

MultiBodySync: Multi-Body Segmentation and Motion Estimation via 3D Scan Synchronization

Jiahui Huang¹ He Wang² Tolga Birdal² Minhyuk Sung³ Federica Arrigoni⁴
Shi-Min Hu¹ Leonidas Guibas²
¹Tsinghua University ²Stanford University ³KAIST ⁴University of Trento

Abstract

We present *MultiBodySync*, a novel, end-to-end trainable multi-body motion segmentation and rigid registration framework for multiple input 3D point clouds. The two non-trivial challenges posed by this multi-scan multi-body setting that we investigate are: (i) guaranteeing correspondence and segmentation consistency across multiple input point clouds capturing different spatial arrangements of bodies or body parts; and (ii) obtaining robust motion-based rigid body segmentation applicable to novel object categories. We propose an approach to address these issues that incorporates spectral synchronization into an iterative deep declarative network, so as to simultaneously recover consistent correspondences as well as motion segmentation. At the same time, by explicitly disentangling the correspondence and motion segmentation estimation modules, we achieve strong generalizability across different object categories. Our extensive evaluations demonstrate that our method is effective on various datasets ranging from rigid parts in articulated objects to individually moving objects in a 3D scene, be it single-view or full point clouds. Code at <https://github.com/huangjh-pub/multibody-sync>.

1. Introduction

Motion analysis in dynamic point clouds is an emerging area, required by various applications such as surveillance, autonomous driving, and robotic manipulation. Our human-made environments are dominated by *rigid body* movements, ranging from articulated objects to solids like furniture or vehicles. These settings require us to address rigid motions of objects or object parts – which is often referred to as the *multi-body* motion estimation problem. Despite its importance, previous work has mainly focused on specific scenarios with known category semantics, like category-level articulated object segmentation [41], indoor scene instance relocalization [65], or car movement detection [72], leaving the literature of generic motion segmentation relatively unexplored.




Figure 1. MultiBodySync. Given an unordered set of point clouds, we simultaneously segment individual moving rigid parts/objects and register them. The transformed point clouds according to the first scan are aggregated and shown in the middle column. Note that the algorithm *does not* use color information and the right column is shown just for visualization.

Different from traditional single scan analysis algorithms like semantic segmentation [39], the most challenging part in multi-body motion analysis is to disambiguate and distinguish rigid bodies. There, we are naturally required to jointly process and relate *multiple* inputs, to effectively find *consistent* motion-based part/object segmentations as well as point correspondences to enable a multi-way registration. It is even more challenging when the capture is not temporally dense, *i.e.*, an intermittent acquisition that does not follow a stream such as a video, and might contain large pose variations, hampering naive temporal tracking.

In this paper, we introduce a multi-scan multi-body segmentation and motion estimation problem, where the goal is to simultaneously discover and register rigid bodies from multiple scans, represented either as full or partial point clouds, where objects come from unseen categories. As an effective solution, we present MultiBodySync, a fully end-to-end trainable deep declarative architecture [23] able to process an arbitrary number of unordered point sets. As shown in Fig. 1, given a set of scans, MultiBodySync begins relating pairs of scans via *3D scene flow* [78, 64] and con-

confidence estimation. Then, the following two differentiable (permutation and segmentation) synchronization modules, which are central to our approach, respectively enforce the consistency of pairwise point correspondences and motion segmentation labelings across different scans. Our design explicitly decouples geometry and motion, making MultiBodySync generalizable to *unseen categories* without sacrificing robustness.

We evaluate MultiBodySync on various datasets composed of full synthetic point clouds and partial real scans with articulated and solid objects. We also contribute a new dataset DynLab with 8 scenes and 64 scan fragments of distinctly moving objects. Our extensive evaluations demonstrate that our algorithm outperforms the state-of-the-art by a large margin on both multi-body motion segmentation and motion estimation. In brief, our contributions are:

1. We introduce a novel end-to-end trainable architecture for solving the multi-scan multi-body motion estimation and segmentation problem.
2. We theoretically analyze the spectral characteristics of the proposed weighted permutation synchronization.
3. To the best of our knowledge, we showcase the first cross-category generalization for the task at hand on both synthetic and real datasets, for both articulated part-level and object-level regimes.

2. Related Works

Dynamic scene understanding. The modeling of 3D dynamic scenes in deep learning literature is often formulated as a 4D data analysis, as done in seminal works like [43, 19]. Ability to infer spatiotemporal geometric properties has recently motivated research in *3D scene flow* as a form of low-level dynamic scene representation [42, 60, 69, 51, 47, 54, 44]. Domain-specific knowledge can be employed to give better predictions as done in autonomous driving [29, 5, 72] or articulated object analysis [77, 68]. The most recent dynamic SLAM works [31, 7, 81, 75] also rely heavily on semantic cues. While some works [47, 54] advocates continuous temporal-dynamics modeling, we instead assume discrete non-sequential input and enforce consistency using synchronization. Similarly, [26, 65] propose to perform instance-level re-localization in a changed scene. Nevertheless, we do not assume a pre-segmentation of the scene, but instead perform joint motion segmentation.

Multi-body motion. Provided point correspondences between two point clouds/images, rigid-body motion segmentation becomes a multi-model fitting problem, amenable for factorization techniques [20, 40, 76], clustering [32], graph optimization [45, 35, 11] or deep learning [38]. Among others, [78] handles raw scans and segments the rigidly moving parts using a Recurrent Neural Network (RNN). [28] fits non-parametric part models to sequential 3D data without needing explicit correspondences. However, to our best

knowledge, no prior work can handle multiple scans while enforcing multi-way consistency like we do.

Synchronization. The art of consistently recovering absolute quantities from a collection of ratios is now a basic component of the classical multi-view/shape analysis pipelines [56, 14, 15]. Various aspects of the problem have been vastly studied: different group structures [25, 24, 12, 2, 1, 33, 27, 1, 66, 18, 59, 62, 4, 6], closed-form solutions [4, 2, 1], robustness [17], certifiability [55], global optimality [13], learning-to-synchronize [34, 50, 22] and uncertainty quantification [61, 10, 9, 12]. In this work, we are concerned with synchronizing correspondence sets, otherwise known as *permutation synchronization* (PS) [48] and motion segmentations [3]. PS is rich in the variety of algorithms: low-rank formulations [80, 67], convex programming [30], distributed optimization [30], multi-graph matching [57] or Riemannian optimization [12]. Out of all those, we are interested in the spectral methods of [2, 46] as they provide efficient, closed-form solutions deployable within a *deep declarative network* [23] like ours.

To the best of our knowledge, synchronization of correspondences [46] or motion segmentation [3] have not been explored in the context of deep learning. This is what we do in this paper to tackle the consistent multi-body motion estimation and segmentation.

