

# Extreme Relative Pose Network under Hybrid Representations

Zhenpei Yang\*  
UT Austin

Siming Yan\*  
UT Austin

Qixing Huang†  
UT Austin

## Abstract

In this paper, we introduce a novel RGB-D based relative pose estimation approach that is suitable for small-overlapping or non-overlapping scans and can output multiple relative poses. Our method performs scene completion and matches the completed scans. However, instead of using a fixed representation for completion, the key idea is to utilize hybrid representations that combine 360-image, 2D image-based layout, and planar patches. This approach of ours adaptively feature representations for relative pose estimation. Besides, we introduce a global-2-local matching procedure, which utilizes initial relative poses obtained during the global phase to detect and then integrate geometric relations for pose refinement. Experimental results justify the potential of this approach across a wide range of benchmark datasets. For example, on ScanNet, the rotation translation errors of the top-1/top-5 predictions of our approach are  $28.6^\circ/0.90m$  and  $16.8^\circ/0.76m$ , respectively. Our approach also considerably boosts the performance of multi-scan reconstruction in few-view reconstruction settings.

## 1. Introduction

Estimating the relative pose between two RGB-D scans is a crucial problem in 3D vision and robotics. In this paper, we are interested in the setting where the overlapping region between two input scans is small, or they may not even overlap. Efficient and robust solutions to this extreme relative pose problem (c.f. [49]) enjoy a wide range of applications. Examples include 3D reconstruction from a few views [12] (e.g., RGB-D scans of an indoor scene captured at a few distinctive locations), enhancing the performance of interactive scanning [42, 35] when there are interruptions during the acquisition process, early detection of loop closure [2], and solving jigsaw puzzles [8].

A standard pipeline in relative pose estimation or object matching, in general, is to first extract features from the input objects and then match these features to derive relative poses or dense correspondences [19, 1]. While early works [15, 21] use hand-crafted features or matching proce-




Figure 1: Results of multi-scan alignment under the few-view setting. Each block shows the results from 5 randomly sampled scans of a 3D scene that show minimal overlaps. (a) RobustRecons [9]. (b) Our approach. (c) Ground-truth.

dures, recent works [16, 33, 39] have focused on applying deep neural networks to aid this pipeline. The challenge in the extreme relative pose setting is that there are insufficient features to match, and one has to leverage priors about the underlying object/scene. In this paper, we focus on two aspects of utilizing deep neural networks to instill priors into feature extraction and feature matching.

The first aspect focuses on completing the input scans and matching completed scans [49]. This strategy learns from data to predict the surrounding regions of each input scan, turning non-overlapping scans into overlapping scans. A critical factor in this context is the data representation for scan completion. This representation not only dictates the extracted features for relative pose estimation but also affects the type of completion networks. Instead of focusing on one data representation, a key contribution of this paper is to combine multiple data representations (i.e., 360-image [49], 2D layout [40], and planar patches [32]) for scan completion. Our approach then extracts a subset of consistent features to compute the relative pose. The proposed hybrid representation exhibits two appealing advantages. First, it alleviates the issue of lacking a particular type of feature on the input scans (e.g., planar elements from

\*Both authors contributed equally

†huangqx@cs.utexas.edu

non-planar scenes). Second, the matched features are adaptive to the input scan type and the relation between the input scans (e.g., dense point-wise features among overlapping regions and primitive features for non-overlapping scans).

The second aspect lies in the type of constraints between features for relative pose estimation. Instead of using a single module, our approach follows a global-2-local procedure. This procedure makes it possible to utilize different constraints at various pipeline stages. Specifically, the global matching module computes initial relative poses by performing spectral matching [28, 20] to align features extracted from the completed scans. The local refinement module, which benefits from having initial relative poses, detects and integrates additional pairwise geometric relations (e.g., parallel planes and distances between planes) for pose refinement. This sequential procedure can be easily modified to output multiple relative poses to model ambiguities in extreme relative pose estimation.

We have evaluated our approach on scan pairs from three benchmark datasets, including SUNCG [44], Matterport [6], and ScanNet [11]. For randomly sampled scan pairs, the mean top-1 rotation/translation errors of our approach are  $18.1^\circ/0.16m$ ,  $22.3^\circ/0.39m$ , and  $28.6^\circ/0.90m$ , on SUNCG, Matterport, and ScanNet, respectively. In contrast, the state-of-the-art approach achieved  $31.1^\circ/0.27m$ ,  $34.2^\circ/0.53m$ , and  $34.1^\circ/0.93m$ , respectively. These improvements show the advantage of using a hybrid 3D representation for relative pose estimation. Moreover, the mean top-5 errors in rotation/translation dropped to  $7.8^\circ/0.21m$ ,  $14.3^\circ/0.41m$ , and  $16.8^\circ/0.76m$ . These performance gains show that ambiguities are abundant in extreme relative pose estimation, and the multiple solutions returned by our approach can effectively address such ambiguities.

Besides, we demonstrate the usefulness of our approach in multi-scan reconstruction. Experimental results show that combining our method and the pose optimizer of MRF-SFM [10] enables faithful 3D reconstructions from a few views and outperforms state-of-the-art multi-scan reconstruction approaches (See Figure 1).

## 2. Related Works

**Relative pose estimation via feature matching.** Global relative pose estimation approaches usually fall into two categories (c.f. [30, 48]). The first category of approaches leverages global descriptors or performs Fourier transform to align two input scans. These methods are most suitable for matching complete objects. The second category of approaches performs feature matching, e.g., RANSAC [13], robust regression [4], and spectral matching [20, 29], to estimate relative poses. However, these approaches still require that the input scans possess considerable overlapping regions. In contrast, this paper focuses on relative pose estimation between potentially non-overlapping scans.

**Relative pose estimation via local optimization.** Besides

global matching, another category of relative pose estimation approaches focuses on local pose refinement [46]. A standard method is geometric alignment [3, 7], which minimizes point-wise distances between a pair of scans in their overlapping region. However, existing approaches typically require that the input scans are overlapping (c.f. [46]). In contrast, this paper proposes to solve the local pose refinement problem by learning and enforcing geometric relation (perpendicular, parallel, co-planar), which is suitable for both overlapping and non-overlapping scans.

**Learning based relative pose estimation.** Thanks to advances in deep learning that provide powerful tools to establish maps between different domains, recent works [33, 47, 51, 24, 39] formulate relative pose estimation as training pose estimation networks. Early works [33, 47, 51] usually follow the procedure of feature extraction and then relative pose estimation using a correlation module. Recent works [24, 39] employ more sophisticated matching techniques such as recurrent neural networks. However, most existing approaches still require that the two input scans possess significant overlaps, with the noticeable exception of [31], which first estimate local layout for partial scans then use hand-coded floorplane priors to assemble all pieces. Our approach differs in that we leverage the priors by explicitly learning a hybrid completion components from data.

