

DeepI2P: Image-to-Point Cloud Registration via Deep Classification

Jiaxin Li
Bytedance

Gim Hee Lee
National University of Singapore

Abstract

This paper presents *DeepI2P*: a novel approach for cross-modality registration between an image and a point cloud. Given an image (e.g. from a rgb-camera) and a general point cloud (e.g. from a 3D Lidar scanner) captured at different locations in the same scene, our method estimates the relative rigid transformation between the coordinate frames of the camera and Lidar. Learning common feature descriptors to establish correspondences for the registration is inherently challenging due to the lack of appearance and geometric correlations across the two modalities. We circumvent the difficulty by converting the registration problem into a classification and inverse camera projection optimization problem. A classification neural network is designed to label whether the projection of each point in the point cloud is within or beyond the camera frustum. These labeled points are subsequently passed into a novel inverse camera projection solver to estimate the relative pose. Extensive experimental results on Oxford Robotcar and KITTI datasets demonstrate the feasibility of our approach. Our source code is available at <https://github.com/lijx10/DeepI2P>.

1. Introduction

Image-to-point cloud registration refers to the process of finding the rigid transformation, i.e., rotation and translation that aligns the projections of the 3D point cloud to the image. This process is equivalent to finding the pose, i.e., extrinsic parameters of the imaging device with respect to the reference frame of the 3D point cloud; and it has wide applications in many tasks in computer vision, robotics, augmented/virtual reality, etc.

Although the direct and easy approach to solve the registration problem is to work with data from the same modality, i.e., image-to-image and point cloud-to-point cloud, several limitations exist in these same-modality registration approaches. For *point cloud-to-point cloud* registration, it is impractical and costly to mount expensive and hard-to-maintain Lidars on large fleet of robots and mobile devices during operations. Furthermore, feature-based point cloud-to-point cloud registration [6, 43, 21, 40] usually requires

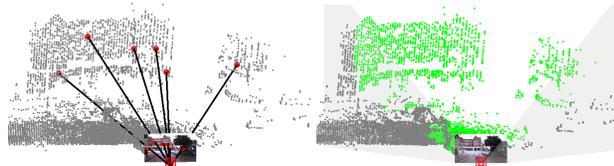


Figure 1. Illustration of feature based registration on the left, e.g., 2D3D-MatchNet, and our feature-free DeepI2P on the right. Instead of detecting and matching features across modalities, we convert the registration problem into a classification problem.

storage of D -dimensional features ($D \gg 3$) in addition to the (x, y, z) point coordinates, which increases the memory complexity. For *image-to-image* registration, meticulous effort is required to perform SfM [37, 36, 12] and store the image feature descriptors [29, 22] corresponding to the reconstructed 3D points for feature matching. Additionally, image features are subjected to illumination conditions, seasonal changes, etc. Consequently, the image features stored in the map acquired in one season/time are hopeless for registration after a change in the season/time.

Cross-modality image-to-point cloud registration can be used to alleviate the aforementioned problems from the same modality registration methods. Specifically, a 3D point cloud-based map can be acquired once with Lidars, and then pose estimation can be deployed with images taken from cameras that are relatively low-maintenance and less costly on a large fleet of robots and mobile devices. Moreover, maps acquired directly with Lidars circumvents the hassle of SfM, and are largely invariant to seasonal/illumination changes. Despite the advantages of cross-modality image-to-point cloud registration, few research has been done due to its inherent difficulty. To the best of our knowledge, 2D3D-MatchNet [11] is the only prior work on general image-to-point cloud registration. This work does cross-modal registration by learning to match image-based SIFT [22] to point cloud-based ISS [45] keypoints using deep metric-learning. However, the method suffers low inlier rate due to the drastic dissimilarity in the SIFT and ISS features across two modalities.