3. Method

Problem setting and notation. Suppose we observe a set of K point clouds $\mathcal{X} = \{\mathbf{X}^k \in \mathbb{R}^{3 \times N}, k \in [1, K]\}$ where each point cloud $\mathbf{X}^k = [\mathbf{x}_1^k, \dots, \mathbf{x}_i^k, \dots, \mathbf{x}_N^k]$ contains N points in \mathbb{R}^3 and sampled from the same object with S *independently moving* rigid parts indexed by s . Each point is assumed to belong to one of the S rigid parts and we denote the binary *point-part association matrices* as $\mathcal{G} = \{\mathbf{G}^k \in [0, 1]^{N \times S}\}$ where $G_{is}^k = 1$ if \mathbf{x}_i^k belongs to the s^{th} rigid part and $G_{is}^k = 0$ otherwise¹. The rigid motions for each part s in each point cloud k is defined as $\mathcal{T} = \{\mathbf{T}_s^k \in \text{SE}(3), k \in [1, K], s \in [1, S]\}$, with the rotational part being $\mathbf{R}_s^k \in \text{SO}(3)$ and the translational part being $\mathbf{t}_s^k \in \mathbb{R}^3$. Our final goal is to infer \mathcal{G} and \mathcal{T} given \mathcal{X} .

Summary. The core of our approach is a fully differentiable deep network fusing rigid dynamic information from multiple 3D scans as outlined in Fig. 2. We begin by explicitly predicting pairwise soft correspondences across all pairs of point clouds while enforcing consistency via a *weighted permutation synchronization* (§ 3.1). Next, the point clouds are segmented using a novel motion-based segmentation network and also further synchronized by a subsequent *motion segmentation synchronization* module (§ 3.2). Finally, the correspondences and segmentations are

¹Throughout our paper we use superscript k, l to index point-clouds, subscript i, j to index points and subscript s to index rigid parts.

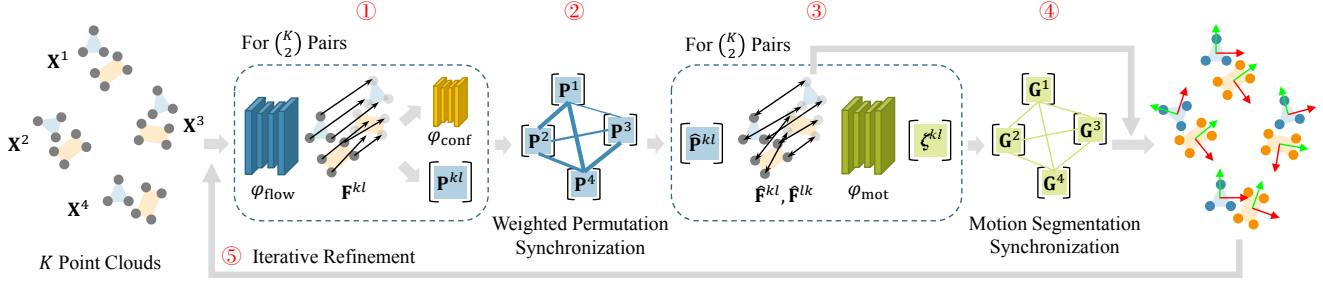


Figure 2. Our pipeline. ① Given multiple input point clouds, we first estimate pairwise scene flows. ② The point correspondences (permutations) computed from the flows are then synchronized in a weighted fashion to enforce consistencies. ③ Pairwise relative segmentations are subsequently estimated from the flows, and ④ further synchronized to get absolute motion segmentation. The pose of each part can be recovered by a weighted Kabsch algorithm. Our pipeline is fully differentiable and can be iterated ⑤ to get improved results.

used to recover the 6-DoF transformation for each of the individual rigid parts. The whole procedure can be iterated to refine the results. The pipeline can be readily trained end-to-end and we describe our training procedure in § 3.3.

3.1. Flow Estimation and Synchronization

Our approach starts with point correspondence estimation between all $\binom{K}{2}$ pairs of point clouds. We tackle this problem by predicting a 3D scene flow $\mathbf{F}^{kl} = [\mathbf{f}_1^{kl}, \dots, \mathbf{f}_i^{kl}, \dots, \mathbf{f}_N^{kl}] \in \mathbb{R}^{3 \times N}$ for each point cloud pair indexed by (k, l) using a deep neural network $\varphi_{\text{flow}}(\cdot)$, i.e. $\mathbf{F}^{kl} = \varphi_{\text{flow}}(\mathbf{X}^k, \mathbf{X}^l)$, so that $\mathbf{X}^k + \mathbf{F}^{kl} \stackrel{P}{=} \mathbf{X}^l$ holds up to a permutation. The architecture of φ_{flow} inspired by *Point PWC-Net* [73] is detailed in the supplementary material.

Flow signals, estimated in a pairwise fashion, are not informed about the multiview configuration at our disposal. To ensure multi-way consistent flows, we employ the weighted variant of permutation synchronization [46] inspired by [22, 34] where a closed-form solution is given under spectral relaxation. We begin by the observation that any flow \mathbf{F}^{kl} would induce a soft assignment matrix $\mathbf{P}^{kl} \in \mathcal{M}^{N \times N}$ based on the nearest-neighbor distances:

$$P_{ij}^{kl} = \frac{\exp(\delta_{ij}^{kl}/\tau)}{\sum_{j=1}^N \exp(\delta_{ij}^{kl}/\tau)}, \quad \delta_{ij}^{kl} = \|\mathbf{x}_i^k + \mathbf{f}_i^{kl} - \mathbf{x}_j^l\|^2 \quad (1)$$

where τ is the temperature of the *softmax*. The *multinomial manifold* \mathcal{M} of row-stochastic matrices is a *continuous relaxation* of the (partial) permutation group \mathcal{P} .

Outlier filtering. To take into account the noise, missing points, or errors in the network, we further associate a confidence value $c_i^{kl} \in \mathbb{R}$ to each point \mathbf{x}_i^k and its corresponding flow vector \mathbf{f}_i^{kl} through another network $\varphi_{\text{conf}}(\cdot) : \mathbb{R}^{7 \times N} \mapsto \mathbb{R}^N$ inspired from OANet [82]. The input to this network are the tuples $\{(\mathbf{x}_i^k, \mathbf{x}_i^k + \mathbf{f}_i^{kl}, \arg\min_j \delta_{ij}^{kl})\} \in \mathbb{R}^7$ and we provide the architectural details in the supplementary. The last dimension of this tuple measures the quality of the flow vector via the distance between the transformed points and their nearest neighbors, thereby detect-

ing spurious flow predictions. The final w^{kl} in Eq (3) reflects the overall quality of the corresponding \mathbf{P}^{kl} . Here we choose w^{kl} as the average confidence of all points, i.e., $w^{kl} = \sum_{i=1}^N c_i^{kl}/N$.

Consistent correspondences. We now use the predictions $\{\mathbf{P}^{kl}, w^{kl}\}_{(k,l)}$ to achieve multiview consistent assignments. To this end, we deploy a *differentiable synchronization* algorithm inspired by [46]. We first introduce *absolute* permutation matrices \mathbf{P}^k which map each point in \mathbf{X}^k to a *universe* space and stack them as $\mathbf{p} = [\dots, (\mathbf{P}^k)^\top, \dots]^\top$. We solve for the best \mathbf{p} minimizing:

$$E(\mathbf{p}) = \sum_{k=1}^K \sum_{l=1}^K w^{kl} \|\mathbf{P}^k - \mathbf{P}^{kl} \mathbf{P}^l\|_F^2. \quad (2)$$

Theorem 1 (Weighted synchronization). *The spectral solution to the weighted synchronization problem in Eq (2) \mathbf{p} is given by the N eigenvectors of \mathbf{L} corresponding to the smallest N eigenvalues, where $\mathbf{L} \in \mathbb{R}^{KN \times KN}$ is the **weighted Graph Connection Laplacian (GCL)** constructed by tiling all \mathbf{P}^{kl} matrices weighted by the related w^{kl} :*

$$\mathbf{L} = \begin{bmatrix} w^1 \mathbf{I}_N & -w^{12} \mathbf{P}^{12} & \dots & -w^{1K} \mathbf{P}^{1K} \\ -w^{21} \mathbf{P}^{21} & w^2 \mathbf{I}_N & \dots & -w^{2K} \mathbf{P}^{2K} \\ \vdots & \vdots & \ddots & \vdots \\ -w^{K1} \mathbf{P}^{K1} & -w^{K2} \mathbf{P}^{K2} & \dots & w^K \mathbf{I}_N \end{bmatrix}, \quad (3)$$

with $w^k = \sum_{l=1, l \neq k}^K w^{kl}$ and $\mathbf{I}_N \in \mathbb{R}^{N \times N}$ is the identity.