This work is most relevant to [49], which presents a relative pose estimation approach that first performs scene completion and then computes a single relative pose prediction via geometric matching on the completed scenes. The difference in this work is that we employ a hybrid 3D scene completion in contrast to the 360-image employed in [49]. Besides, our approach introduces a novel geometric matching procedure that outputs multiple solutions and a learned local pose refinement module.

**Scene completion.** Our approach is also motivated by recent advances in inferring complete environments from partial observations [37, 23, 45, 22, 52]. However, our approach differs from these approaches in two ways. First, in contrast to returning the completion as the final output [45, 52] or utilizing it for learning feature representations [37, 22] or motion policies [23], our approach treats completions as an intermediate representation for feature extraction and detecting geometric relations. More importantly, instead of using a single representation for scene completion, our approach utilizes multiple representations and lets the geometric matching module to pick the most suitable representations for relative pose estimation.

## 3. Approach

This section describes the technical details of our approach. Section 3.1 presents an overview of our approach. Section 3.2 to Section 3.4 discuss each component in depth.

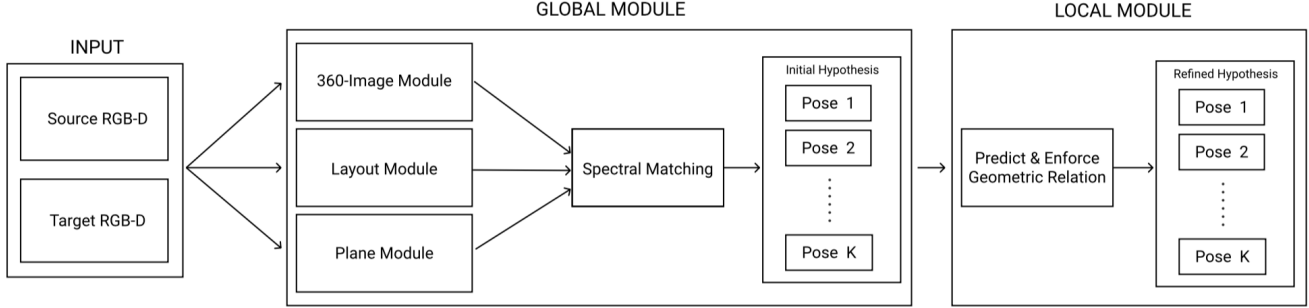


Figure 2: Our approach takes two potentially non-overlapping RGB-D scans as input and outputs one or multiple relative poses between them. The proposed pipeline combines a global module and a local module. The global module takes a pair of RGB-D scans as input and outputs multiple relative poses between them. The local module refines each relative pose by enforcing predicted geometric relations.

### 3.1. Problem Statement and Approach Overview

Our approach takes two RGB-D scans  $S_1$  and  $S_2$  of the same scene as input and outputs multiple rigid transformations  $\mathcal{T}_{12} = \{T_{12}\} \subset SE(3)$  that align the two input scans. As illustrated in Figure 2, our approach combines a global module, which generates a set of candidate relative poses, and a local module, which refines each candidate relative pose. In doing so, our approach can output multiple relative poses.

**Global module.** This module follows the standard pipeline of first extracting features from the input scans and then matching the extracted features. It begins with performing scan completion to generate sufficient shared features between small-overlapping or non-overlapping scans. In particular, our approach utilizes three intermediate representations, i.e., 360-image [49], 2D layout [40], and planar primitives [32], for scan completion. Given the completed scans, this module then establishes consistent correspondences between extracted features and fit a rigid transformation to each consistent set. To this end, we apply spectral matching, which extracts consistent subsets of feature correspondences by computing leading eigenvectors of a consistency matrix among candidate correspondences. Each resulting correspondence set generates an initial relative pose by rigid pose regression.

**Local module.** In the same spirit as using ICP [3] to optimize relative poses, our approach employs a local module to refine each output of the global module. The motivation comes from the fact initial relative poses enable additional constraints, such as iteratively computed dense correspondences and geometric relations (i.e., co-planar, parallel planes, and perpendicular planes), for relative pose estimation. Our approach also employs robust norms to remove outliers in dense correspondences and geometric relations.

**Network training.** A fundamental property of our approach is that the local module solves a robust optimization problem. Its optimal solution is generally insensitive to small perturbations to the outputs of the global module. Thus, we introduce two objectives for network training;

namely, the optimal solution of the local module is close to the underlying ground-truth, and the global module provides effective initial solutions. The total objective function for network training consists of three terms. The first term trains the completion networks. The second term jointly trains the completion networks and the spectral matching procedure. The third term is enforced on the output of the local module.

### 3.2. Global Module

The global module computes a set of relative poses  $\mathcal{T} = \{(R_{12}, \mathbf{t}_{12})\}$  between the input scans  $S_1$  with  $S_2$ . This module consists of two sub-modules. The first sub-module performs scan completion and feature extraction from each input scan  $S_i$ . The result is summarized into a collection of features  $\mathcal{F}_{\theta_g}(S_i) = \{f\}$ , where  $\theta_g$  denotes network parameters. Each  $f = \{\mathbf{p}_f, \mathbf{n}_f, \mathbf{d}_f\}$  consists of a 3D position  $\mathbf{p}_f$ , a 3D normal  $\mathbf{n}_f$ , and a feature descriptor  $\mathbf{d}_f$ . Note that although all extracted features share the same encoding, this sub-module combines the strengths of scan completion networks that are available under multiple representations (e.g., 360-image, 2D-layout, and 3D planar patches).

The second sub-module computes initial relative poses from correspondences between these extracted features. Since the local module will refine each relative pose, this sub-module uses spectral matching [28, 20], which is efficient and can output multiple relative poses.

**Scan completion sub-module.** Our approach considers three data representations (See Figure 3). Since we primarily use state-of-the-art scan completion results, we only highlight the main characteristics of each completion network and defer the details to the supp. material.

The first representation is 360-image [45, 49], which represents a given 3D scene as a collection of multi-channel images. The same as [49], we encode for each pixel its position  $\mathbf{p}_f$ , normal  $\mathbf{n}_f$ . The completion network admits an encoder-decoder architecture and outputs predicted position, normal, and descriptor channels. This representation can generate dense completions. Besides, it is a generic representation that does not place any assumptions about

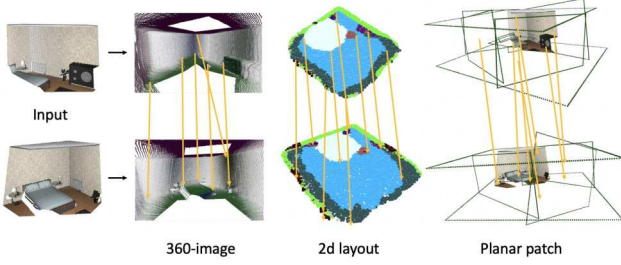


Figure 3: Matched correspondences between two scans across different representations. All representations help. Note that to make the visualization uncluttered, we only show a subset of computed correspondences and draw edges of the walls.

the underlying scene. However, the quality of the predicted features drops substantially among regions that are far from the visible area of each scan (c.f. [49]).