In this paper, we propose the DeepI2P: a novel approach for cross-modal registration of an image and a point cloud

without explicit feature descriptors as illustrated in Fig. 1. Our method requires lesser storage memory, i.e., $\mathcal{O}(3N)$ for the reference point cloud since we do not rely on feature descriptors to establish correspondences. Furthermore, the images captured by cameras can be directly utilized without SfM. We solve the cross-modal image-to-point cloud registration problem in two stages. In the first stage, we design a two-branch neural network that takes the image and point cloud as inputs, and outputs a label for every point that indicates whether the projection of this point is within or beyond the image frustum. The second stage is formulated as an unconstrained continuous optimization problem. The objective is to find the optimal camera pose, i.e., the rigid transformation with respect to the reference frame of the point cloud, such that 3D points labeled as within the camera frustum is correctly projected into the image. Standard solvers such as the Gauss-Newton algorithm can be used to solve our camera pose optimization problem. Extensive experimental results on the open-source Oxford Robotcar and KITTI datasets show the feasibility of our approach.

The main contributions of this paper are listed as follow:

- We circumvent the challenging need to learn cross-modal feature descriptor for registration by casting the problem into a two-stage classification and optimization framework.
- A two-branch neural network with attention modules to enhance cross-modality fusion is designed to learn labels of whether a 3D point is within or beyond the camera frustum.
- The inverse camera projection optimization is proposed to solve for the camera pose with the classification labels of the 3D points.
- Our method and the experimental results show a proof-of-concept that cross-modal registration can be achieved with deep classification.

2. Related Works

Image-to-Image Registration. Images-to-image registrations [33, 32] are done in the \mathbb{P}^2 space because of the lack of depth information. This is usually the first step to the computation of the projective transformation or SfM. Typical methods are usually based on feature matching. A set of features such as SIFT [22] or ORB [29] are extracted from both source and target images. Correspondences are then established based on the extracted features, which can be used to solve for the rotation, translation using Bundle Adjustment [36, 16], Perspective-n-Point solvers [12], etc. Such techniques have been applied in modern SLAM systems [10, 25, 9]. However, such methods are based on feature descriptors in the image modality to establish correspondences, and do not work for our general image-to-point cloud registration task.

Point Cloud-to-Point Cloud Registration. The availability of 3D information enables direct registration between point clouds without establishing feature correspondences. Methods like ICP [2, 5], NDT [3] work well with proper initial guess, and global optimization approaches such as Go-ICP [39] work without initialization requirements. These methods are widely used in point cloud based SLAM algorithms like LOAM [44], Cartographer [18], etc. Recently data driven methods like DeepICP [23], DeepClosestPoint [38], RPM-Net [41], etc, are also proposed. Although these approaches do not require feature correspondences, they still rely heavily on the geometrical details of the point structures in the same modality to work well. Consequently, these approaches cannot be applied to our task on cross-modal registration. Another group of common approaches is the two-step feature-based registration. Classical point cloud feature detectors [35, 45, 31, 8] and descriptors [34, 30] usually suffer from noise and clutter environments. Recently deep learning based feature detectors like USIP [21], 3DFeatNet [40], and descriptors like 3DMatch [43], PPF-Net [7], PPF-FoldNet [6], PerfectMatch [15], have demonstrated improved performances in point cloud-based registration. Similar to image-to-image registration, these approaches require feature descriptors that are challenging to obtain in cross-modality registration.

Image-to-Point Cloud Registration. To the best of our knowledge, 2D3D-MatchNet [11] is the only prior work for general image-point cloud registration. It extracts image keypoints with SIFT [22], and point cloud keypoints with ISS [45]. The image and point cloud patches around the keypoints are fed into each branch of a Siamese-like network and trained with triplet loss to extract cross-modal descriptors. At inference, it is a standard pipeline that consists of RANSAC-based descriptor matching and EPnP [19] solver. Despite its greatly simplified experimental settings where the point clouds and images are captured at nearby timestamps with almost zero relative rotation, the low inlier rate of correspondences reveals the struggle for a deep network to learn common features across the drastically different modalities. Another work [42] establishes 2D-3D line correspondences between images and prior Lidar maps, but they requires accurate initialization, e.g., from a SLAM/Odometry system. In contrast, the general image-to-point cloud registration, including 2D3D-MatchNet [11] and our DeepI2P do not rely on another accurate localization system. Some other works [26, 4] focus on image-to-point cloud place recognition / retrieval without estimating the relative rotation and translation.

