significant performance boost, reaching a 21.1% and 16.2% lower error compared to the nearest baseline under full supervision and self-supervised scenarios, respectively.

Referring to the visualizations in Fig. 4, SyNoRiM is most effective in point clouds with large changes, where traditional methods fail either due to wrong associations caused by matching ambiguity or insufficient degrees of warps. For learning-based methods: 3D-CODED and PTF rely on pre-defined canonical human shape prior, which prevents them from guaranteeing a tight fitting to the input points. CZO builds strict point-level correspondences. Hence, it is not capable of *densifying* the points and is error-prone once the number of outliers from φ_{desc} is large. Moreover, the approximate point cloud bases, used also in GeomFMaps and DeepShells, introduce further drifts. While PWC and FLOT are trained using dense flow annotations and are free from large systematic errors, they still yield noisy estimates and result in shrunk artifacts in the detailed region. Comparably, we use the limited-band bases to effectively regularize the motions, while preserving local rigidity and details. By explicitly projecting point-level *matches* onto smooth functions, the generated warping is more accurate and robust to noise. This is in contrast to EmbAlign, where the descriptors are globally projected onto the bases, largely limiting the distinctiveness of the point features. In the unsupervised scenario, we are able to outperform PWC-U and RMA thanks to the superiority of our matching representation

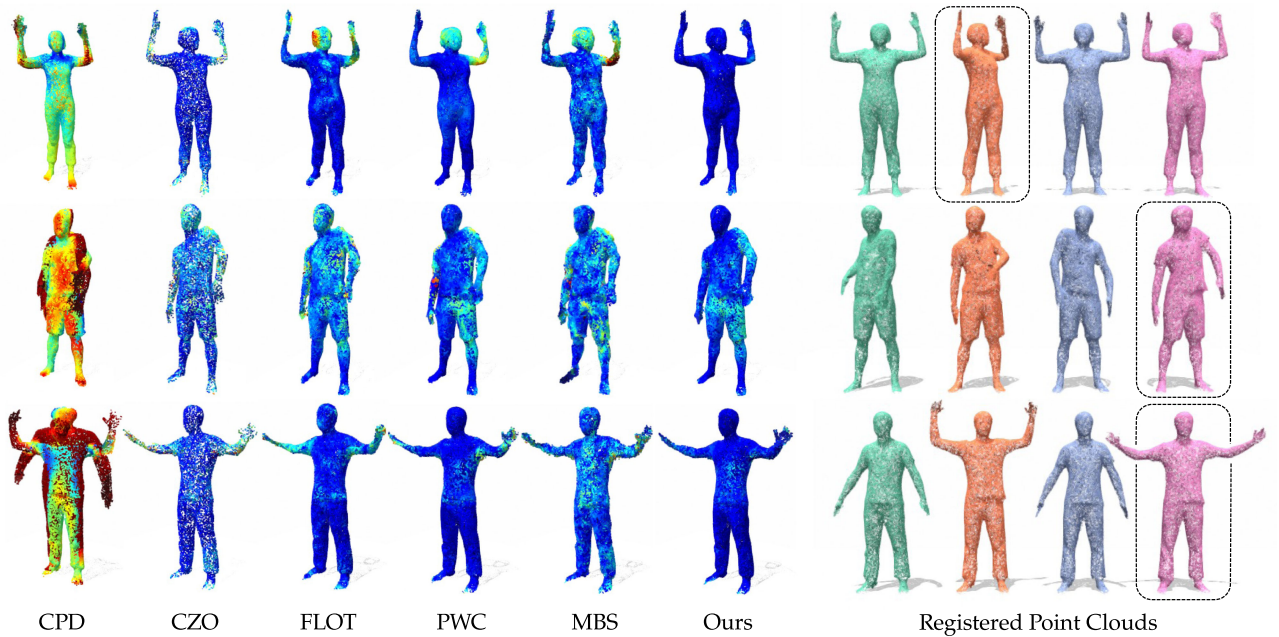


Fig. 4. Comparison of fully-supervised MPC-CAPE. Left: Color-coded per-point errors (0cm \rightarrow 10cm) are shown on the registered point clouds by warping all points to a selected pose. Right: Registered point clouds using our inferred 3D flows. Dashed rectangles are the views to which we warp other point clouds on the left side.

TABLE 3
FAUST Intra-Subject Challenge [78] Results (Top9)

	DHNN [79]	DVM [80]	Ours	SP [81]	AtlasNetV2 [82]	3D-CODED [72]	BPS [83]	LBS-AE [84]	FMNet [35]
Error (cm)	1.011	1.185 [†] / 1.417	<u>1.486</u>	1.568	1.626	1.985	2.010	2.161	2.436

Methods are ordered from left to right according to their rankings.

[†]: Requires multi-view (72) input that requires multiple forward passes of the model. Costs $\sim 7s$ /pair.

and the proposed learning strategy. Moreover, with the help of the synchronization module, we reach a better estimation by pruning inconsistent matches from the graph, reducing the standard deviation of the L2 error by over 5.8%.

FAUST Challenge. To reinforce our results on CAPE, we test our pairwise registration module (without synchronization) on the online benchmark FAUST [78], where ground-truth information is not accessible. We perform the evaluation on the intra-subject challenge, which measures the registration error between 60 pairs of shapes, where each pair consists of two full meshes of the same subject in different poses. We directly test our pretrained model without further finetuning on the points sampled from the input meshes. The results are shown in Table 3. Our method ranks 3rd (as of Mar. 2022) among all published methods. The two methods that surpass ours either require manually-labelled landmarks *during test time* (DHNN [79]), or rely on artist-designed canonicalization of a predefined human model (DVM [80]). We are free of the above assumptions and output the more general scene flow representation applicable to different subject categories.

6.3 Evaluations on MPC-DT4D

Scanned by a virtual camera, MPC-DT4D contains challenging partialities and occlusions, the handling which is an important yet often missing ingredient from many methods. Input shapes that overlap only partially can cause

catastrophic failure for some methods that are incapable of modeling the flow into *empty space*, especially when large deformations are present. Somewhat surprisingly, Table 4, shows that none of the SoTA methods can effectively deal with the hard problem setting presented in this dataset. On the one hand, methods like CZO and FLOT build the point-level correlation matrix (e.g. permutations), which results in large errors on the boundaries of the point cloud. On the other hand, the motion regularization implicitly introduced in CPD and GUIP cannot be easily adapted to the various kinds of input shapes. Moreover, the utilization of consistency on the level of points renders MBS inapplicable under such scenarios.