The second representation is 2D layout. We first to predict the plane of the floor of each scan. We then turn each input scan into an top-down view image by projecting points of the scan onto a grid representation of the floor. Similar to [27], the feature descriptor of each pixel collects the averaged feature among 4 vertical bins. We train the feature descriptor jointly with contrastive loss same as [49] and semantic segmentation loss. The completion network again employs an encoder-decoder architecture.

The third representation is planar patches [45, 43, 32]. To generate this representation, we perform RANSAC [14] to extract planar patches from each input scan and the corresponding complete 3D scene. We then perform PCA on the points associated with each patch to extract its patch center. The position  $\mathbf{p}_f$  and normal  $\mathbf{n}_f$  of this planar feature are simply the patch center and plane normal, respectively. Its feature descriptor  $\mathbf{d}_f$  is the mean feature descriptor among the points associated with this planar patch. The completion network employs a variant of PointNet [38] by treating each plane as a generalized point. Experimentally, we found that this planar patch representation leads to better generalization performance than the first two representations, but it only involves planar patches.

**Spectral matching sub-module.** The output of the completion sub-module is a feature set  $\mathcal{F}_{\theta_g}(S_i)$  for each scan  $S_i$  (that include dense pixels and planar features). Let  $\overline{\mathcal{M}}_{\theta_g,p} \subset \mathcal{F}_{\theta_g}(S_1) \times \mathcal{F}_{\theta_g}(S_2)$  collect all correspondences between the same type of features for the scan pair  $p = (S_1, S_2)$ . For representations of 360-image and 2D layout, we match SIFT keypoint among the visible region of one scan to the closest point on the other scan (in terms of feature descriptors) to reduce the number of feature correspondences. As the global module only provides initial solutions for the local module, we found that increasing the number of feature matches has minor impacts on the results of the local module.

The spectral matching sub-module extracts consistent correspondence sets  $\mathcal{M}_k \subset \overline{\mathcal{M}}_{\theta_g,p}$ ,  $1 \leq k \leq K$  ( $K = 5$  in

our experiments). Each correspondence set generates one initial relative pose. This step follows [28, 49], which construct a consistency matrix  $C_{\theta_g,\alpha_g,p} \in \mathbb{R}^{|\overline{\mathcal{M}}_{\theta_g,p}| \times |\overline{\mathcal{M}}_{\theta_g,p}|}$ . The value of each element

$$C_{\theta_g,\alpha_g,p}(c, c') = s_c s_{c'} g_{c,c'}$$

where  $c$  and  $c'$  represent a pair of correspondences,  $s_c$  models the descriptor similarity of  $c$ ,  $g_{c,c'}$  measures the preservation of distances and angles between  $c$  and  $c'$ , and  $\alpha_g$  represent hyper-parameters of  $s_c$  and  $g_{c,c'}$ . Due to the space constraint, please refer to the supp. material for details.

Given  $C_{\theta_g,\alpha_g,p}$ , we follow [28, 49] to extract top  $K$  matches, each of which is given by an indicator vector  $\mathbf{u}_k \in [0, 1]^{|\overline{\mathcal{M}}_{\theta_g,p}|}$ . We employ the rigid regression formulation in [18] to obtain the initial relative pose  $R_k^*, \mathbf{t}_k^*$ :

$$\min_{R, \mathbf{t}} \sum_{c=(f_1, f_2)} u_k(c) \|R\mathbf{p}_{f_1} + \mathbf{t} - \mathbf{p}_{f_2}\|^2 \quad (1)$$

### 3.3. Local Module

The local module refines each initial relative pose obtained from the global module. Benefited from having initial relative poses, this module employs different sets of features for refinement. Specifically, it performs robust regression on iteratively updated closest point pairs and predicted geometric relations (e.g., point pairs that lie on the same plane) between the completed scans. In the following, we first introduce how to predict geometric relations. We then describe how to perform robust regression.

**Geometric relations.** We consider three types of geometric relations that frequently occur in man-made scenes: point pairs that share the same plane, point pairs that lie on perpendicular planes, and point pairs that lie on parallel planes. Consider a down-sampled set of features  $\mathcal{F}_{\theta_g}^g(S_i) \subset \mathcal{F}_{\theta_g}(S_i)$  for each scan  $S_i$ . Our implementation uses 6000 random samples for each scan pair. Given all feature pairs  $\hat{\mathcal{M}}_{\theta_g,p}^l$  between these features the local module trains a variant of PointNet++ [38] to classify each pair  $c \in \hat{\mathcal{M}}_{\theta_g,p}^l$  into either no relation or one of the relations described above, i.e., four categories in total. In the following, we denote the probability of  $c$  belonging to the categories of co-planar, parallel, and perpendicular as  $w_{\theta_l}^1(c) \in [0, 1]$ ,  $w_{\theta_l}^2(c) \in [0, 1]$ , and  $w_{\theta_l}^3(c) \in [0, 1]$ , respectively, where  $\theta_l$  denotes the network parameters of the relation prediction network.

**Robust regression.** We begin with defining consistency scores between a predicted geometric relation  $c = (f_1, f_2)$  and the current rigid transformation  $(R, \mathbf{t})$ :

$$f_{R, \mathbf{t}}^1(c) := ((R\mathbf{p}_{f_1} + \mathbf{t} - \mathbf{p}_{f_2})^T \mathbf{n}_{f_2})^2 + ((R\mathbf{p}_{f_1} + \mathbf{t} - \mathbf{p}_{f_2})^T R\mathbf{n}_{f_1})^2 \quad (2)$$

$$f_R^2(c) := 1 - ((R\mathbf{n}_{f_1})^T \mathbf{n}_{f_2})^2 \quad (3)$$

$$f_R^3(c) := ((R\mathbf{n}_{f_1})^T \mathbf{n}_{f_2})^2 \quad (4)$$

In other words,  $f_1$ ,  $f_2$ , and  $f_3$  force a point pair to satisfy the co-planar, parallel, and perpendicular relations, respectively. In addition,  $f_1$  also serves as a symmetric point-to-plane measure for matching closest point pairs (c.f. [41]).