3. Overview of DeepI2P

We denote an image as $I \in \mathbb{R}^{3 \times W \times H}$, where W and H are the image width and height, and a point cloud as $P =$

$\{\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_N \mid \mathbf{P}_n \in \mathbb{R}^3\}$. The cross-modal image-to-point cloud registration problem is to solve for the rotation matrix $R \in \text{SO}(3)$ and translation vector $\mathbf{t} \in \mathbb{R}^3$ between the coordinate frames of the camera and point cloud. The problem is difficult because standard approaches such as ICP, PnP and Bundle Adjustment (BA) algorithms cannot be used due to the lack of point-to-pixel correspondences. Unlike the point cloud obtained from SfM, our point cloud is obtained from a point cloud scanner and does not contain any image-based feature descriptors. Establishing cross-modal point-to-pixel correspondence is non-trivial. This is because the points in the \mathbb{R}^3 space shares very little appearance and geometric correlations with the image in the \mathbb{P}^2 space. We circumvent the problem by designing our cross-modality image-to-point cloud registration approach to work without point-to-pixel correspondences.

To this end, we propose a two-stage ‘‘Frustum classification + Inverse camera projection’’ pipeline. The first stage classifies each point in the point cloud into within or beyond the camera frustum. We call this the frustum classification, which is done easily by a deep network shown in Section 4. In the second stage, we show that it is sufficient to solve the pose between camera and point cloud using only the frustum classification result. This is the inverse camera projection problem in Section 5.1. In our supplementary materials, we propose *another* cross-modality registration method ‘‘Grid classification + PnP’’ as our baseline for experimental comparison. In the grid classification, the image is divided into a tessellation of smaller regular grids, and we predict the cell each 3D point projects into. The pose estimation problem can then be solved by applying RANSAC-based PnP to the grid classification output.

4. Classification

The input to the network is a pair of image I and point cloud P , and the output is a per-point classification for P . There are two classification branches: frustum and grid classification. The frustum classification assign a label to each point, $L^c = \{l_1^c, l_2^c, \dots, l_N^c\}$, where $l_n^c \in \{0, 1\}$. $l_n^c = 0$ if the point \mathbf{P}_n is projected to outside the image I , and vice versa. Refer to the supplementary for the details of the grid classification branch used in our baseline.

4.1. Our Network Design

As shown in Fig. 2, our per-point classification network consists of four parts: point cloud encoder, point cloud decoder, image encoder and image-point cloud attention fusion. The point cloud encoder/decoder follows the design of SO-Net [20] and PointNet++ [28], while the image encoder is a ResNet-34 [17]. The classified points are then used in our inverse camera projection optimization in Section 5.1 to solve for the unknown camera pose.

Point Cloud Encoder. Given an input point cloud denoted as $P \in \mathbb{R}^{3 \times N}$, a set of nodes $\mathfrak{P}^{(1)} \in \mathbb{R}^{3 \times M_1}$ is sampled by Farthest Point Sampling (FPS). A point-to-node grouping [20] is performed to obtain M_1 clusters of points. Each cluster is processed by a PointNet [27] to get M_1 feature vectors of length C_1 , respectively, i.e. $P^{(1)} \in \mathbb{R}^{C_1 \times M_1}$. The point-to-node grouping is adaptive to the density of points. This is beneficial especially for point clouds from Lidar scans, where points are sparse at far range and dense at near range. The above sampling-grouping-PointNet operation is performed again to obtain another set of feature vectors $P^{(2)} \in \mathbb{R}^{C_2 \times M_2}$. Finally, a PointNet is applied to obtain the global point cloud feature vector $P^{(3)} \in \mathbb{R}^{C_3 \times 1}$.