Proof. Please refer to the supplementary material. \square

This *spectral solution* requires only an eigen-decomposition lending itself to easy differentiation [34, 22]. The synchronized soft correspondence $\hat{\mathbf{P}}^{kl}$ is then extracted as the (k, l) -th $N \times N$ block of $\hat{\mathbf{p}} \hat{\mathbf{p}}^\top$. As a consequence of the relaxation, we cannot ensure that each sub-matrix of $\hat{\mathbf{p}} \hat{\mathbf{p}}^\top$ would be a valid permutation. To preserve differentiability we avoid Hungarian-like projection operators [46] and propose to directly compute the induced flow




Figure 3. Illustration of our motion segmentation network φ_{mot} .

$\hat{\mathbf{F}}^{kl} = [\dots, \hat{f}_i^{kl}, \dots]$ using a softmax normalization on the synchronized soft correspondences:

$$\hat{f}_i^{kl} = \frac{\sum_{j=1}^N \bar{P}_{ij}^{kl} (\mathbf{x}_j^l - \mathbf{x}_i^k)}{\sum_{j=1}^N \bar{P}_{ij}^{kl}}, \quad \bar{P}_{ij}^{kl} = \exp(\hat{P}_{ij}^{kl}/t). \quad (4)$$

Intuitively, this amounts to using the normalized synchronized result as a soft-assignment matrix, diminishing the effect of non-corresponding matches (false positives).

3.2. Motion Segmentation

Based upon the multiview consistent flow output $\hat{\mathbf{F}}^{kl}$, we now predict the point-part associations \mathcal{G} . Since we are not provided with consistent labeling of the parts, instead of predicting \mathbf{G}^k directly, we estimate for all $\binom{K}{2}$ point cloud pairs a relative *motion segmentation matrix* $\zeta^{kl} \in [0, 1]^{N \times N}$, where ζ_{ij}^{kl} is 1 when \mathbf{x}_i^k and \mathbf{x}_j^l belong to the same rigid body, and 0 otherwise.

Our motion segmentation network $\varphi_{\text{mot}}(\cdot) : \mathbb{R}^{12 \times N} \mapsto \mathbb{R}^{N \times N}$ illustrated in Fig. 3 takes the point cloud pair $\mathbf{X}^k, \mathbf{X}^l$ as well as flow $\hat{\mathbf{F}}^{kl}, \hat{\mathbf{F}}^{lk}$ estimated from the last step as input and outputs the matrix ζ^{kl} . It begins with a PointNet++ [53] predicting a transformation $\tilde{\mathbf{T}}_i^{kl}$ for each point in $\mathbf{x}_i^k \in \mathbf{X}^{k2}$. The predictions map the part in \mathbf{X}^k containing \mathbf{x}_i^k to \mathbf{X}^l . We then compute a residual matrix $\beta^{kl} \in \mathbb{R}^{3 \times N \times N}$ based on $\tilde{\mathbf{T}}_i^{kl}$, whose element is:

$$\beta_{ij}^{kl} = (\tilde{\mathbf{T}}_i^{kl})^{-1} \circ \mathbf{x}_j^l - (\mathbf{x}_j^l + \hat{f}_j^{lk}), \quad (5)$$

where \circ denotes the action of $\tilde{\mathbf{T}}$. One can easily verify that the smaller the norm of the (i, j) -th entry of β^{kl} is, the more likely that \mathbf{x}_i^k and \mathbf{x}_j^l are in the same rigid part. Therefore, it contains valuable information for deducing the motion segmentation ζ^{kl} . We apply N denoising mini-PointNet [52] $\varphi_{\text{mlp}}(\cdot)$ to each horizontal $3 \times N$ slice of β^{kl} , concatenated with \mathbf{X}^l to get a likelihood score for each pair of points $(\mathbf{x}_i^k \in \mathbf{X}^k, \mathbf{x}_j^l \in \mathbf{X}^l)$. The network output ζ_{net}^{kl} is subsequently computed by applying a sigmoid on the output:

$$(\zeta_{\text{net}}^{kl})_{i,:} = \text{sigmoid}(\varphi_{\text{mlp}}([\beta_{i,:}^{kl}, \mathbf{X}^l])). \quad (6)$$

²In practice, instead of predicting $\tilde{\mathbf{T}}_i^{kl}$ directly, we estimate a residual motion w.r.t. the already obtained flow vectors similar to the method in [78]. This procedure is detailed in our supplementary material.

Motion segmentation consistency. Given all pairwise motion information ζ^{kl} , we adopt the method of Arrigoni and Pajdla [3] to compute an absolute motion segmentation $\mathbf{g} \in \mathbb{R}^{KN \times S}$ as a stack of matrices in \mathcal{G} . Once again, this is an instance of a synchronization problem, with the stacked relative and absolute motion segmentation matrices being:

$$\mathbf{Z} = \begin{bmatrix} \mathbf{0} & \zeta^{12} & \dots & \zeta^{1K} \\ \zeta^{21} & \mathbf{0} & \dots & \zeta^{2K} \\ \vdots & \vdots & \ddots & \vdots \\ \zeta^{K1} & \zeta^{K2} & \dots & \mathbf{0} \end{bmatrix}, \quad \mathbf{g} = \begin{bmatrix} \mathbf{G}^1 \\ \mathbf{G}^2 \\ \vdots \\ \mathbf{G}^K \end{bmatrix}, \quad (7)$$

A spectral approach similar to the one in § 3.1 optimizes for \mathbf{g} so that $\mathbf{Z} = \mathbf{g}\mathbf{g}^\top$ is best satisfied. Then, \mathbf{g} is just the S leading eigenvectors of \mathbf{Z} , scaled by the square root of its S largest eigenvalues. Here, the point-part association matrices \mathbf{G}^k are relaxed to *fuzzy* segmentations by allowing its entries to take real values. As a subsequent step similar to § 3.1, we replace the projection step with a row-wise softmax on \mathbf{g} to maintain differentiability.