The robust regression formulation generalizes the standard approach of using iteratively reweighted non-linear least squares for rigid alignment [46]. Specifically, at each iteration, it solves the following non-linear least squares:

$$\begin{aligned} \min_{R, \mathbf{t}} \sum_{c \in \mathcal{M}_{\theta_c}^{nn}} w_c^0 f_{R, \mathbf{t}}^1(c) + \alpha_r \left( \sum_{c \in \mathcal{M}_{\theta_c}^1} w_c^1 w_{\theta_g}^1(c) f_{R, \mathbf{t}}^1(c) \right. \\ \left. + \sum_{c \in \mathcal{M}_{\theta_c}^2} w_c^2 w_{\theta_g}^2(c) f_{R, \mathbf{t}}^2(c) + \sum_{c \in \mathcal{M}_{\theta_c}^3} w_c^3 w_{\theta_g}^3(c) f_{R, \mathbf{t}}^3(c) \right) \quad (5) \end{aligned}$$

where  $\mathcal{M}_{\theta_c}^{nn}$  collects nearest neighbor pairs from features of  $\mathcal{F}_{\theta_c}(S_1)$  to  $\mathcal{F}_{\theta_c}(S_2)$ , and vice-versa, based on the relative pose at the previous iteration.  $w_c^0$ ,  $w_c^1$ ,  $w_c^2$ , and  $w_c^3$  are updated adaptively based on the current relative pose  $R, \mathbf{t}$ :

$$\begin{aligned} w_c^0 &:= \alpha_0^2 / (\alpha_0^2 + f_{R, \mathbf{t}}^1(c)), \quad w_c^1 := \alpha_1^2 / (\alpha_1^2 + f_{R, \mathbf{t}}^1(c)) \\ w_c^2 &:= \alpha_2^2 / (\alpha_2^2 + f_{R, \mathbf{t}}^2(c)), \quad w_c^3 := \alpha_3^2 / (\alpha_3^2 + f_{R, \mathbf{t}}^3(c)) \quad (6) \end{aligned}$$

where  $\alpha_l = \{\alpha_r, \alpha_0, \alpha_1, \alpha_2, \alpha_3\}$  are hyper-parameters, which will be trained together with  $\theta_g$  and  $\theta_l$ .

With this setup, the robust regression procedure first applies (6) to set the term weights using the initial relative pose. It then alternates among computing the nearest neighbors, applying the Gauss-Newton method to solve (5), and updating the term weights. In our implementation, we run 3 alternating iterations.

### 3.4. Network Training

Since the local module solves a robust optimization problem, its optimal solution is insensitive to small perturbations to the outputs of the global module. Thus, instead of training the combination of the local module and the global module end-to-end, we utilize a decoupled approach in this paper. Specifically, our approach optimizes all variables  $\Theta = \{\theta_g, \alpha_g, \theta_l, \alpha_l\}$  by minimizing the following loss term:

$$\min_{\Theta} l_1(\theta_g) + \lambda_2 l_2(\theta_l) + \lambda_3 l_3(\alpha_g) + \lambda_4 l_4(\alpha_l) \quad (7)$$

where  $l_1$ ,  $l_2$ ,  $l_3$ , and  $l_4$  train scan completion networks, the geometric relation network, the spectral matching sub-module, and the local module, respectively.  $\lambda_i$  are hyper-parameters determined via 10-fold cross validation.

The scan completion loss  $l_1(\theta_g)$  combines the losses for training the scan completion network under each data representation. The geometric relation loss  $l_2(\theta_l)$  utilizes initial relative poses that randomly perturbed from the underlying ground-truth. Since the formulations of these two losses are standard for network training, we leave the details to the supp. material. In the following, we focus on the loss terms for spectral matching sub-module and the local module.

**Loss term for spectral matching.** Consider the consistency matrix  $C_p$  for a scan pair  $p = (S_1, S_2)$ . Our goal for training the spectral matching sub-module is to enforce that the normalized indicator vector  $\mathbf{u}_p^{gt}$  of ground-truth correspondences is a top eigenvector of  $C_p$ . Note that  $C_p$  is dependent on  $\alpha_g$ . We omit them to make the notations uncluttered.

To avoid training a recurrent neural network converted from power iteration for eigen-vector computation (c.f. [36]), our approach utilizes three properties of a top eigen-vector, i.e., a normalized vector  $\mathbf{u}$  is a top eigen-vector of a matrix  $C$  if (1) the corresponding eigenvalue  $\lambda = \mathbf{u}^T C \mathbf{u}$ , (2)  $C \mathbf{u} = \lambda \mathbf{u}$ , and (3)  $\lambda$  is maximized. Given a collection of scan pairs  $\mathcal{P} = \{p\}$  as, we define  $l_3(\alpha_g) :=$

$$\sum_{p \in \mathcal{P}} (\|C_p \mathbf{u}_p^{gt} - (\mathbf{u}_p^{gt T} C_p \mathbf{u}_p^{gt}) \mathbf{u}_p^{gt}\|^2 - \mathbf{u}_p^{gt T} C_p \mathbf{u}_p^{gt}).$$

**Loss term for the local module.** Define  $\mathcal{I}_p(R, \mathbf{t}) \in \mathbb{R}^{3 \times 4}$  as one iteration of the IRNLS procedure from the current pose  $(R, \mathbf{t})$ , i.e., optimization and reweighting. Similar to  $l_3$ , we avoid training recurrent networks of the robust regression procedure by directly penalizing the difference between the ground-truth relative pose  $R_p^{gt}, \mathbf{t}_p^{gt}$  and  $\mathcal{I}_p(R^{gt}, \mathbf{t}^{gt})$ :

$$l_4(\alpha_l) = \sum_{p \in \mathcal{P}} \|(R_p^{gt}, \mathbf{t}_p^{gt}) - \mathcal{I}_p(R^{gt}, \mathbf{t}^{gt})\|_{\mathcal{F}}^2 \quad (8)$$

where  $\|\cdot\|_{\mathcal{F}}$  is the matrix Frobenius-norm. One intuition of (8) is that when  $(R_p^{gt}, \mathbf{t}_p^{gt}) = \mathcal{I}_p(R^{gt}, \mathbf{t}^{gt})$ , then  $(R_p^{gt}, \mathbf{t}_p^{gt})$  becomes the stationary point of the IRNLS procedure.

Minimizing  $l_4(\alpha_l)$  requires computing the derivatives of  $\mathcal{I}_p$ . Since  $\mathcal{I}_p$  is based on a local optimum of (5), we use the implicit function theorem [26] to back-propagate the gradients from the optimized poses to the predicted features and geometric relations. Again, we defer the details to the supp. material for the interest of space.