Comparably, as also shown in Fig. 5, we remark that in our method, the learning of the basis functions constructs a shared pattern of similar shapes under different views, hence producing a good functional extrapolation of point coordinates in a learned manner. By not operating on the level of points, our synchronization module bypasses the occlusion problem via aligning the globally-defined basis functions, successfully reducing the error by 5.4% for the full point cloud. Further challenging examples that can be successfully tackled by our method are visualized in Fig. 6.

6.4 Evaluations on MPC-DD

As shown in Fig. 7, our method consistently outperforms scene-flow based methods (19.4% / 29.2% lower L2 Error and

TABLE 4
Quantitative Comparison on MPC-CAPE Dataset

	<i>Non-Occluded</i>				<i>Full</i>			
	L2 Error ↓	AccS ↑	AccR ↑	Outlier ↓	L2 Error ↓	AccS ↑	AccR ↑	Outlier ↓
CPD [20]	10.12 ± 4.28	11.7 ± 8.9	43.2 ± 18.1	67.3 ± 12.8	10.91 ± 4.69	11.4 ± 8.6	41.8 ± 17.8	68.4 ± 12.3
GUIP [70]	16.43 ± 7.42	20.3 ± 8.9	43.0 ± 13.5	73.6 ± 10.6	17.22 ± 7.66	19.4 ± 8.8	41.2 ± 13.5	74.6 ± 10.2
CZO [54]	4.93 ± 1.62	42.4 ± 12.1	75.2 ± 10.3	37.8 ± 11.3	6.44 ± 2.45	39.8 ± 11.9	70.7 ± 11.2	40.8 ± 11.0
PWC [28]	5.39 ± 2.84	47.5 ± 14.0	77.0 ± 12.2	36.8 ± 12.5	6.35 ± 3.43	44.9 ± 13.9	73.2 ± 13.1	39.3 ± 12.5
FLOT [29]	5.14 ± 1.94	43.2 ± 11.7	75.3 ± 10.0	39.8 ± 10.9	6.57 ± 2.92	40.3 ± 11.6	70.5 ± 11.2	43.0 ± 10.9
EmbAlign [40]	6.97 ± 2.75	36.3 ± 11.4	66.1 ± 11.3	48.1 ± 10.5	8.22 ± 3.56	34.1 ± 11.3	62.3 ± 12.1	50.7 ± 10.4
MBS [7]	5.55 ± 2.45	43.9 ± 11.5	76.3 ± 11.1	38.4 ± 11.7	6.70 ± 3.13	41.2 ± 11.6	72.1 ± 12.1	41.1 ± 11.8
Ours (φ_{desc} only)	3.77 ± 1.54	63.0 ± 9.8	86.5 ± 6.0	28.2 ± 9.9	5.40 ± 2.78	58.2 ± 10.6	81.2 ± 8.4	31.9 ± 10.3
Ours (w/o sync.)	3.04 ± 1.54	69.2 ± 10.7	89.5 ± 6.4	24.1 ± 10.1	3.73 ± 2.05	65.5 ± 11.5	86.3 ± 8.0	26.0 ± 10.4
Ours	2.87 ± 1.25	69.2 ± 10.2	89.8 ± 5.7	23.8 ± 9.7	3.53 ± 1.71	65.5 ± 10.9	86.6 ± 7.2	25.6 ± 9.9

We report results for both the non-occluded set of points and the full point cloud. See caption of table 2 for other legends.




Fig. 5. Comparison of different methods on the MPC-DT4D dataset. The rightmost four columns show color-coded per-point errors (0cm 10cm) on the registered point clouds by warping all points to a selected pose.

2.9% / 11.3% higher AccS than FLOT / PWC, resp.), even with the presence of non-negligible real-world noise, holes, and large motions. We show the application of dynamic reconstruction from sparse views (usually hundreds of frames apart as detailed in the appendix, available in the online supplemental material,) in Fig. 8 and found that our learned basis functions are smooth and robust, despite the difficulties in MPC-DD.

6.5 Evaluations on MPC-SAPIEN

Quantitative and qualitative evaluations on MPC-SAPIEN are shown in Table 5 and Fig. 9. Note, MPC-SAPIEN contains no occlusions and only small *local* deformations. This allows [58] to accurately estimate the point cloud bases, making CZO the most effective out of all the competitors. MBS is specially designed to handle multi-scan multi-body case, and already

decreases the PWC backbone error by half thanks to the rigidity prior. However, the permutation and segmentation synchronization modules of MBS limit the maximum number of input points to hundreds of points, causing the segmentation and causing a large overall error. As visualized earlier in Fig. 3, the basis functions learned from our method already contain motion cues whose patterns are similar to moving part segmentations. Hence, we conclude that our good performance stems mainly from the explicit usage of bases and the synchronization module, which itself contributes an 11.2% and 30.3% improvement over mean L2 Error and its variance, providing a better and more consistent flow estimation.

6.6 Affinity Bases on LidarKITTI

We now conduct an experiment on the cluttered scenes such as roads in autonomous driving to demonstrate the broad

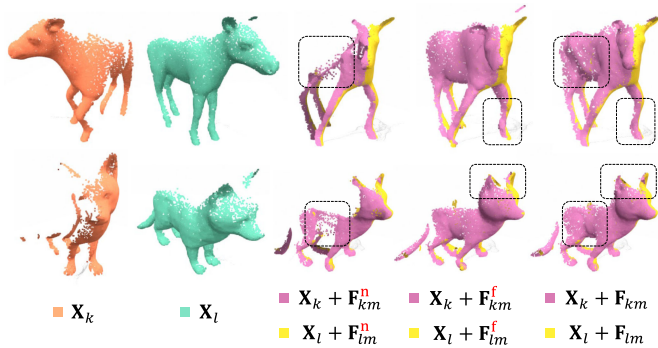


Fig. 6. *Handling low-overlap samples.* We warp two views (X_k, X_l) to the third view (X_m) using different strategies for flow computation from the functional map matrix (See Section 4.4 for notations and analysis). Dashed rectangles highlight the differences.