Note that the output ζ_{net}^{kl} of φ_{mot} is *unnormalized*, meaning that any submatrix in \mathbf{Z} can be written as $\zeta^{kl} = \sigma^{kl} \zeta_{\text{net}}^{kl}$, where σ acts as a normalizer. This is akin to encoding a *confidence* in the norm of the matrix ζ_{net}^{kl} and requires us to solve a weighted synchronization. However, as we prove in the following theorem, such a solution would involve an anisotropic scaling in the eigenvectors as a function of the number of points belonging to each part. As this piece of information is not available in runtime, we take an alternative approach and approximate the scaling factor as $q^{kl} = \text{mean}(\zeta_{\text{net}}^{kl})$ and *pre-factor* it out of ζ_{net}^{kl} , by letting $\zeta^{kl} = \zeta_{\text{net}}^{kl} / q^{kl}$. In this way, we ensure that the eigenvectors yield the synchronized motion segmentation.

Theorem 2. *Under mild assumptions, the solution to the segmentation synchronization problem using a non-uniformly weighted matrix will result in a proportionally scaled version of the solution obtained by the eigenvectors of the unweighted matrix \mathbf{Z} .*

Proof. Please refer to the supplementary material. \square

As we show in our supplement, entry k in the decomposed eigenvalues is related to the number of points belonging to motion k . To compute the number of rigid bodies S , i.e., determine how many eigenvectors to use in \mathbf{g} , the spectrum of \mathbf{Z} is analyzed during test time: We estimate S as the number of eigenvalues that are larger than α -percent of the sum of the first 10 eigenvalues. For training, we just fix $S = 6$ as an over-parametrization.

Pose Computation and Iterative Refinement. We finally estimate the motion for each part using a weighted Kabsch algorithm [37, 22] followed by a joint pose estimation. During test time we also iterate our pipeline several times to

gradually refine the correspondence and segmentation estimation by transforming input point clouds according to the estimated \mathcal{T} and adding back the residual flow onto the flow predicted at the previous iteration. The details are provided in our supplementary.

3.3. Network Training

We propose to train each learnable component of our pipeline separately in a pairwise manner and then fine-tune their parameters using the full pipeline. Specifically, we first train the flow estimation network φ_{flow} supervised with ground-truth flow: $\mathcal{L}_{\text{flow}}^{kl} = \|\mathbf{F}^{kl} - \mathbf{F}^{kl, \text{gt}}\|_F^2$. Given the trained φ_{flow} , the confidence estimation network φ_{conf} is trained based on its output using a binary cross-entropy (BCE) loss supervised by comparing whether the error of the predicted flow is under a certain threshold:

$$\mathcal{L}_{\text{conf}}^{kl} = \sum_{i=1}^N \text{BCE}(c_i^{kl, \text{gt}}, c_i^{kl}), \quad (8)$$

with $c_i^{kl, \text{gt}} = 1$ if we have $\|\mathbf{f}_i^{kl} - \mathbf{f}_i^{kl, \text{gt}}\|_2 < \epsilon_f$ and 0 otherwise. The motion segmentation network φ_{mot} is trained using joint supervision over the estimated transformation residual and the final motion segmentation matrix: $\mathcal{L}_{\text{seg}}^{kl} = \mathcal{L}_{\text{trans}}^{kl} + \mathcal{L}_{\text{group}}^{kl}$ where each term is defined as:

$$\mathcal{L}_{\text{trans}}^{kl} = \sum_{i=1}^N \frac{\sum_{j=1}^N \mathbb{I}(c_{ij}^{kl} = 1) \|\beta_{ij}^{kl}\|_2^2}{\sum_{j=1}^N \mathbb{I}(c_{ij}^{kl} = 1)}, \quad (9)$$

$$\mathcal{L}_{\text{group}}^{kl} = \sum_{i=1}^N \sum_{j=1}^N \text{BCE}(\zeta_{ij}^{kl, \text{gt}}, \zeta_{ij}^{kl}). \quad (10)$$

After we train all the networks (*i.e.*, φ_{flow} , φ_{conf} and φ_{mot}), the entire pipeline is trained end-to-end with the supervision on both the pairwise flow $\sum_{k=1}^K \sum_{l=1}^K \mathcal{L}_{\text{flow}}^{kl}$ and the IoU (Intersection-over-union) loss, defined as:

$$\mathcal{L}_{\text{iou}} = \underset{\mathbf{A}}{\text{argmin}} \sum_{s, s'=1}^{S \times S} \frac{\mathbf{A}(s, s') \cdot (\mathbf{g}_{:,s}^{\text{gt}})^\top \mathbf{g}_{:,s'}}{\|\mathbf{g}_{:,s}^{\text{gt}}\|_2^2 + \|\mathbf{g}_{:,s'}\|_2^2 - (\mathbf{g}_{:,s}^{\text{gt}})^\top \mathbf{g}_{:,s'}},$$

where \mathbf{A} is an $S \times S$ binary assignment matrix which we found using the Hungarian algorithm. The flow supervision is added to both the output of flow network, and the final pairwise *rigid flow* computed as $\mathbf{f}_i^{kl} = \mathbf{T}_s^l (\mathbf{T}_s^k)^{-1} \circ \mathbf{x}_i - \mathbf{x}_i$.

4. Experiments

Datasets. Our algorithm is tested on two main datasets: SAPIEN [74] and *DynLab* dataset contributed by this work: SAPIEN consists of realistic simulated articulated models with part mobility annotated. We ensure that the categories used for training and validation do not overlap with the test set, finally leading to 720 articulated objects with 20 dif-

Table 1. Rigid flow estimation on SAPIEN. Table reports mean and std. dev. (+/-) of the EPE3D over all pairwise flows, with (S) or w/o (NS) Synchronization and with (W) or w/o (NW) Weighting.

	Deep Part [78]	NPP [28]	Ours		
			NS, NW	S, NW	S, W
Mean	5.95	21.22	6.20	6.08	5.03
+/-	3.57	6.29	4.06	3.47	2.00

ferent categories. We then perform K virtual 3D scans of the models, with each scan capturing the same object with a different camera (and hence object) pose and object articulating state. Later, furthest point sampling is applied to down-sample the number of points to N . **DynLab** (Dynamic Laboratory) contains 8 different scenes in a laboratory, each with 2-3 rigidly moving *solid* objects from various categories. Each of the scenes is captured 8 times, reconstructed using ElasticFusion [71] and between each capture, the object positions are randomly changed. The dataset also contains manual annotations of the object segmentation mask and rigid absolute transformations. For benchmarking, in each scene we choose different combinations of the 8 captures, leading to a total of $8 \cdot \binom{8}{4} = 560$ dataset items. We believe the two different scenarios (articulated single object and moving rigid bodies) reflected in the test sets are sufficient to verify the robustness and the general applicability of our algorithm.

The training data for articulated objects are generated using the dataset from [79], containing manually annotated semantic segmentation of 16 categories. Similar to [78], we generate K random motions for each connected semantic part of the shapes. For the training data of solid objects, we randomly sample independent motions for multiple objects taken from ShapeNet [16] as if they are floating and rotating in the air. Please refer to supplementary material for detailed data specifications and visualizations.

Metrics. Two main metrics are used: (1) **EPE3D** (End-Point Error in 3D) of all $\binom{K}{2}$ pairs of point clouds. The mean and standard deviation (+/-) measures the rigid 3D flow estimation quality: While the mean reflects an overall error in the transformation, the standard deviation shows how consistent the estimate is among all pairs - a desirable property in the multi-scan setting. (2) Segmentation accuracy assesses the motion segmentation quality. We use **mIoU** (mean Intersection-over-Union) and **RI** (Rand Index) to score the output based on ‘Multi-Scan’ and ‘Per-Scan’ segmentations. For ‘**Multi-Scan**’, we evaluate the points from all K clouds altogether, revealing the consistency of the labeling across multiple scans. For ‘**Per-Scan**’, we compute the score for each of the clouds separately and evaluate the mean and standard deviation across all scans.