**Training procedure.** We train (7) through three phases. The first phase trains each prediction network in isolation. The second phase fixes the prediction networks and trains the hyper-parameters of each module in isolation. The final phase fine-tunes all the variables together by solving (7). Network training involves the ADAM optimizer [25], and please refer to the supp. material for details.

## 4. Experimental Evaluation

This section presents an experimental evaluation of the proposed approach. Section 4.1 describes the experimental setup. Section 4.2 analyzes the experimental results and compares the proposed approach with baseline approaches. Section 4.3 and Section 4.4 analyze the global module and the local module, respectively.




Figure 4: Qualitative comparison between our approach and baseline approaches. From left to right, we show the results of Super4PCS [34], RobustGR [50], ScanComplete [49], Ours-Top1, Ours-Top3, and Ground Truth figures.

#### 4.1. Experimental Setup

**Datasets.** We perform experimental evaluation using three datasets: *SUNCG* [44], *Matterport* [6], and *ScanNet* [11], where *SUNCG* provides synthetic data, *Matterport* and *ScanNet* provide real data. For each room in each dataset, we randomly sample 25 camera locations. For real data sets such as *ScanNet*, the camera locations are sampled from the recording sequences. For each dataset, we collect around 50k scan pairs for training, and 1k scan pairs for testing. Note that the scenes of the training and testing scans do not overlap.

**Baseline approaches.** We consider the five baseline approaches: *Super4PCS* [34], *RobustGR* [50], *ScanComplete* [49], *Ours-Global*, and *Ours-Local*. Here, *Super4PCS* is a widely used global scan matching method; *RobustGR* is a state-of-the-art global scan matching method that is based on robust regression; *ScanComplete* is a state-of-the-art approach for indoor RGB-D registration that performs

scan completion using the 360-image representation. Moreover, *Ours-Global* is our approach without the local module. *Ours-Local* is our approach without the global module, i.e., applying the local module directly on the completed scans.

**Evaluation protocol.** We follow the standard protocol of reporting the relative angular error  $\|\log(R^* R^{gtT})\|/\sqrt{2}$  and the relative translation error  $\|\mathbf{t}^{gt} - \mathbf{t}^*\|$ . Here  $(R^{gt}, \mathbf{t}^{gt})$  and  $(R^*, \mathbf{t}^*)$  denote the ground-truth and the output of a method, respectively. In the presence of multiple outputs, we define the best match as the one that minimizes the rotation error. In this section, we report the rotation and the translation errors among the top 1, 3, and 5 matches. The order of the matches are given by the eigenvalues of the consistency matrix  $C_{\theta_g, \alpha_g, p}$  and the final scores of (5), respectively. In addition, we divide scan pairs into two categories: overlapping and non-overlapping. The category of overlapping scans consists of scans pairs that overlap by at least 10% points (with respect to the scan with fewer num-

	SUNCG						Matterport						ScanNet					
	Rotation			Trans			Rotation			Trans			Rotation			Trans		
	top1	top3	top5	top1	top3	top5	top1	top3	top5	top1	top3	top5	top1	top3	top5	top1	top3	top5
360-image ( $\geq 10\%$ )	9.8	4.5	3.5	0.17	0.16	0.16	19.9	10.4	8.5	0.39	0.38	0.38	18.4	11.0	8.7	0.71	0.54	0.50
Plane ( $\geq 10\%$ )	100.3	65.9	45.8	0.58	0.70	0.59	123.6	92.4	78.6	0.79	0.99	0.99	123.6	91.4	79.1	3.03	2.63	2.65
2D-layout ( $\geq 10\%$ )	52.0	23.1	16.4	0.85	0.50	0.52	73.1	42.2	33.8	1.57	1.38	1.36	45.5	25.2	18.3	1.31	1.07	1.01
Ours ( $\geq 10\%$ )	<b>8.0</b>	<b>3.6</b>	<b>3.3</b>	<b>0.16</b>	<b>0.14</b>	<b>0.14</b>	<b>13.9</b>	<b>9.7</b>	<b>7.8</b>	<b>0.32</b>	<b>0.31</b>	<b>0.31</b>	<b>16.7</b>	<b>10.5</b>	<b>8.7</b>	<b>0.69</b>	<b>0.53</b>	<b>0.50</b>
360-image ( $\leq 10\%$ )	79.7	30.4	19.7	0.49	0.44	0.40	70.0	42.6	34.1	0.88	0.83	0.80	66.4	46.9	36.1	1.37	1.34	1.35
Plane ( $\leq 10\%$ )	123.4	80.4	57.8	0.65	0.86	0.72	135.8	77.8	68.9	0.71	0.96	0.74	119.5	86.4	81.5	2.56	2.60	2.67
2D-layout ( $\leq 10\%$ )	99.9	55.6	41.4	1.18	0.76	0.82	107.0	63.8	52.5	1.42	1.69	1.74	73.7	46.3	34.2	1.94	1.90	1.89
Ours ( $\leq 10\%$ )	<b>62.8</b>	<b>25.6</b>	<b>18.1</b>	<b>0.40</b>	<b>0.38</b>	<b>0.38</b>	<b>58.9</b>	<b>39.9</b>	<b>32.3</b>	<b>0.69</b>	<b>0.68</b>	<b>0.68</b>	<b>59.3</b>	<b>40.4</b>	<b>33.5</b>	<b>1.36</b>	<b>1.33</b>	<b>1.33</b>
360-image (all)	31.1	12.4	8.5	0.27	0.24	0.23	34.2	19.0	15.3	0.53	0.50	0.49	34.1	22.7	17.7	0.93	0.79	0.78
Plane(all)	107.3	70.3	49.4	0.60	0.75	0.63	126.4	89.1	76.4	0.78	0.99	0.95	122.4	89.9	79.8	2.89	2.60	2.65
2D-layout(all)	66.6	33.0	24.0	0.95	0.58	0.61	82.1	47.9	38.8	1.53	1.46	1.46	56.3	32.1	23.5	1.50	1.33	1.30
Ours (all)	<b>24.7</b>	<b>10.3</b>	<b>7.8</b>	<b>0.23</b>	<b>0.21</b>	<b>0.21</b>	<b>26.0</b>	<b>17.8</b>	<b>14.3</b>	<b>0.41</b>	<b>0.41</b>	<b>0.41</b>	<b>30.7</b>	<b>20.3</b>	<b>16.8</b>	<b>0.91</b>	<b>0.76</b>	<b>0.76</b>

Table 1: Ablation study on global module. We show the comparison between our hybrid representation module and the single representation module. The 10% refers to the overlap ratio. For each module, we also show the top1/top3/top5 error for rotation and translation.