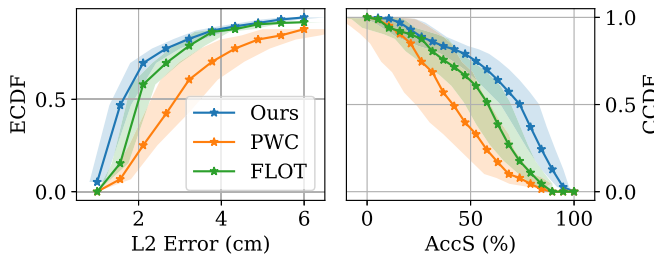


Fig. 7. *Quantitative comparison on MPC-DD dataset.* Left and right sub-figures show the empirical cumulative distribution function (ECDF) and the complementary cumulative distribution function (CCDF) of L2 Error and AccS metrics. Higher curve is better. Shaded areas denote the standard deviations of the respective metrics among views.

applicability of SyNoRiM. To this end, we use the Lidar-KITTI [8] dataset with ground points removed. Note that this dataset contains only pairwise Lidar frames, hence the synchronization module is not applied. The correspondence estimator \mathfrak{g} is trained on FlyingThings3D [85] dataset using the method described in Section 4.2 and the bases are computed according to the *affinity bases* strategy in Section 4.6 with the foreground segmentation module from [8]. We also train our full pairwise registration module using [85] with *learned bases*.

The flow estimation results are listed in Table 6. Note that here the metric ratios use the original thresholds defined in [69]. Compared to our baselines, ours with affinity bases reach a better performance, showing that the basis structure can effectively introduce meaningful priors to the scene flow estimation. Visualized in Fig. 10, the main error comes from the affinely-warped background due to wrong raw correspondences estimated by \mathfrak{g} . In fact, scene flow estimation in the wild achieves improved performance mainly due to more accurate ego-motion estimation as already pointed out in, e.g., [27] and our result can be further improved if such rigidity constraint is added. Though, this is out of our current scope, as our aim is to validate the effectiveness of our affinity bases.

Nevertheless, we find that in such large-scale multi-object scenes, the learned global bases are less accurate than hand-crafted ones. This is because in such scenes, each object moves *independently* and each basis should thus be concentrated on a *single* rigid body. However, the number of bodies vary between scenes and the boundaries are hard




Fig. 8. *Dynamic reconstruction on MPC-DD dataset* by accumulating the input point clouds. Left: RGB image of the reference frame (ref), which are *not* taken as input. Middle: Unprojected depth point cloud from ref . Right: Accumulated points by warping other 3 frames to ref .

TABLE 5
Quantitative Comparison on MPC-SAPIEN Dataset

	L2 Error ↓	AccS ↑	AccR ↑	Outlier ↓
CPD [20]	9.18 ± 4.50	7.4 ± 5.8	37.1 ± 22.7	81.3 ± 13.0
GUIP [70]	10.85 ± 4.28	10.0 ± 8.0	30.4 ± 16.8	84.7 ± 6.9
CZO [54]	3.40 ± 2.12	62.3 ± 14.5	86.9 ± 12.4	24.5 ± 11.8
PWC [28]	8.95 ± 3.86	7.3 ± 7.0	35.6 ± 20.9	81.3 ± 11.9
FLOT [29]	5.17 ± 2.69	32.9 ± 16.0	71.3 ± 18.6	45.6 ± 11.9
EmbAlign [40]	11.12 ± 4.94	7.1 ± 6.1	28.2 ± 17.7	88.6 ± 6.0
MBS [7]	4.93 ± 2.27	34.8 ± 14.4	72.8 ± 16.2	43.6 ± 10.5
Ours (φ_{desc} only)	4.66 ± 2.45	38.4 ± 13.0	78.4 ± 13.9	39.1 ± 9.7
Ours (w/o sync.)	3.48 ± 2.41	61.0 ± 19.0	86.2 ± 13.2	25.1 ± 11.9
Ours	3.09 ± 1.68	62.8 ± 15.3	88.1 ± 9.7	23.7 ± 9.6

See caption of Table 2 for other legends.

to fit. These properties counteract our φ_{basis} network and are challenging to learn. Such problems do not exist in object-level datasets where the geometric structures are simpler.

6.7 Ablation Studies

A. Number of Bases. The number of basis functions M reflects the prior knowledge on the intrinsic motion variability of the input shape. Table 7 shows that as the number of bases increases, the flow estimation error first decreases, but after passing a certain point, rises significantly. This is because the noise in the descriptors mistakenly incorporated into the learned bases start to dominate the increased representation power brought by adding more bases. In the meantime, computing the pseudo-inverse of a wider Φ_k matrix demands stronger numerical stability hampering the convergence of the entire model. In practice, we found that selecting $M = 24$ strikes a good balance between accuracy and computation time of the synchronization, which grows exponentially as M increases.

B. Number of Frames. Our method can be easily applied to more frames ($K > 4$) without re-training. To demonstrate this, we additionally sample $K = 6, 8, 10$ frames (under the same time interval) for MPC-CAPE dataset and test our trained model on them. As shown in Fig. 11, the accuracy of our method only drops marginally when K increases. This is because all of our networks are trained in the pairwise setting and there is no learnable component in the multi-frame synchronization module.




Fig. 9. Comparison of fully-supervised MPC-SAPIEN. Left: Color-coded per-point errors (0.0 to 0.1) are shown on the registered point clouds by warping all points to a selected pose. Right: Registered point clouds using our inferred scene flows for all input poses.

C. Choice of Basis. Different basis functions defined on point clouds are available and we compare SyNoRiM to one of the non-learned methods [58]. Despite being proven by [86] to be *optimal* in representing smooth functions on manifold (in the sense that its gradient magnitudes are bounded), such bases functions are not task-tailored, whereas ours are learned to maximize the scene flow estimation performance and can work robustly under noise. As shown in Table 8, even with 80 bases, the flow recovered using [58] is worse than ours with only 24 bases. Moreover, the basis computation from a sparse convolution backbone accelerated with GPU is highly efficient and benefits from parallel implementation. The reduced number of bases not only makes

the eigen-decomposition faster, but also boosts the efficiency of subsequent synchronization module.

D. Training Scheme of φ_{basis} . We validate the effectiveness of the training strategies applied for φ_{basis} , including (a) the use of the robust term $\rho(\cdot)$, (b) the addition of the consistency loss \mathcal{L}_c and (c) the mixing strategy of \mathbf{g}_{gt} and \mathbf{g} . As demonstrated in Table 9, adding the robust term gives the framework more tolerance towards noise while the consistency loss acts as an efficient regularizer and speeds up network convergence. Apart from the performance gain from the mixing strategy, we empirically found that using only the ground-truth \mathbf{g} causes singularities in the network prediction which gives ill-posed scene flow when two frames are too close.