Training. φ_{flow} , φ_{mot} and φ_{conf} are trained using Adam optimizer with initial learning rate of 10^{-3} and a 0.5/0.7/0.7 decay every 400K iterations for the three networks. The

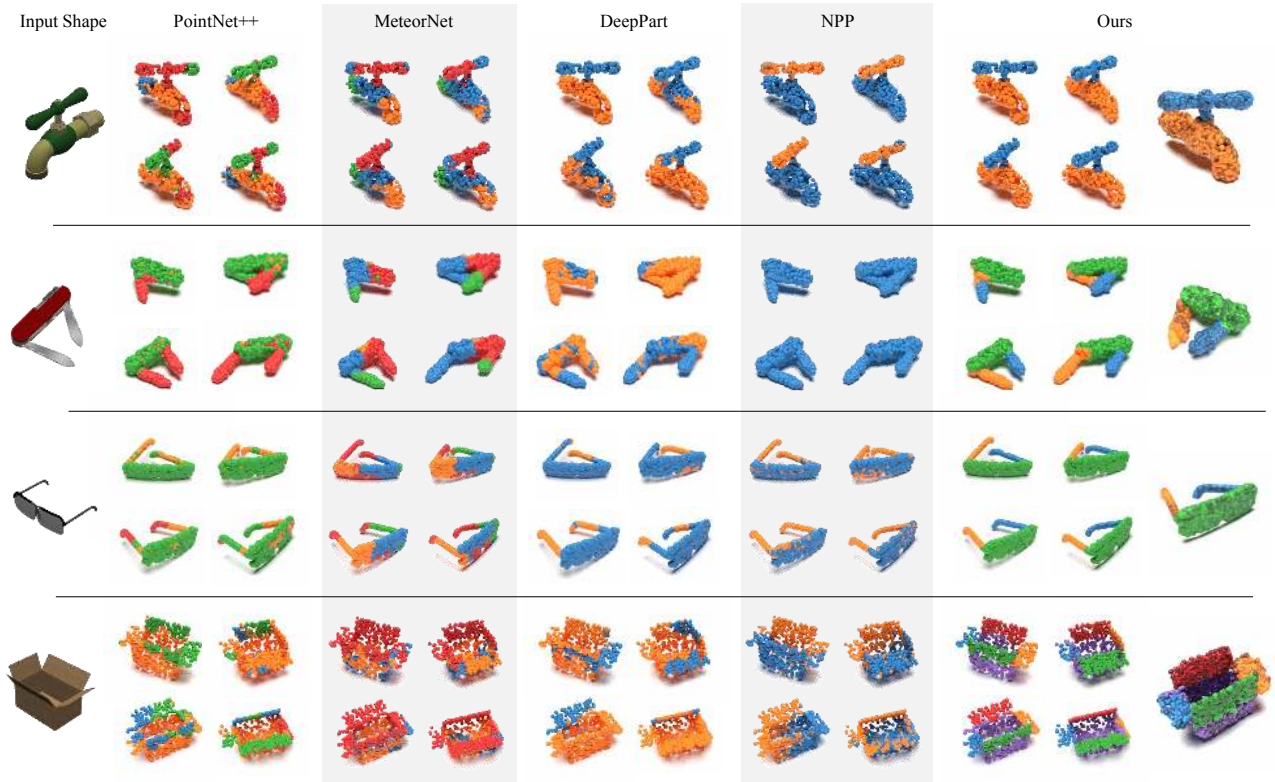


Figure 4. Qualitative results on SAPIEN [74] dataset. On the left-most column, we show reference rendering of the objects we tackle. Each method span two columns and the point colors show the motion segmentation. For our method, we additionally show the registration result as the last column, where the darkness of the points shows the point cloud index it comes from.

Table 2. Segmentation Accuracy on SAPIEN Dataset.

	Multi-Scan		Per-Scan	
	mIoU	RI	mIoU	RI
PointNet++ [53]	47.5	0.62	51.2±12.1	0.65±0.09
MeteorNet [43]	43.7	0.59	45.7±5.4	0.60±0.03
DeepPart [78]	49.2	0.64	53.0±8.9	0.67±0.06
NPP [28]	48.2	0.63	51.5±6.6	0.66±0.05
Ours (4 iters)	66.7	0.76	67.3±4.3	0.77±0.03

batch sizes are set to 32/8/32, respectively. The entire pipeline is trained end-to-end using $K = 4$ point clouds, with a learning rate of 10^{-6} . The gradient computation for eigen-decomposition will sometimes lead to numerical instabilities [21], so we roll back that iteration when the gradient norm is large. Our algorithm is implemented using PyTorch [49] with $N = 512$, $\tau = 0.01$, $\epsilon_f = 0.1$. We set $\alpha = 0.05$ for articulated objects and $\alpha = 0.15$ for solid objects.

4.1. Results on Articulated Objects

Baselines. Given our new multi-scan multi-body setting, we made adaptations to previous methods and compared to the following 4 baselines: (1) **PointNet++** [53]: We use the segmentation backbone to directly predict \mathbf{G}^k matri-

ces. We aggregate the bottleneck features by taking the max before feeding it to the individual K feature propagation modules. (2) **MeteorNet** [43]: We use the *MeteorNet-seg* model proposed to directly predict the segmentations. Both PointNet++ and MeteorNet are supervised with the IoU loss (Eq (11)) which counts in the ambiguity of rigid body labeling. (3) **DeepPart** [78]: As this method only allows pairwise input, we associate multiple point clouds using sequential label propagation. (4) **NPP** (Non-Parametric Part) [28]: This algorithm does not need training and a grid search is conveyed for its many tunable parameters.

Flow Accuracy. Tab. 1 shows that despite being based on [78], our method gives the lowest flow error and variance across different view pairs. This is thanks to the correspondence consistency among the provided K scans enforced by our synchronization module. The NPP method suffers from a surprisingly high flow error mainly because the point-level correspondence is not explicitly modeled. Note that PointNet++ and MeteorNet are excluded because they only output point-wise segmentations.

Segmentation Accuracy. For the segmentation benchmark, we achieve a significantly better result than all the baselines as shown in Tab. 2. Among the baselines, Me-




Figure 5. Qualitative segmentation results on two articulated sequences from [63] (‘Donkey’ and ‘Pipe 3/4’, first two rows) and [58] (‘Alex’ and ‘Duck’, last two rows).

torNet fails because it assumes proximity of relevant data in the given point clouds, which is not robust to SAPIEN dataset because of change in both object pose and articulated parts. Even though PointNet++ reaches a relatively high mean score, the standard deviation and Multi-Scan score show the segmentation is not consistent across different input scans. DeepPart is specially designed for part-based motion segmentation, but only operates on two-views, which can cause drastic performance degradation if the input two-views have a large difference. Also the RNN they proposed for part segmentation tends to generate short sequences and most of the shapes are only divided into two parts. Despite the large error in flow estimation, NPP behaves reasonably in terms of segmentation. Qualitative comparisons are visualized in Fig. 4.