	SUNCG		Matterport		ScanNet	
	Rotation	Trans	Rotation	Trans	Rotation	Trans
4PCS-Overlap ( $\geq 10\%$ )	45.2	0.22	49.7	0.53	31.3	0.94
RobustGR ( $\geq 10\%$ )	64.9	0.57	38.1	0.69	45.4	1.20
ScanComp. ( $\geq 10\%$ )	9.8	0.17	19.9	0.39	18.4	0.71
Ours-Local ( $\geq 10\%$ )	36.0	0.22	40.8	0.49	25.4	0.93
Ours-Global ( $\geq 10\%$ )	8.0	0.16	13.9	0.31	16.7	0.69
Ours ( $\geq 10\%$ )	<b>5.0</b>	<b>0.11</b>	<b>10.7</b>	<b>0.29</b>	<b>14.1</b>	<b>0.67</b>
ScanComp. ( $\leq 10\%$ )	79.7	0.49	70.0	0.88	66.4	1.37
Ours-Local ( $\leq 10\%$ )	105.0	<b>0.31</b>	115.6	<b>0.49</b>	67.2	<b>1.34</b>
Ours-Global ( $\leq 10\%$ )	62.8	0.40	59.3	0.69	59.3	1.36
Ours ( $\leq 10\%$ )	<b>60.9</b>	0.41	<b>54.1</b>	0.68	<b>58.4</b>	1.36
ScanComp.(all)	31.1	0.27	34.2	0.53	34.1	0.93
Ours-Local (all)	56.9	0.24	60.7	0.49	36.6	0.97
Ours-Global (all)	24.7	0.23	26.0	0.41	30.7	0.91
Ours (all)	<b>18.1</b>	<b>0.16</b>	<b>22.3</b>	<b>0.39</b>	<b>28.6</b>	<b>0.90</b>

Table 2: Benchmark evaluation on our approach and baseline approaches. Ours-Local and Ours-Global stand for our method with global module and local module removed, respectively. We show the mean error for rotation and translation components for overlapping scan pairs ( $\geq 10\%$ ) and non-overlapping scan pairs ( $\leq 10\%$ ), respectively.

ber of points). The category of non-overlapping contains the remaining scans.

## 4.2. Analysis of Results

Table 2 and Figure 4 provide quantitative and qualitative results of our approach and baseline methods. Overall, our approach outperforms baseline approaches considerably. The mean rotation/translation errors of the best-match of our approach are  $18.1^\circ/0.16\text{m}$ ,  $22.3^\circ/0.39\text{m}$ , and  $28.6^\circ/0.90\text{m}$  on SUNCG, Matterport, and ScanNet, respectively. In contrast, the state of the art method only achieved  $31.1^\circ/0.27\text{m}$ ,  $34.2^\circ/0.53\text{m}$ , and  $34.1^\circ/0.93\text{m}$ , respectively.

**Overlapping v.s. non-overlapping.** In the overlapping regime, the relative improvements of mean rotation/translation errors are  $40.8\%/35.3\%$ ,  $46.2\%/25.6\%$ , and  $55.0\%/15.6\%$  on the three datasets described above, respec-

tively. Meanwhile, the relative improvements in the non-overlapping regime are  $23.6\%/16.3\%$ ,  $22.9\%/22.7\%$ , and  $10.5\%/0.7\%$ , respectively. These statistics show the consistency of our approach across different overlapping rates.

## 4.3. Analysis of the Global Module

We focus on analyzing two important aspects of the global module, namely the effects of using hybrid representation and multiple outputs.

**Single representation versus hybrid representation.** As shown in Table 1, using hybrid representations leads to noticeable performance gains when compared to using a single representation. For mean rotation/translation errors, the relative improvements from the top performing single representation (*ScanComplete* in this case) are  $20.4\%/14.8\%$ ,  $24.0\%/22.6\%$ , and  $11.7\%/2\%$  on SUNCG, Matterport, and ScanNet, respectively. It should be noted that the 2D layout and Plane representations individually do not provide sufficient features for relative pose estimation. However, they provide complementary features to the 360-image representation that boost the overall performance of our approach. This is also consistent with Figure 3, which shows that the matched features consist of a mixture of planar features and point features from 360-image and 2D layout.

**Best match versus K-best.** With multiple outputs, the top-5 rotation/translation errors of the global module drop significantly. The relative improvements on the three datasets are  $68.4\%/8.7\%$ ,  $45.0\%/0.1\%$ , and  $45.2\%/16.5\%$ , respectively. Besides, the relative improvements on non-overlapping scans are significantly bigger than those on overlapping scans. An explanation is that indoor scenes exhibit approximate discrete symmetries, e.g., rotational symmetries of a box-shape, and such symmetries incur ambiguities for non-overlapping scans (i.e., they lack detailed geometric features to lock their relative poses). On the other hand, the orders of these symmetries are relatively small, and top matches tend to capture them.

	SUNCG		Matterport		ScanNet	
	Rot	Trans	Rot	Trans	Rot	Trans
Base-Global+Base-Local	31.9	0.45	33.6	0.52	34.1	0.93
Base-Global+Ours-Local	27.8	0.24	30.1	0.51	31.2	0.92
Ours-Global+Base-Local	25.9	0.44	25.4	0.40	30.4	0.91
Ours-Global+Ours-Local	<b>18.1</b>	<b>0.16</b>	<b>22.3</b>	<b>0.34</b>	<b>28.6</b>	<b>0.90</b>

Table 3: Benchmark evaluation on the different combination of global initialization and local refinement. We use the strongest baseline [49] to serve as global module, and SparseICP [5] as baseline local module. We show the mean rotation and translation errors on each dataset.

In addition, the probability distribution of the rank of the matches with the best rotational error are 42%, 17%, 17%, 10%, and 14%, respectively. In other words, 58% of the best matches are not derived from the first eigen-vector. This indicates the importance of leveraging multiple outputs. In Section 4.5, we show that using multiple outputs for multi-scan reconstruction is superior to using a single output.

#### 4.4. Analysis of the Local Module

Overall, the local module leads to noticeable performance gains from the outputs of the global module. As shown in Table 2, the mean rotation/translation errors of the top 1 match reduced from 24.7°/0.23m to 18.1°/0.16m on SUNCG, from 26.0°/0.41m to 22.3°/0.39m on Matterport, and from 30.7°/0.91m to 28.6°/0.90m on ScanNet, respectively. These statistics indicate the robustness of our local module.

**Baseline comparison.** Our approach outperforms the robust version of ICP [5] when applied on the output of the global module (See Table 3). The mean rotation/translation errors reduce from 25.9°/0.44m, 25.4°/0.40m, and 30.4°/0.91m to 18.1°/0.16m, 22.3°/0.34m, and 28.6°/0.90m. Such improvements mainly come from the predicted geometric relations. Our approach also outperforms the baseline that combines the output of Scan-Complete and our local module. The relative improvements are salient across all the datasets. These statistics indicate that having good initial poses for local refinement is key to the success of relative pose estimation.