TABLE 6
Comparisons on LidarKITTl Dataset

	EPE3D ↓	AccS ↑	AccR ↑	Outliers ↓
PointPWC-Net [28]	39.0	38.7	55.0	65.3
FLOT [29]	65.3	15.5	31.3	83.7
Ours (φ_{desc} only)	57.0	31.3	45.6	73.4
Ours (learned bases)	46.2	38.1	59.6	66.5
Ours (affinity bases)	31.5	34.8	55.2	70.4
Rigid3DSceneFlow [8]	15.3	48.5	73.0	46.6

Rigid3DSceneFlow [8] exploits explicit rigid ego-motion estimation for the background scene flow. (φ_{desc} only) simply maps each point to its nearest neighbor in the feature space embedded by φ_{desc} .

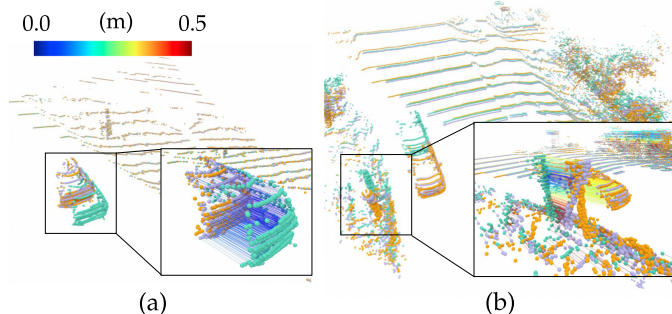


Fig. 10. Results on LidarKITTl dataset. (a) and (b) show the best and worst result estimated using our method, respectively (L2 Errors are 1.7cm and 87.7cm). ■ Source points; ■ Target points; ■ Warped points.

Authorized licensed use limited to: Johns Hopkins University. Downloaded on October 18, 2023 at 01:32:42 UTC from IEEE Xplore. Restrictions apply.

TABLE 7
Performance Using Different Number of Bases on MPC-CAPE Dataset

# Basis M	L2 Error ↓	Run Time (sec./iter.)
8	1.14	6.0×10^{-3}
16	1.11	1.2×10^{-2}
24	1.08	2.6×10^{-2}
32	1.05	7.0×10^{-2}
64	1.27	1.9×10^0
80	1.55	5.9×10^0

For run time we measure the average time for each iteration of our synchronization algorithm, in the unit of seconds.




Fig. 11. Influence of frame count K . Changes of L2 Error and AccS w.r.t. the number of frames K are shown separately in the two plots. Dark green thin curves show the average metric (Avg.) while light green bold curves show the standard deviation (+/-). The black arrows highlight the changes of the average metrics from $K = 4$ to $K = 10$.

TABLE 8
Comparison of Our Learned Bases With [58]
on MPC-CAPE Dataset

	# Basis	L2 Error ↓	AccS ↑	Run Time (sec.)	
				Basis	Sync. per iter.
[58]	24	2.96 ± 0.42	48.4 ± 8.2	(1.6)	0.03
	80	1.53 ± 0.21	78.4 ± 5.1	(3.2)	5.9
Ours	24	1.08 ± 0.17	90.1 ± 4.3	0.24	0.03

Run time in (parentheses) denotes no GPU acceleration.

TABLE 9
Comparisons of Different φ_{basis} Training Strategies
on MPC-DT4D Dataset

(a) $\rho(\cdot)$	(b) \mathcal{L}_c	(c) Mixed g	L2 Error F^{n} ↓	L2 Error F^{f} ↓
			3.64 / 5.62	4.33 / 6.05
✓			3.58 / 5.48	4.13 / 5.78
✓	✓		3.52 / 5.41	4.16 / 5.78
✓		✓	3.49 / 5.38	4.11 / 5.74
✓	✓	✓	3.23 / 4.76	4.19 / 5.49
✓	✓	✓	3.14 / 4.59	3.86 / 5.09

See text for the explanation of (a),(b),(c). The values of EPE3D are measured on (non-occluded / full) point clouds.

E. Refinement Strategy. We show that the combination of flows in Eq. (7) is indeed helpful, especially in the occluded regions. As illustrated in Table 10, feeding the three pieces of information altogether helps the refinement module achieve the best result. X acts as a regularizer by reducing gross errors. Despite its decisive effect on flow accuracy, F^{n} is, by definition, not capable of transferring points to the empty region, where F^{f} comes to the rescue by extending the target mapped positions to the full 3D space (c.f. Fig. 6 for visual comparisons).

F. Generalization. To further demonstrate how our model generalizes to unseen domains, we re-split the animals part of the MPC-DT4D dataset into two disjoint category sets \mathcal{C}_1 and \mathcal{C}_2 (detailed in the appendix). Results in Table 11 show that our model trained only on \mathcal{C}_1 can reach a similar performance to the model trained on $\mathcal{C}_1 \cup \mathcal{C}_2$ when tested on \mathcal{C}_2 (note the actions of the training and test set never overlap, denoted by $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$). This empirically verifies that our method can generalize across different animal breeds. However, our method fails to generalize from humanoids (denoted as \mathcal{C}_{H}) to novel categories with drastically different geometries like \mathcal{C}_2 . As echoed in Section 7, this probably stems from the fact that our descriptor network and basis network overfit to the global structural information.

G. Effect of Synchronization. Synchronization links the information from multiple input point clouds exploiting cycle consistency and reduces the overall error. Here, we ablate different optimization strategies in Table 12 using MPC-SAPIEN dataset. The use of the robust term $\rho(\cdot)$, different from the one used for pairwise training, now jointly considers multi-scan cues and prunes the inconsistent outliers which is hard to detect in the pairwise setting. Although the basis preconditioning (Basis Precond.) strategy does not result in a large performance difference at convergence, it has a positive impact on the convergence rate. As a direct

TABLE 10
Comparisons of Different Refinement Strategies on MPC-DT4D
Dataset to Generate the Final 3D Scene Flow

Input + F^{n}	Input + F^{f}	Input + X	L2 Error ↓
(F ⁿ directly as output without φ_{refine})			3.14 / 4.59
✓			3.09 / 3.91
	✓		3.47 / 4.37
	✓	✓	3.48 / 4.26
✓		✓	3.04 / 3.83
✓	✓		3.04 / 3.81
✓	✓	✓	3.04 / 3.73

L2 Errors are measured on (non-occluded / full) point clouds.