Image-Point Cloud Attention Fusion. The goal of the classification is to determine whether a point projects to the image plane (frustum classification) and which region it falls into (grid classification). Hence, it is intuitive that the classification requires fusion of information from both modalities. To this end, we design an Attention Fusion module to combine the image and point cloud information. The input to the Attention Fusion module consists of three parts: a set of node features P_{att} ($P^{(1)}$ or $P^{(2)}$), a set of image features $I_{att} \in \mathbb{R}^{C_{img} \times H_{att} \times W_{att}}$ ($I^{(1)}$ or $I^{(2)}$), and the global image feature vector $I^{(3)}$. As shown in Fig. 2, the image global feature is stacked and concatenated with the node features P_{att} , and fed into a shared MLP to get the attention score $S_{att} \in \mathbb{R}^{H_{att} W_{att} \times M}$. S_{att} provides a weighting of the image features I_{att} for M nodes. The weighted image features are obtained by multiplying I_{att} and S_{att} . The weighted image features can now be concatenated with the node features in the point cloud decoder.

Point Cloud Decoder. The decoder takes the image and point cloud features as inputs, and outputs the per-point classification result. In general, it follows the interpolation idea of PointNet++ [28]. At the beginning of the decoder, the global image feature $I^{(3)}$ and global point cloud feature $P^{(3)}$ are stacked M_2 times, so that they can be concatenated with the node features $P^{(2)}$ and the Attention Fusion output $\tilde{I}^{(2)}$. The concatenated $[I^{(3)}, \tilde{I}^{(2)}, P^{(3)}, P^{(2)}]$ is processed by a shared MLP to get M_2 feature vectors denoted as $\tilde{P}^{(2)} \in \mathbb{R}^{C_2 \times M_2}$. We perform interpolation to get $\tilde{P}_{(itp)}^{(2)} \in \mathbb{R}^{C_2 \times M_1}$, where the M_2 features are upsampled to $M_1 \geq M_2$ features. Note that $P^{(2)}$ and $\tilde{P}^{(2)}$ are associated with node coordinates $\mathfrak{P}^{(2)} \in \mathbb{R}^{3 \times M_2}$. The interpolation is based on k -nearest neighbors between node coordinates $\mathfrak{P}^{(1)} \in \mathbb{R}^{3 \times M_1}$, where $M_1 \geq M_2$. For each C_2 channel, the interpolation is denoted as:

$$\tilde{P}_{(itp)_i}^{(2)} = \frac{\sum_{j=1}^k w_j \tilde{P}_j^{(2)}}{\sum_{j=1}^k w_j}, \text{ where } w_j = \frac{1}{d(\mathfrak{P}_i^{(1)}, \mathfrak{P}_j^{(2)})}, \quad (1)$$