**The quality of geometric relations prediction.** The average accuracy of geometric relations across all three datasets are 46%, 69%, 75%, and 93% for no-relation, perpendicular, parallel, co-planar respectively. The statistics are collected from scan pairs perturbed from the underlying ground-truth by a random rotation with degree in  $[0, 30^\circ]$  and a random translation in  $[-0.5m, 0.5m]^3$ . Although the predictions are imperfect, the success of our approach lies in using robust norms to filter out wrong predictions automatically. Due to space constraints, we leave a detailed analysis of the predicted geometric relations to the supp. material.

	5 Scans		10 Scans		15 Scans		20 Scans	
	Rot	Trans	Rot	Trans	Rot	Trans	Rot	Trans
Fine2C.[17]	96.5	1.41	62.4	0.93	69.6	1.57	68.6	1.43
RobustRec[9]	33.1	0.78	25.1	0.52	16.5	0.33	12.3	0.26
Our-top-1	8.39	0.20	3.55	0.15	2.97	0.134	2.31	0.12
Our-top-3	5.77	<b>0.18</b>	3.12	<b>0.12</b>	2.26	0.112	1.88	0.10
Our-top-5	<b>5.13</b>	<b>0.18</b>	<b>2.73</b>	<b>0.12</b>	<b>2.13</b>	<b>0.10</b>	<b>1.81</b>	<b>0.09</b>

Table 4: Comparisons on multi-scan alignment using our approach + MRF-SFM[10] and two baseline approaches. Our approach shows a clear advantage in the sparse setting, e.g., using 5 scans.

#### 4.5. Application in Few-View Reconstruction

We have studied the impacts of our relative pose estimation approach in multi-scan alignment. To this end, we pair our approach with the multi-view pose optimizer of MRF-SFM [10]. The goal is to find the most consistent subset of relative poses among the outputs of our approach. To analyze the benefits of having multiple relative pose estimations between each scan pair, we report the performance of using top-1, top-3, top-5 outputs of our approach, respectively. We performed an experimental study on ScanNet, using 5/10/15/20 randomly sampled scans. Regarding the evaluation protocol, we report absolute errors in rotations and translations of aligned scans using the same metric for evaluating relative poses. Since MRF-SFM extracts consistent relative transformations, the mean rotation/translation errors of the resulting absolute transformations are significantly reduced. The baseline approaches include RobustRecons [9] and Fine-to-Coarse [17], which are two state-of-the-art multi-scan alignment approaches. Fine-to-Coarse additionally utilizes the scanning order among the input scans.

As shown in Figure 1 and Table 4, our approach significantly outperforms baseline approaches. The improvements are salient in the sparse setting, where our approach shows high-quality outputs. In such a sparse setting, RobustRecons also outperforms Fine-to-Coarse, as the latter has to utilize dense inputs to detect initial planes. Utilizing top-3 and top-5 matches also improves the performance. One explanation is that there are more correct inputs for multi-scan alignment when utilizing top-3 and top-5 matches.

### 5. Conclusions and Future Work

In this paper, we have introduced an approach for estimating the relative pose between two potentially non-overlapping RGB-D scans. Our approach combines a global module for computing candidate relative poses and a local module for refining each candidate pose. Our method can output multiple relative poses, and the resulting pose quality is considerably better than state-of-the-art approaches for both overlapping and non-overlapping scans.



## References

- [1] Henrik Andreasson and Todor Stoyanov. Real time registration of rgb-d data using local visual features and 3d-ndt registration. In *SPME Workshop at Int. Conf. on Robotics and Automation (ICRA)*, 2012. [1](#)
- [2] Adrien Angeli, David Filliat, Stéphane Doncieux, and Jean-Arcady Meyer. A fast and incremental method for loop-closure detection using bags of visual words. *IEEE Transactions on Robotics*, pages 1027–1037, 2008. [1](#)
- [3] Paul J. Besl and Neil D. McKay. A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(2):239–256, Feb. 1992. [2, 3](#)
- [4] Sofien Bouaziz, Andrea Tagliasacchi, and Mark Pauly. Sparse iterative closest point. In *Proceedings of the Eleventh Eurographics/ACMSIGGRAPH Symposium on Geometry Processing*, pages 113–123. Eurographics Association, 2013. [2](#)
- [5] Sofien Bouaziz, Andrea Tagliasacchi, and Mark Pauly. Sparse iterative closest point. In *Proceedings of the Eleventh Eurographics/ACMSIGGRAPH Symposium on Geometry Processing*, SGP '13, pages 113–123, Aire-la-Ville, Switzerland, Switzerland, 2013. Eurographics Association. [8](#)
- [6] Angel X. Chang, Angela Dai, Thomas A. Funkhouser, Maciej Halber, Matthias Nießner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from RGB-D data in indoor environments. In *2017 International Conference on 3D Vision, 3DV 2017, Qingdao, China, October 10-12, 2017*, pages 667–676, 2017. [2, 6](#)
- [7] Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. *Image Vision Comput.*, 10(3):145–155, Apr. 1992. [2](#)
- [8] Taeg Sang Cho, Shai Avidan, and William T Freeman. A probabilistic image jigsaw puzzle solver. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 183–190. IEEE, 2010. [1](#)
- [9] Sungjoon Choi, Qian-Yi Zhou, and Vladlen Koltun. Robust reconstruction of indoor scenes. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 5556–5565, 2015. [1, 8](#)
- [10] David Crandall, Andrew Owens, Noah Snavely, and Daniel Huttenlocher. SfM with MRFs: Discrete-continuous optimization for large-scale structure from motion. *Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2013. [2, 8](#)
- [11] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas A. Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. *CoRR*, abs/1702.04405, 2017. [2, 6](#)
- [12] P. Doubek and T. Svoboda. Reliable 3d reconstruction from a few catadioptric images. In *Proceedings of the IEEE Workshop on Omnidirectional Vision 2002. Held in conjunction with ECCV'02*, pages 71–78, June 2002. [1](#)
- [13] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. [2](#)
- [14] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, June 1981. [4](#)
- [15] Natasha Gelfand, Niloy J. Mitra, Leonidas J. Guibas, and Helmut Pottmann. Robust global registration. In *Proceedings of the Third Eurographics Symposium on Geometry Processing*, SGP '05, Aire-la-Ville, Switzerland, Switzerland, 2005. Eurographics Association. [1](#)
- [16] Hunter Goforth, Yasuhiro Aoki, Arun Srivatsan Rangaprasad, and Simon Lucey. Pointnetlk: Robust and efficient point cloud registration using pointnet. In *Proceedings of IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [1](#)
- [17] Maciej Halber and Thomas A. Funkhouser. Fine-to-coarse global registration of RGB-D scans. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 6660–6669, 2017. [8](#)
- [18] Berthold K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America*, 4:629–642, 1987. [4](#)
- [19] Albert S Huang, Abraham Bachrach, Peter Henry, Michael Krainin, Daniel Maturana, Dieter Fox, and Nicholas Roy. Visual odometry and mapping for autonomous flight using an rgb-d camera. In *Robotics Research*, pages 235–252. Springer, 2017. [1](#)
- [20] Qi-Xing Huang, Bart Adams, Martin Wicke, and Leonidas J Guibas. Non-rigid registration under isometric deformations. In *Computer Graphics Forum*, volume 27, pages 1449–1457. Wiley Online Library, 2008. [2, 3](#)
- [21] Qi-Xing Huang, Simon Flöry, Natasha Gelfand, Michael Hofer, and Helmut Pottmann. Reassembling fractured objects by geometric matching. *ACM Trans. Graph.*, 25(3):569–578, July 2006. [1](#)
- [22] Dinesh Jayaraman, Ruohan Gao, and Kristen Grauman. Shapecodes: Self-supervised feature learning by lifting views to viewgrids. In *ECCV*, 2018. [2](#)
- [23] D. Jayaraman and K. Grauman. Learning to look around: Intelligently exploring unseen environments for unknown tasks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)
- [24] Seungryong Kim, Stephen Lin, SANG RYUL JEON, Dongbo Min, and Kwanghoon Sohn. Recurrent transformer networks for semantic correspondence. In *Advances in Neural Information Processing Systems*, pages 6126–6136, 2018. [2](#)
- [25] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [5](#)
- [26] Steven G. G. Krantz and Harold R. Parks. *The Implicit Function Theorem: History, Theory, and Applications*. Modern Birkhauser Classics. Springer, 2013. [5](#)
- [27] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12697–12705, 2019. [4](#)