One important aspect of our network is that it can generalize to different objects and motions *without re-training*. To qualitatively showcase this, we use two additional dynamic RGB-D sequences from [63] and [58]. For each sequence, we use four views and back-project the depth map into point clouds for inference. As shown in Fig. 5, our model trained on full objects of synthetic SAPIEN dataset, can generalize to real dynamic depth sequences producing consistent motion-based segmentation. This is possible thanks to the property that our network anchors on the motion and not on the specific geometry.

4.2. Results on Full Objects

In DynLab, each rigid body (*i.e.* object) is now semantically meaningful, so apart from the 4 baseline methods from § 4.1, we additionally compare to the following two alternatives: (5) **InstSeg** (Instance Segmentation): We take the state-of-the-art indoor semantic instance module PointGroup [36] trained on ScanNet dataset to segment for each input cloud. (6) **Geometric**: We use the Ward-linkage [70] to agglomeratively cluster the points in each scan. In order to obtain consistent segmentation across multiple inputs, we

Table 3. Segmentation Accuracy on DynLab dataset.

	Multi-Scan		Per-Scan	
	mIoU	RI	mIoU	RI
PointNet++ [53]	37.2	0.53	39.4±7.1	0.54±0.03
MeteorNet [43]	58.5	0.69	71.8±9.7	0.76±0.06
DeepPart [78]	60.7	0.70	66.3±17.2	0.75±0.13
NPP [28]	65.7	0.74	71.6±7.7	0.78±0.05
Geometric	83.1	0.87	88.6±5.8	0.91±0.04
InstSeg [36]	56.5	0.66	72.4±12.5	0.78±0.09
Ours	90.7	0.95	94.0±3.1	0.96±0.02




Figure 6. Example comparisons to baselines on DynLab. In the leftmost column we compare the warped point clouds without (‘NS’) and with (‘S’) synchronization. The right three sub-figures show the segmentation in different colors. For clarity we exclude the computed pose for the *geometric* & *InstSeg* approaches because the inaccuracy in segmentation leads to very noisy poses.

associate the segmentations between two different scans using a Hungarian search over the object assignment matrix, whose element is the *root mean squared error* measuring the fitting quality between any combinations of the object associations.

Interestingly, as listed in Tab. 3, all the previous deep methods lead to unsatisfactory results on this dataset. PointNet++ and MeteorNet are found to be inaccurate because by design they associate labels in the level of semantics (not motion) and no explicit consistencies across scans are considered. Even though the InstSeg method is trained on large-scale scene dataset, it is impossible for it to cover all real-world categories so wrong detections are observed in some scenes. The geometric approach is less robust in cluttered scenes where no obvious geometric cues can be used. Our method is motion-induced and is hence robust to geometric variations and out-of-distribution semantics, outperforming all baselines. A typical failure scenario for these approaches is visualized in Fig. 6. We show additional qualitative results in Fig. 7, demonstrating our ability to accurately segment, associate, and compute correct object transformations even if there are large pose changes.

Tab. 4 shows the rigid flow estimation result against the baselines. Apart from the influence of wrong per-scan segmentation and cross-scan associations, the iterative closest point (ICP) [8] method used to register object scans can also suffer from poor initializations. Our approach not only




Figure 7. Results on DynLab dataset. Note that we detect and remove the ground for all baselines except InstSeg so only the points on the moving furniture are input. Point colors indicate segmentation and the bounding boxes show relative transformations.

Table 4. Rigid flow estimation result on DynLab dataset. The numbers represent mean and standard deviation (+/-) of the EPE3D from all pairwise flows. Note the value here does not reflect real-world metrics because the scales are uniformly normalized.

	DeepPart [78]	NPP [28]	Geometric	InstSeg [36]	Ours
Mean	16.89	51.14	21.61	46.40	11.01
+/-	11.39	18.38	9.76	20.73	6.65

reaches the lowest mean error, but also respects the motion consistency across multiple scans.

4.3. Ablation Study and System Analysis

Effect of synchronization. For permutation synchronization (§ 3.1), we can directly feed the network-predicted flow vector \mathbf{F}^{kl} to subsequent steps instead of using synchronized $\hat{\mathbf{F}}^{kl}$ (Ours: NS, NW), or use an unweighted version of the synchronization by setting all $w^{kl} = 1$ (Ours: S, NW). However, as shown quantitatively in Tab. 1, both variants result in higher flow error due to the failure to find consistent correspondences. Similar results can be observed on DynLab dataset as demonstrated in the two sub-figures of Fig. 6, where direct flow prediction failed because the geometric variation is too large between two scans.

Effect of K . Our method can be naturally applied to an arbitrary number of views K even if we train using 4 views, because by design the learnable parameters are unaware of the input counts. As shown in Fig. 9, the segmentation accuracy improves given more views. This is because the introduction of additional scans helps build the connection between existing scans and benefits the ‘co-segmentation’ process.




Figure 8. Iterative refinement on SAPIEN dataset. We show transformed and segmented point clouds according to recovered motions over iterations (shown in the middle).




Figure 9. Influence of number of iterations (**left**) and number of views (**right**) on the final segmentation accuracy.

Number of iterations. As pointed out in § 3.2, our pipeline can be run multiple iterations to refine the results and an example is given in Fig. 8. Shown in Fig. 9, our method works better with more iterations because we estimate more accurate flows. Moreover, more iterations are demonstrated to be unnecessary because previous iterations already lead to converged results.

Timing. Our experiments are conducted using an Nvidia GeForce GTX 1080 card. For the input of 4 scans, the running time of our full model is $\sim 870\text{ms}$ per iteration. The entirety of a 4-iteration scheme hence takes $\sim 3.5\text{s}$, while [78] and [28] take 11.5s and 60s resp. in comparison.

5. Conclusion

We presented *MultiBodySync*, a pipeline for simultaneously segmenting and registering multiple dynamic scans with multiple rigid bodies. We, for the first time, incorporated weighted permutation synchronization and motion segmentation synchronization into a fully-differentiable pipeline for generating consistent results across all input point clouds. However, currently *MultiBodySync* is not scalable to a large number (like hundreds) of scans or rigid bodies. Future directions include improvement of the pipeline’s scalability and robustness in more complicated and dynamic settings.

Acknowledgements. We ack support from the China Scholarship Council, the Natural Science Foundation of China (No. 61521002), the Joint NSFC-DFG Research Program (No. 61761136018), a grant from Tsinghua-Tencent Joint Laboratory for Internet Innovation Technology, NSF grant CHS-1528025, a Vannevar Bush Faculty fellowship, a TUM/IAS Hans Fischer senior fellowship, and gifts from the Adobe, Amazon AWS, and Snap corporations. Arrigoni was supported by the EU Horizon 2020 Research and Innovation Programme under project SPRING (No. 871245).