TABLE 11
Performance of Generalization on Different Splits
of the Animals Subset in MPC-DT4D Dataset

Training set			Test on ($\mathcal{C}_2, \mathcal{A}_3$)			
($\mathcal{C}_1, \mathcal{A}_1$)	($\mathcal{C}_2, \mathcal{A}_2$)	($\mathcal{C}_{\text{H}}, \mathcal{A}_{\text{H}}$)	L2 Error ↓	Δ ↓	AccS ↑	Δ ↑
✓			3.16	0.0%	63.5	0.0
✓	✓		3.32	+5.06%	61.9	-1.6
		✓	10.22	+223%	34.3	-29.2

Checkmarks indicate inclusion for training. L2 Error and AccS are measured on full point clouds w/o synchronization.

TABLE 12
Comparison of Different Synchronization Methods
on MPC-SAPIEN Dataset

	Fixed Iter. L2 Error ↓	Converged Iter.		
		L2 Error ↓	Time (s)	# Iter (k)
Init. (w/o sync.)	3.48 ± 2.41	-	-	-
Spectral [7]	3.24 ± 1.57	-	-	-
w/o $\rho(\cdot)$	3.31 ± 1.90	3.18 ± 1.79	4.2	0.2
w/o Basis Precond.	3.17 ± 1.87	3.03 ± 1.57	16.7	0.72
w/o Alter. Optim.	3.27 ± 1.87	3.00 ± 1.55	(45.5)	1.0
Ours Full	3.09 ± 1.68	3.03 ± 1.56	7.0	0.32

'Fixed Iter.' denotes optimizing for a fixed number of iterations, while 'Converged Iter.' means to iterate until convergence. 'Time' measures the full run time of the synchronization module.

replacement to our alternating optimization scheme (Alter. Optim.), we borrow the solver from [87] to jointly optimize H and C on the product manifold of Stiefel and Euclidean. At convergence, this direct optimizer reaches a slightly lower error. Nevertheless, a lower flow error does not necessarily indicate a better convergence of Eq. (16) and vice versa. We conclude, though, that our optimization scheme can reach a decent performance with much better efficiency than its counterparts. Please see Fig. 12 for visual comparisons.

Furthermore, to show the effect under large view differences, we create a MPC-CAPE-360 dataset, where a moving camera is circling around the performing human actor, and captures 5 frames in between. All individual views are partial, but combined, they cover the full body. After an additional post-processing step of rigidly aligning the point clouds from camera space to world space, our method can successfully align different views and assemble a complete body, as illustrated in Fig. 13. The synchronization module

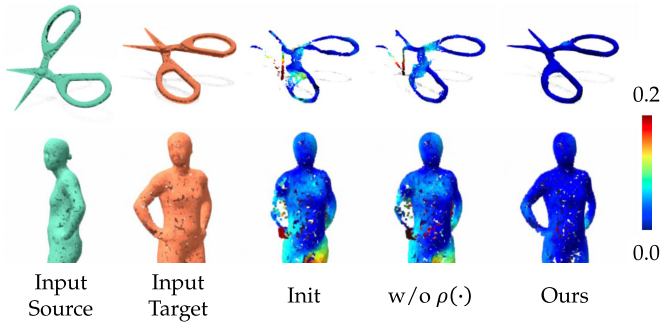


Fig. 12. *Visual comparison of the synchronization effect.* Flow error is shown on the warped source point cloud. The last 2 columns are computed using additional inputs (not shown here) for synchronization.





Fig. 13. *Qualitative results on MPC-CAPE-360.* Input column shows overlaid point clouds from all 5 views (different colors denote different views), while *Before Sync.* and *After Sync.* show the warped point clouds obtained from our full pipeline, before and after synchronization, respectively.

closes the loop by jointly considering multiple views and corrects the alignment errors, which are hard to tackle under the pairwise setting.

6.8 Timing and Memory Analysis



We evaluate the efficiency of SyNoRiM on a workstation with a 3.70GHz Intel i9 CPU with a single GeForce RTX 3090 graphics card. On average it takes 1.22s, 1.86s, 2.80s, 3.89s to generate 12, 20, 30, 42 pairs of flow vectors for $K = 4, 5, 6, 7$ (i.e., roughly 100ms/pair), respectively. The detailed break-down of the relative time cost by each component is shown in the pie chart above: while the refinement module (■) takes more time than φ_{basis} (■) and computing g (●) when K gets larger (because the number of pairs grow quadratically), the synchronization module (■) is always the one that dominates the computation. We find that the training of both φ_{basis} and φ_{refine} usually reaches a full converge within two days, while baselines like PWC take more than 5 days. The parameters for φ_{basis} , φ_{desc} and φ_{refine} are 8.8M, 5.0M and 0.7M, respectively. Runtime memory footprint for the entire pipeline is around 2.8GiB for $K = 4$.

Authorized licensed use limited to: Johns Hopkins University. Downloaded on




Fig. 14. *Non-rigid reconstruction.* We show demonstrative results on the dataset from [88]. Reconstructions are obtained via [89].

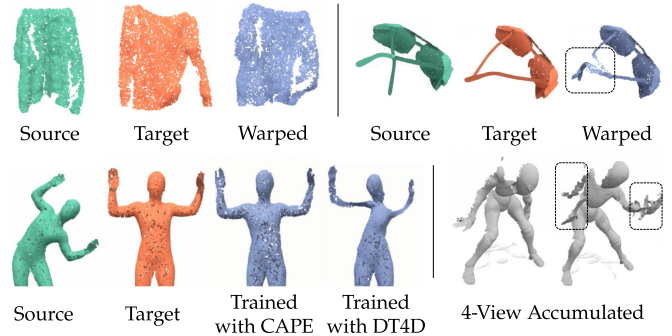


Fig. 15. *Failure Cases.* Large topological changes and matching ambiguities fail SyNoRiM, suffering limited cross-domain generalizability.

7 CONCLUSION

In this paper we present SyNoRiM, a novel method for multiway non-rigid point cloud registration. In our framework, we learn the basis functions defined on the domain of point clouds and represent the registration under the framework of functional maps, endowing us with both efficiency and flexibility. The maps are then digested by our robust synchronization module to incorporate cycle-consistency constraints into the solution, producing more accurate scene flow estimates. A potential application of non-rigid reconstruction is shown in Fig. 14. Our method leverages ideas from computer graphics community concerned with processing clean geometry and applies these ideas to the task of computer vision where occlusions and noise are inevitable. The good performance of SyNoRiM on a full spectrum of cases gives hope for future works in bridging the gap between these two schools.