- [28] Marius Leordeanu and Martial Hebert. A spectral technique for correspondence problems using pairwise constraints. In *IEEE International Conference on Computer Vision (ICCV)*, volume 2, pages 1482–1489, 2005. 2, 3, 4
- [29] Marius Leordeanu and Martial Hebert. A spectral technique for correspondence problems using pairwise constraints. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pages 1482–1489. IEEE, 2005. 2
- [30] Xin Li and SS Iyengar. On computing mapping of 3d objects: A survey. *ACM Computing Surveys (CSUR)*, 47(2):34, 2015. 2
- [31] Cheng Lin, Changjian Li, and Wenping Wang. Floorplan-jigsaw: Jointly estimating scene layout and aligning partial scans. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2
- [32] Chen Liu, Jimei Yang, Duygu Ceylan, Ersin Yumer, and Yasutaka Furukawa. Planenet: Piece-wise planar reconstruction from a single rgb image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2579–2588, 2018. 1, 3, 4
- [33] Iaroslav Melekhov, Juha Ylioinas, Juho Kannala, and Esa Rahtu. Relative camera pose estimation using convolutional neural networks. In *International Conference on Advanced Concepts for Intelligent Vision Systems*, pages 675–687. Springer, 2017. 1, 2
- [34] Nicolas Mellado, Dror Aiger, and Niloy J. Mitra. Super 4pcs fast global pointcloud registration via smart indexing. *Computer Graphics Forum*, 33(5):205–215, 2014. 6
- [35] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohli, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *Proceedings of the 2011 10th IEEE International Symposium on Mixed and Augmented Reality, ISMAR '11*, pages 127–136, Washington, DC, USA, 2011. IEEE Computer Society. 1
- [36] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28, ICML'13*, pages III–1310–III–1318. JMLR.org, 2013. 5
- [37] Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei A. Efros. Context encoders: Feature learning by inpainting. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2536–2544, 2016. 2
- [38] 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*, pages 652–660, 2017. 4
- [39] René Ranftl and Vladlen Koltun. Deep fundamental matrix estimation. In *European Conference on Computer Vision*, pages 292–309. Springer, 2018. 1, 2
- [40] Daniel Ritchie, Kai Wang, and Yu-an Lin. Fast and flexible indoor scene synthesis via deep convolutional generative models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6182–6190, 2019. 1, 3
- [41] Szymon Rusinkiewicz. A symmetric objective function for ICP. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 38(4), July 2019. 5
- [42] Szymon Rusinkiewicz, Olaf Hall-Holt, and Marc Levoy. Real-time 3d model acquisition. *ACM Trans. Graph.*, 21(3):438–446, July 2002. 1
- [43] Yifei Shi, Kai Xu, Matthias Nießner, Szymon Rusinkiewicz, and Thomas Funkhouser. Planematch: Patch coplanarity prediction for robust rgb-d reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 4
- [44] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. *Proceedings of 30th IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2, 6
- [45] Shuran Song, Andy Zeng, Angel X Chang, Manolis Savva, Silvio Savarese, and Thomas Funkhouser. Im2pano3d: Extrapolating 360 structure and semantics beyond the field of view. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2, 3, 4
- [46] Gary K. L. Tam, Zhi-Quan Cheng, Yu-Kun Lai, Frank C. Langbein, Yonghuai Liu, A. David Marshall, Ralph R. Martin, Xianfang Sun, and Paul L. Rosin. Registration of 3d point clouds and meshes: A survey from rigid to nonrigid. *IEEE Trans. Vis. Comput. Graph.*, 19(7):1199–1217, 2013. 2, 5
- [47] Benjamin Ummenhofer, Huizhong Zhou, Jonas Uhrig, Nikolaus Mayer, Eddy Ilg, Alexey Dosovitskiy, and Thomas Brox. Demon: Depth and motion network for learning monocular stereo. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 2
- [48] Oliver Van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. A survey on shape correspondence. In *Computer Graphics Forum*, volume 30, pages 1681–1707. Wiley Online Library, 2011. 2
- [49] Zhenpei Yang, Jeffrey Z. Pan, Linjie Luo, Xiaowei Zhou, Kristen Grauman, and Qixing Huang. Extreme relative pose estimation for rgb-d scans via scene completion. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4531–4540, June 2019. 1, 2, 3, 4, 6, 8, 11, 14
- [50] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Fast global registration. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II*, pages 766–782, 2016. 6
- [51] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised learning of depth and ego-motion from video. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 2
- [52] Chuhan Zou, Alex Colburn, Qi Shan, and Derek Hoiem. Layoutnet: Reconstructing the 3d room layout from a single rgb image. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2