References

- [1] Federica Arrigoni and Andrea Fusiello. Synchronization problems in computer vision with closed-form solutions. In *International Journal of Computer Vision*, Sep 2019. 2
- [2] Federica Arrigoni, Eleonora Maset, and Andrea Fusiello. Synchronization in the symmetric inverse semigroup. In *International Conference on Image Analysis and Processing*, pages 70–81. Springer, 2017. 2
- [3] Federica Arrigoni and Tomas Pajdla. Motion segmentation via synchronization. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2019. 2, 4
- [4] Federica Arrigoni, Beatrice Rossi, and Andrea Fusiello. Spectral synchronization of multiple views in se (3). *SIAM Journal on Imaging Sciences*, 9(4):1963–1990, 2016. 2
- [5] Aseem Behl, Despoina Paschalidou, Simon Donné, and Andreas Geiger. Pointflownet: Learning representations for rigid motion estimation from point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7962–7971, 2019. 2
- [6] Florian Bernard, Johan Thunberg, Peter Gemmar, Frank Hertel, Andreas Husch, and Jorge Goncalves. A solution for multi-alignment by transformation synchronisation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2
- [7] Berta Bescos, Carlos Campos, Juan D Tardós, and José Neira. Dynaslam ii: Tightly-coupled multi-object tracking and slam. *arXiv preprint arXiv:2010.07820*, 2020. 2
- [8] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor fusion IV: control paradigms and data structures*, volume 1611, pages 586–606. International Society for Optics and Photonics, 1992. 7
- [9] Tolga Birdal, Michael Arbel, Umut Simsekli, and Leonidas J Guibas. Synchronizing probability measures on rotations via optimal transport. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [10] Tolga Birdal, Umut Şimşekli, M. Onur Eken, and Slobodan Ilic. Bayesian Pose Graph Optimization via Bingham Distributions and Tempered Geodesic MCMC. In *Advances in Neural Information Processing Systems*, 2018. 2
- [11] Tolga Birdal and Slobodan Ilic. Cad priors for accurate and flexible instance reconstruction. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017. 2
- [12] Tolga Birdal and Umut Simsekli. Probabilistic permutation synchronization using the riemannian structure of the birkhoff polytope. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [13] Jesus Briales and Javier Gonzalez-Jimenez. Cartan-sync: Fast and global se (d)-synchronization. *IEEE Robotics and Automation Letters*, 2(4):2127–2134, 2017. 2
- [14] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on robotics*, 32(6):1309–1332, 2016. 2
- [15] Luca Carlone, Roberto Tron, Kostas Daniilidis, and Frank Dellaert. Initialization techniques for 3d slam: a survey on rotation estimation and its use in pose graph optimization. In *International Conference on Robotics and Automation*, pages 4597–4604. IEEE, 2015. 2
- [16] 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. 5
- [17] Avishek Chatterjee and Venu Madhav Govindu. Robust relative rotation averaging. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):958–972, 2017. 2
- [18] Kunal N Chaudhury, Yuehaw Khoo, and Amit Singer. Global registration of multiple point clouds using semidefinite programming. *SIAM Journal on Optimization*, 25(1), 2015. 2
- [19] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [20] João Paulo Costeira and Takeo Kanade. A multibody factorization method for independently moving objects. *International Journal of Computer Vision*, 29(3), 1998. 2
- [21] Zheng Dang, Kwang Moo Yi, Yinlin Hu, Fei Wang, Pascal Fua, and Mathieu Salzmann. Eigendecomposition-free training of deep networks with zero eigenvalue-based losses. In *European Conference on Computer Vision*, pages 768–783, 2018. 6
- [22] Zan Gojcic, Caifa Zhou, Jan D Wegner, Leonidas J Guibas, and Tolga Birdal. Learning multiview 3d point cloud registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2, 3, 4
- [23] Stephen Gould, Richard Hartley, and Dylan Campbell. Deep declarative networks: A new hope. Technical report, Australian National University (arXiv:1909.04866), Sep 2019. 1, 2
- [24] Venu Madhav Govindu. Lie-algebraic averaging for globally consistent motion estimation. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages I–I. IEEE, 2004. 2
- [25] Venu Madhav Govindu and A Pooja. On averaging multiview relations for 3d scan registration. *IEEE Transactions on Image Processing*, 23(3):1289–1302, 2014. 2
- [26] Maciej Halber, Yifei Shi, Kai Xu, and Thomas Funkhouser. Rescan: Inductive instance segmentation for indoor rgb-d scans. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2541–2550, 2019. 2
- [27] Richard Hartley, Jochen Trumpf, Yuchao Dai, and Hongdong Li. Rotation averaging. *International journal of computer vision*, 103(3), 2013. 2
- [28] David S Hayden, Jason Pacheco, and John W Fisher. Non-parametric object and parts modeling with lie group dynamics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2, 5, 6, 7, 8
- [29] Hou-Ning Hu, Qi-Zhi Cai, Dequan Wang, Ji Lin, Min Sun, Philipp Krahenbuhl, Trevor Darrell, and Fisher Yu. Joint monocular 3d vehicle detection and tracking. In *Proceedings of the IEEE international conference on computer vision*, pages 5390–5399, 2019. 2

- [30] Nan Hu, Qixing Huang, Boris Thibert, UG Alpes, and Leonidas Guibas. Distributable consistent multi-object matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [31] Jiahui Huang, Sheng Yang, Tai-Jiang Mu, and Shi-Min Hu. Clustervo: Clustering moving instances and estimating visual odometry for self and surroundings. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2168–2177, 2020. 2
- [32] Jiahui Huang, Sheng Yang, Zishuo Zhao, Yu-Kun Lai, and Shi-Min Hu. Clusterslam: A slam backend for simultaneous rigid body clustering and motion estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5875–5884, 2019. 2
- [33] Qixing Huang, Zhenxiao Liang, Haoyun Wang, Simiao Zuo, and Chandrajit Bajaj. Tensor maps for synchronizing heterogeneous shape collections. *ACM Trans. Graph.*, 38(4):106, 2019. 2
- [34] Xiangru Huang, Zhenxiao Liang, Xiaowei Zhou, Yao Xie, Leonidas J Guibas, and Qixing Huang. Learning transformation synchronization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8082–8091, 2019. 2, 3
- [35] Hossam Isack and Yuri Boykov. Energy-based geometric multi-model fitting. *International journal of computer vision*, 97(2):123–147, 2012. 2
- [36] Li Jiang, Hengshuang Zhao, Shaoshuai Shi, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Pointgroup: Dual-set point grouping for 3d instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4867–4876, 2020. 7, 8
- [37] Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography*, 1976. 4
- [38] Florian Kluger, Eric Brachmann, Hanno Ackermann, Carsten Rother, Michael Ying Yang, and Bodo Rosenhahn. Consac: Robust multi-model fitting by conditional sample consensus. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [39] Loïc Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4558–4567, 2018. 1
- [40] Ting Li, Vinutha Kallem, Dheeraj Singaraju, and René Vidal. Projective factorization of multiple rigid-body motions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–6. IEEE, 2007. 2
- [41] Xiaolong Li, He Wang, Li Yi, Leonidas J Guibas, A Lynn Abbott, and Shuran Song. Category-level articulated object pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 1
- [42] Xingyu Liu, Charles R Qi, and Leonidas J Guibas. Flownet3d: Learning scene flow in 3d point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 529–537, 2019. 2
- [43] Xingyu Liu, Mengyuan Yan, and Jeannette Bohg. Meteor-net: Deep learning on dynamic 3d point cloud sequences. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9246–9255, 2019. 2, 6, 7
- [44] Wei-Chiu Ma, Shenlong Wang, Rui Hu, Yuwen Xiong, and Raquel Urtasun. Deep rigid instance scene flow. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3614–3622, 2019. 2
- [45] Luca Magri and Andrea Fusiello. Fitting multiple heterogeneous models by multi-class cascaded t-linkage. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7460–7468, 2019. 2
- [46] Eleonora Maset, Federica Arrigoni, and Andrea Fusiello. Practical and efficient multi-view matching. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4568–4576, 2017. 2, 3
- [47] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 2
- [48] Deepti Pachauri, Risi Kondor, and Vikas Singh. Solving the multi-way matching problem by permutation synchronization. In *Advances in neural information processing systems*, pages 1860–1868, 2013. 2
- [49] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Álché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2019. 6
- [50] Pulak Purkait, Tat-Jun Chin, and Ian Reid. Neurora: Neural robust rotation averaging. *arXiv preprint arXiv:1912.04485*, 2019. 2
- [51] Gilles Puy, Alexandre Boulch, and Renaud Marlet. FLOT: Scene Flow on Point Clouds Guided by Optimal Transport. In *European Conference on Computer Vision*, 2020. 2
- [52] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 4
- [53] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017. 4, 6, 7
- [54] Davis Rempe, Tolga Birdal, Yongheng Zhao, Zan Gojcic, Srinath Sridhar, and Leonidas J. Guibas. Caspr: Learning canonical spatiotemporal point cloud representations. In *Advances in Neural Information Processing Systems*, 2020. 2
- [55] David M Rosen, Luca Carlone, Afonso S Bandeira, and John J Leonard. Se-sync: A certifiably correct algorithm for synchronization over the special euclidean group. *The International Journal of Robotics Research*, 38, 2019. 2
- [56] Renato F Salas-Moreno, Richard A Newcombe, Hauke Strasdat, Paul HJ Kelly, and Andrew J Davison. Slam++: Simultaneous localisation and mapping at the level of objects.