Limitations. Despite the state-of-the-art performance, SyNoRiM also has some limitations which we point out in the following and illustrate in Fig. 15: (1) the estimated bases sometimes overly-regularize the motion under large displacements and topological changes, causing wrong parts to be stitched together; (2) SyNoRiM is less robust to handling motion ambiguities caused by severe occlusions — in the extreme case C will be ill-conditioned if the corresponding functional coordinates amount to zero; and (3) we find that the trained network has difficulties to generalize to new domains (e.g. from humans to animals/objects). This is partially due to the global field-of-view we employed in our basis network being easily fitted to the entire object structure. (4) Due to the large and possibly ambiguous solution space of non-rigid registration, we limit all the datasets to a

relatively small view change, in line with the weakly-supervised setting in [90]. Future studies on transformation in-/equi-variant backbones may help relax such an assumption. (5) Performance-wise, the synchronization module requires $O(K^2)$ pairwise maps as input and the computational cost grows quadratically w.r.t. the number of views. The root cause of this is the fully-connected graph \mathcal{G} we have assumed: A principled way to build and solve on a sparsified graph \mathcal{G} , or to use batch-wise or incremental synchronization algorithms are good directions to take.

Future Work. Besides addressing aforementioned limitations, SyNoRiM leaves ample room for future works, including better representation of the deformation flow and incorporating more priors to capture complicated motions.

REFERENCES

- [1] M. Dou, H. Fuchs, and J.-M. Frahm, "Scanning and tracking dynamic objects with commodity depth cameras," in *Proc. IEEE Int. Symp. Mixed Augmented Reality*, 2013, pp. 99–106.
- [2] A. Collet *et al.*, "High-quality streamable free-viewpoint video," *ACM Trans. Graph.*, vol. 34, pp. 1–13, 2015.
- [3] R. A. Newcombe, D. Fox, and S. M. Seitz, "Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 343–352.
- [4] M. Slavcheva, M. Baust, and S. Ilic, "Sobolevfusion: 3D reconstruction of scenes undergoing free non-rigid motion," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 2646–2655.
- [5] T. Birdal, U. Şimşekli, M. O. Eken, and S. Ilic, "Bayesian pose graph optimization via bingham distributions and tempered geodesic MCMC," in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst.*, 2018, Art. no. 306317.
- [6] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon, "Bundle adjustment—A modern synthesis," in *Proc. Int. Workshop Vis. Algorithms*, 1999, pp. 298–372.
- [7] J. Huang *et al.*, "Multibodysync: Multi-body segmentation and motion estimation via 3D scan synchronization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 7108–7118.
- [8] Z. Gojcic, O. Litany, A. Wieser, L. J. Guibas, and T. Birdal, "Weakly supervised learning of rigid 3D scene flow," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 5692–5703.
- [9] O. Sorkine and M. Alexa, "As-rigid-as-possible surface modeling," in *Proc. Symp. Geometry Process.*, 2007, pp. 109–116.
- [10] A. L. Yuille and N. M. Grzywacz, "The motion coherence theory," in *Proc. IEEE Int. Conf. Comput. Vis.*, 1988, pp. 344–353.
- [11] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas, "Functional maps: A flexible representation of maps between shapes," *ACM Trans. Graph.*, vol. 31, no. 4, pp. 1–11, 2012.
- [12] R. Huang, F. Chazal, and M. Ovsjanikov, "On the stability of functional maps and shape difference operators," *Comput. Graph. Forum*, vol. 37, no. 1, pp. 145–158, 2018.
- [13] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-squares fitting of two 3-D point sets," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 9, no. 5, pp. 698–700, Sep. 1987.
- [14] S. Ao, Q. Hu, B. Yang, A. Markham, and Y. Guo, "Spinnet: Learning a general surface descriptor for 3D point cloud registration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 11 753–11 762.
- [15] S. Huang, Z. Gojcic, M. Usvyatsov, A. Wieser, and K. Schindler, "Predator: Registration of 3D point clouds with low overlap," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 4267–4276.
- [16] Y. Zhao, T. Birdal, J. E. Lenssen, E. Menegatti, L. Guibas, and F. Tombari, "Quaternion equivariant capsule networks for 3D point clouds," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 1–19.
- [17] H. Deng, T. Birdal, and S. Ilic, "PPFNet: Global context aware local features for robust 3D point matching," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 195–205.
- [18] H. Deng, T. Birdal, and S. Ilic, "3D local features for direct pairwise registration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 3244–3253.
- [19] B. Amberg, S. Romdhani, and T. Vetter, "Optimal step nonrigid ICP algorithms for surface registration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2007, pp. 1–8.
- [20] A. Myronenko and X. Song, "Point set registration: Coherent point drift," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 12, pp. 2262–2275, Dec. 2010.
- [21] Y. Yao, B. Deng, W. Xu, and J. Zhang, "Quasi-newton solver for robust non-rigid registration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 7600–7609.
- [22] R. W. Sumner, J. Schmid, and M. Pauly, "Embedded deformation for shape manipulation," *ACM Trans. Graph.*, vol. 26, no. 3, 2007, Art. no. 80–es.
- [23] A. Bozic, P. Palafox, M. Zollhofer, J. Thies, A. Dai, and M. Nießner, "Neural deformation graphs for globally-consistent non-rigid reconstruction," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 1450–1459.
- [24] J. P. Costeira and T. Kanade, "A multibody factorization method for independently moving objects," *Int. J. Comput. Vis.*, vol. 29, no. 3, pp. 159–179, 1998.
- [25] J. Huang, S. Yang, Z. Zhao, Y.-K. Lai, and S.-M. Hu, "Clusterslam: A slam backend for simultaneous rigid body clustering and motion estimation," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 5875–5884.
- [26] J. Huang, S. Yang, T.-J. Mu, and S.-M. Hu, "Clustervo: Clustering moving instances and estimating visual odometry for self and surroundings," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 2168–2177.
- [27] S. A. Baur, F. Moosmann, P. Pinggera, B. Ommer, and A. Geiger, "SLIM: Self-supervised LiDAR scene flow and motion segmentation," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 13 126–13 136.
- [28] W. Wu, Z. Y. Wang, Z. Li, W. Liu, and L. Fuxin, "PointPWC-net: Cost volume on point clouds for (self-) supervised scene flow estimation," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 88–107.
- [29] G. Puy, A. Boulch, and R. Marlet, "FLOT: Scene flow on point clouds guided by optimal transport," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 527–544.
- [30] R. Li, G. Lin, T. He, F. Liu, and C. Shen, "HCRF-flow: Scene flow from point clouds with continuous high-order CRFs and position-aware flow embedding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 364–373.
- [31] B. Ouyang and D. Raviv, "Occlusion guided self-supervised scene flow estimation on 3D point clouds," 2021, *arXiv:2104.04724*.
- [32] S. Melzi, J. Ren, E. Rodolà, A. Sharma, P. Wonka, and M. Ovsjanikov, "Zoomout: Spectral upsampling for efficient shape correspondence," *ACM Trans. Graph.*, vol. 38, no. 6, 2019, Art. no. 155.
- [33] O. Litany, E. Rodolà, A. M. Bronstein, and M. M. Bronstein, "Fully spectral partial shape matching," *Comput. Graph. Forum*, vol. 36, no. 2, pp. 247–258, May 2017.
- [34] N. Donati, A. Sharma, and M. Ovsjanikov, "Deep geometric functional maps: Robust feature learning for shape correspondence," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 8592–8601.
- [35] O. Litany, T. Remez, E. Rodolà, A. M. Bronstein, and M. M. Bronstein, "Deep functional maps: Structured prediction for dense shape correspondence," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 5659–5667.
- [36] M. Eisenberger, A. Toker, L. Leal-Taixé, and D. Cremers, "Deep shells: Unsupervised shape correspondence with optimal transport," in *Proc. Annu. Conf. Neural Inf. Process. Syst.*, 2020, pp. 10491–10502.
- [37] J.-M. Roufousse, A. Sharma, and M. Ovsjanikov, "Unsupervised deep learning for structured shape matching," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 1617–1627.
- [38] O. Halimi, O. Litany, E. Rodolà, A. M. Bronstein, and R. Kimmel, "Unsupervised learning of dense shape correspondence," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4370–4379.
- [39] R. Magnat and M. Ovsjanikov, "Dwks: A local descriptor of deformations between meshes and point clouds," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 3793–3802.
- [40] R. Marin, M.-J. Rakotosaona, S. Melzi, and M. Ovsjanikov, "Correspondence learning via linearly-invariant embedding," in *Proc. Annu. Conf. Neural Inf. Process. Syst.*, 2020, pp. 1608–1620.
- [41] N. Sharp, S. Ataiiki, K. Crane, and M. Ovsjanikov, "DiffusionNet: Discretization agnostic learning on surfaces," *ACM Trans. Graph.*, vol. 41, 2021, Art. no. 27.
- [42] C. Tang, L. Yuan, and P. Tan, "LSM: Learning subspace minimization for low-level vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 6234–6245.
- [43] N. Ye, C. Wang, H. Fan, and S. Liu, "Motion basis learning for unsupervised deep homography estimation with subspace projection," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 13 117–13 125.

- [44] R. Hartley, J. Trumpf, Y. Dai, and H. Li, "Rotation averaging," *Int. J. Comput. Vis.*, vol. 103, no. 3, pp. 267–305, 2013.
- [45] T. Birdal, M. Arbel, U. Simsekli, and L. J. Guibas, "Synchronizing probability measures on rotations via optimal transport," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 1569–1579.
- [46] F. Wang, Q. Huang, and L. J. Guibas, "Image co-segmentation via consistent functional maps," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2013, pp. 849–856.
- [47] T. Birdal and U. Simsekli, "Probabilistic permutation synchronization using the riemannian structure of the birkhoff polytope," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 11 105–11 116.
- [48] T. Birdal, V. Golyanik, C. Theobalt, and L. J. Guibas, "Quantum permutation synchronization," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 13 122–13 133.
- [49] X. Huang, Z. Liang, and Q. Huang, "Uncertainty quantification for multi-scan registration," *ACM Trans. Graph.*, vol. 39, no. 4, pp. 130:1–130:24, 2020.
- [50] Z. Gojic, C. Zhou, J. D. Wegner, L. J. Guibas, and T. Birdal, "Learning multiview 3D point cloud registration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 1759–1769.
- [51] H. Zhang, O. Van Kaick, and R. Dyer, "Spectral mesh processing," *Comput. Graph. Forum*, vol. 29, no. 6, pp. 1865–1894, 2010.
- [52] Q. Huang, F. Wang, and L. Guibas, "Functional map networks for analyzing and exploring large shape collections," *ACM Trans. Graph.*, vol. 33, no. 4, pp. 1–11, 2014.
- [53] R. Huang, P. Achlioptas, L. Guibas, and M. Ovsjanikov, "Limit shapes—a tool for understanding shape differences and variability in 3D model collections," *Comput. Graph. Forum*, vol. 38, no. 5, 2019, pp. 187–202.
- [54] R. Huang, J. Ren, P. Wonka, and M. Ovsjanikov, "Consistent zoomout: Efficient spectral map synchronization," *Comput. Graph. Forum*, vol. 39, no. 5, pp. 265–278, 2020.
- [55] M. Gao, Z. Lahner, J. Thunberg, D. Cremers, and F. Bernard, "Isometric multi-shape matching," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 14 183–14 193.
- [56] M. Ovsjanikov et al., "Computing and processing correspondences with functional maps," in *Proc. SIGGRAPH ASIA Courses*, 2016, pp. 1–60.
- [57] C. Choy, J. Gwak, and S. Savarese, "4D spatio-temporal convnets: Minkowski convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 3075–3084.
- [58] N. Sharp and K. Crane, "A laplacian for nonmanifold triangle meshes," *Comput. Graph. Forum*, vol. 39, no. 5, 2020, pp. 69–80.
- [59] S. Bengio, O. Vinyals, N. Jaitly, and N. Shazeer, "Scheduled sampling for sequence prediction with recurrent neural networks," in *Proc. 28th Int. Conf. Neural Inf. Process. Syst.*, 2015, pp. 1171–1179.
- [60] P. J. Huber, "Robust estimation of a location parameter," in *Breakthroughs in Statistics*. Berlin, Germany: Springer, 1992, pp. 492–518.
- [61] F. R. Hampel, "The influence curve and its role in robust estimation," *J. Amer. Statist. Assoc.*, vol. 69, no. 346, pp. 383–393, 1974.
- [62] D. Nogneng and M. Ovsjanikov, "Informative descriptor preservation via commutativity for shape matching," *Comput. Graph. Forum*, vol. 36, no. 2, pp. 259–267, 2017.
- [63] Q. Ma et al., "Learning to dress 3D people in generative clothing," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 6469–6478.
- [64] G. Pons-Moll, S. Pujades, S. Hu, and M. Black, "Clothcap: Seamless 4D clothing capture and retargeting," *ACM Trans. Graph.*, vol. 36, no. 4, pp. 1–15, 2017.
- [65] Y. Li, H. Takehara, N. Taketomi, B. Zheng, and M. Nießner, "4DComplete: Non-rigid motion estimation beyond the observable surface," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 12 706–12 716.
- [66] A. Bozic, M. Zollhöfer, C. Theobalt, and M. Nießner, "Deepdeform: Learning non-rigid RGB-D reconstruction with semi-supervised data," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 7000–7010.
- [67] F. Xiang et al., "SAPIEN: A simulated part-based interactive environment," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 11 097–11 107.
- [68] A. Jacobson, Z. Deng, L. Kavan, and J. Lewis, "Skinning: Real-time shape deformation," in *Proc. ACM SIGGRAPH Courses*, 2014, Art. no. 24.
- [69] X. Liu, C. R. Qi, and L. J. Guibas, "Flownet3D: Learning scene flow in 3D point clouds," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 529–537.
- [70] L. G. S. Giraldo, E. Hasanbelliu, M. Rao, and J. C. Príncipe, "Group-wise point-set registration based on rényi's second order entropy," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 2454–2462.
- [71] D. S. Hayden, J. Pacheco, and J. W. Fisher, "Nonparametric object and parts modeling with lie group dynamics," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 7424–7433.
- [72] T. Groueix, M. Fisher, V. G. Kim, B. C. Russell, and M. Aubry, "3D-coded: 3D correspondences by deep deformation," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 230–246.
- [73] S. Wang, A. Geiger, and S. Tang, "Locally aware piecewise transformation fields for 3D human mesh registration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 7639–7648.
- [74] W. Feng, J. Zhang, H. Cai, H. Xu, J. Hou, and H. Bao, "Recurrent multi-view alignment network for unsupervised surface registration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 10 297–10 307.
- [75] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," 2017, *arXiv:1711.05101*.
- [76] C. Choy, W. Dong, and V. Koltun, "Deep global registration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 2514–2523.
- [77] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance normalization: The missing ingredient for fast stylization," 2016, *arXiv:1607.08022*.
- [78] F. Bogo, J. Romero, M. Loper, and M. J. Black, "FAUST: Dataset and evaluation for 3D mesh registration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 3794–3801.
- [79] B. Jiang, J. Zhang, J. Cai, and J. Zheng, "Disentangled human body embedding based on deep hierarchical neural network," *IEEE Trans. Vis. Comput. Graphics*, vol. 26, no. 8, pp. 2560–2575, Aug. 2020.
- [80] H. Kim, J. Kim, J. Kam, J. Park, and S. Lee, "Deep virtual markers for articulated 3d shapes," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 11 615–11 625.
- [81] S. Zuffi and M. J. Black, "The stitched puppet: A graphical model of 3d human shape and pose," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 3537–3546.
- [82] T. Deprelle, T. Groueix, M. Fisher, V. Kim, B. Russell, and M. Aubry, "Learning elementary structures for 3d shape generation and matching," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 7435–7445.
- [83] S. Prokudin, C. Lassner, and J. Romero, "Efficient learning on point clouds with basis point sets," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 4332–4341.
- [84] C.-L. Li, T. Simon, J. Saragih, B. Póczos, and Y. Sheikh, "LBS autoencoder: Self-supervised fitting of articulated meshes to point clouds," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 11 967–11 976.
- [85] N. Mayer et al., "A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 4040–4048.
- [86] Y. Aflalo, H. Brezis, and R. Kimmel, "On the optimality of shape and data representation in the spectral domain," *SIAM J. Imag. Sci.*, vol. 8, no. 2, pp. 1141–1160, 2015.
- [87] J. Townsend, N. Koepf, and S. Weichwald, "Pymanopt: A python toolbox for optimization on manifolds using automatic differentiation," *J. Mach. Learn. Res.*, vol. 17, no. 137, pp. 1–5, 2016.
- [88] M. Slavcheva, M. Baust, D. Cremers, and S. Ilic, "Killingfusion: Non-rigid 3D reconstruction without correspondences," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1386–1395.
- [89] M. Kazhdan and H. Hoppe, "Screened poisson surface reconstruction," *ACM Trans. Graph.*, vol. 32, no. 3, pp. 1–13, 2013.
- [90] A. Sharma and M. Ovsjanikov, "Weakly supervised deep functional maps for shape matching," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 19 264–19 275.