- In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1352–1359, 2013. 2
- [57] Michele Schiavinato and Andrea Torsello. Synchronization over the birkhoff polytope for multi-graph matching. In *International Workshop on Graph-Based Representations in Pattern Recognition*, pages 266–275. Springer, 2017. 2
- [58] M. Slavcheva, M. Baust, D. Cremers, and S. Ilic. Killing-Fusion: Non-rigid 3D Reconstruction without Correspondences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 7
- [59] Johan Thunberg, Florian Bernard, and Jorge Goncalves. Distributed methods for synchronization of orthogonal matrices over graphs. *Automatica*, 80:243–252, 2017. 2
- [60] Ivan Tishchenko, Sandro Lombardi, Martin R Oswald, and Marc Pollefeys. Self-supervised learning of non-rigid residual flow and ego-motion. *arXiv preprint arXiv:2009.10467*, 2020. 2
- [61] Roberto Tron and Kostas Daniilidis. Statistical pose averaging with non-isotropic and incomplete relative measurements. In *European Conference on Computer Vision*. Springer, 2014. 2
- [62] Roberto Tron and Rene Vidal. Distributed 3-d localization of camera sensor networks from 2-d image measurements. *IEEE Transactions on Automatic Control*, 59(12), 2014. 2
- [63] Dimitrios Tzionas and Juergen Gall. Reconstructing articulated rigged models from rgb-d videos. In *European Conference on Computer Vision Workshops*, 2016. 7
- [64] Sundar Vedula, Simon Baker, Peter Rander, Robert Collins, and Takeo Kanade. Three-dimensional scene flow. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2, pages 722–729. IEEE, 1999. 1
- [65] Johanna Wald, Armen Avetisyan, Nassir Navab, Federico Tombari, and Matthias Nießner. Rio: 3d object instance re-localization in changing indoor environments. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7658–7667, 2019. 1, 2
- [66] Lanhui Wang and Amit Singer. Exact and stable recovery of rotations for robust synchronization. *Information and Inference: A Journal of the IMA*, 2(2):145–193, 2013. 2
- [67] Qianqian Wang, Xiaowei Zhou, and Kostas Daniilidis. Multi-image semantic matching by mining consistent features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [68] Xiaogang Wang, Bin Zhou, Yahao Shi, Xiaowu Chen, Qingping Zhao, and Kai Xu. Shape2motion: Joint analysis of motion parts and attributes from 3d shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8876–8884, 2019. 2
- [69] Zirui Wang, Shuda Li, Henry Howard-Jenkins, Victor Prisacariu, and Min Chen. Flownet3d++: Geometric losses for deep scene flow estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 91–98, 2020. 2
- [70] Joe H Ward Jr. Hierarchical grouping to optimize an objective function. *Journal of the American statistical association*, 58(301):236–244, 1963. 7
- [71] Thomas Whelan, Stefan Leutenegger, R Salas-Moreno, Ben Glocker, and Andrew Davison. Elasticfusion: Dense slam without a pose graph. In *Robotics: Science and Systems*, 2015. 5
- [72] Pengxiang Wu, Siheng Chen, and Dimitris N Metaxas. Motionnet: Joint perception and motion prediction for autonomous driving based on bird’s eye view maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11385–11395, 2020. 1, 2
- [73] Wenxuan Wu, Zhi Yuan Wang, Zhuwen Li, Wei Liu, and Li Fuxin. Pointpwc-net: Cost volume on point clouds for (self-) supervised scene flow estimation. In *European Conference on Computer Vision*, pages 88–107. Springer, 2020. 3
- [74] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, and Hao Su. SAPIEN: A simulated part-based interactive environment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2020. 5, 6
- [75] Binbin Xu, Wenbin Li, Dimos Tzoumanikas, Michael Bloesch, Andrew Davison, and Stefan Leutenegger. Mid-fusion: Octree-based object-level multi-instance dynamic slam. In *International Conference on Robotics and Automation*, pages 5231–5237. IEEE, 2019. 2
- [76] Xun Xu, Loong Fah Cheong, and Zhuwen Li. 3d rigid motion segmentation with mixed and unknown number of models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 2
- [77] Zihao Yan, Ruizhen Hu, Xingguang Yan, Luanmin Chen, Oliver van Kaick, Hao Zhang, and Hui Huang. Rpm-net: Recurrent prediction of motion and parts from point cloud. *ACM Trans. Graph.*, 38(6):240:1–240:15, 2019. 2
- [78] Li Yi, Haibin Huang, Difan Liu, Evangelos Kalogerakis, Hao Su, and Leonidas Guibas. Deep part induction from articulated object pairs. *ACM Trans. Graph.*, 37(6), 2018. 1, 2, 4, 5, 6, 7, 8
- [79] Li Yi, Vladimir G Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A scalable active framework for region annotation in 3d shape collections. *ACM Trans. Graph.*, 35(6):1–12, 2016. 5
- [80] Jin-Gang Yu, Gui-Song Xia, Ashok Samal, and Jinwen Tian. Globally consistent correspondence of multiple feature sets using proximal gauss–seidel relaxation. *Pattern Recognition*, 51:255–267, 2016. 2
- [81] Jun Zhang, Mina Henein, Robert Mahony, and Viorela Ila. Vdo-slam: A visual dynamic object-aware slam system. *arXiv preprint arXiv:2005.11052*, 2020. 2
- [82] Jiahui Zhang, Dawei Sun, Zixin Luo, Anbang Yao, Lei Zhou, Tianwei Shen, Yurong Chen, Long Quan, and Hongen Liao. Learning two-view correspondences and geometry using order-aware network. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 3