Jiahui Huang received the BS degree in computer science and technology from Tsinghua University, in 2018. He is currently working toward the PhD degree with the Department of Computer Science and Technology, Tsinghua University. His research interests include computer vision, robotics and computer graphics. For more information, please visit <https://cg.cs.tsinghua.edu.cn/people/~huangjh/>.



Tolga Birdal is a postdoctoral researcher with the Geometric Computing Group of Stanford University. He has defended his masters and Ph.D. theses at the Computer Vision Group, CAMP Chair, TU Munich. He was also a Doktorand at Siemens AG. His current foci of interest involve geometric machine learning and 3D computer vision. More theoretical work is aimed at investigating the limits in geometric computing, non-Euclidean inference and principles of deep learning. For more information, please visit <https://www.tbirdal.me>.



Leonidas J. Guibas (Senior member, IEEE) received the PhD degree from Stanford University, in 1976, under the supervision of Donald Knuth. His main subsequent employers were Xerox PARC, MIT, and DEC/SRC. Since 1984, he has been at Stanford University, where he is a professor of computer science. His research interests include computational geometry, geometric modeling, computer graphics, computer vision, sensor networks, robotics, and discrete algorithms. He is a senior member of the IEEE Computer Society.



Zan Gojcic received the PhD degree from ETH Zurich, in 2021, under the supervision of Andreas Wieser. He is currently a research scientist with the NVIDIA AI Lab, which he joined in 2021. His research interests include revolve around 3D computer vision, domain adaptation, and reducing the supervision. For more information, please visit <https://zgojcic.github.io/>.



Shi-Min Hu (Senior Member, IEEE) received the PhD degree from Zhejiang University, in 1996. He is currently a professor in computer science with Tsinghua University. His research interests include geometry processing, image & video processing, rendering, computer animation and CAD. He has published more than 100 papers in journals and refereed conferences. He is editor-in-chief of Computational Visual Media, and on the editorial boards of several journals, including Computer Aided Design and Computer & Graphics.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.