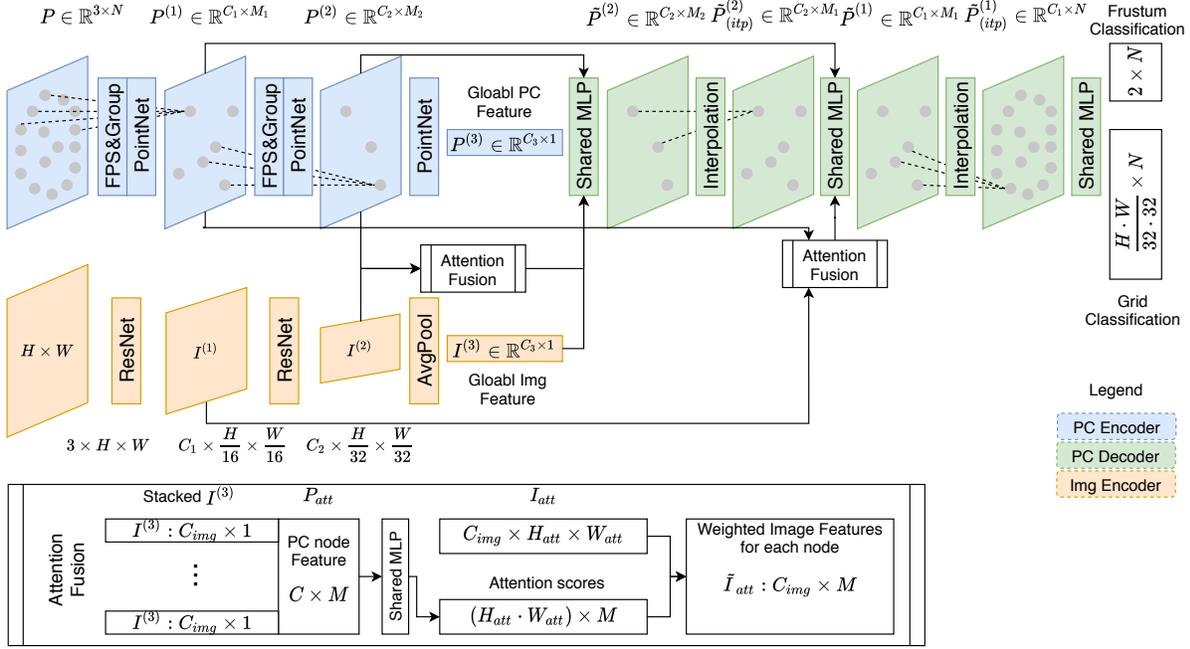


Figure 2. Our network architecture for the classification problem.

and $\mathfrak{P}_j^{(2)}$ is one of the k -nearest neighbors of $\mathfrak{P}_i^{(1)}$ in $\mathfrak{P}^{(2)}$. We get $\tilde{P}_{(itp)}^{(2)} \in \mathbb{R}^{C_2 \times M_1}$ with the concatenate-sharedMLP-interpolation process. Similarly, we obtain $\tilde{P}_{(itp)}^{(1)} \in \mathbb{R}^{C_1 \times M_1}$ after another round of operations. Lastly, we obtain the final output $(2 + HW/(32 \times 32)) \times N$, which can be reorganized into the frustum prediction scores $2 \times N$ and grid prediction scores $(HW/(32 \times 32)) \times N$.

4.2. Training Pipeline

The generation of the frustum labels is simply a camera projection problem. During training, we are given the camera intrinsic matrix $K \in \mathbb{R}^{3 \times 3}$ and the pose $G \in SE(4)$ between the camera and point cloud. The 3D transformation of a point $\mathbf{P}_i \in \mathbb{R}^3$ from the point cloud coordinate frame to the camera coordinate frame is given by:

$$\tilde{\mathbf{P}}'_i = [X'_i, Y'_i, Z'_i, 1]^\top = G\tilde{\mathbf{P}}_i = \begin{bmatrix} R & \mathbf{t} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{P}_i \\ 1 \end{bmatrix}, \quad (2)$$

and the point $\tilde{\mathbf{P}}'_i$ is projected into the image coordinate:

$$\tilde{\mathbf{p}}'_i = \begin{bmatrix} x'_i \\ y'_i \\ z'_i \end{bmatrix} = K\mathbf{P}'_i = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \mathbf{P}'_i. \quad (3)$$

Note that homogeneous coordinate is represented by a tilde symbol, e.g., $\tilde{\mathbf{P}}'_i$ is the homogeneous representation of \mathbf{P}'_i . The inhomogeneous coordinate of the image point is:

$$\mathbf{p}'_i = [p'_{x_i}, p'_{y_i}]^\top = [x'_i/z'_i, y'_i/z'_i]^\top. \quad (4)$$

Frustum Classification. For a given camera pose G , we define the function:

$$f(\mathbf{P}_i; G, K, H, W) = \begin{cases} 1 & : 0 \leq p'_{x_i} \leq W - 1, 0 \leq p'_{y_i} \leq H - 1, z'_i > 0 \\ 0 & : \text{otherwise} \end{cases}, \quad (5)$$

which assigns a label of 1 to a point \mathbf{P}_i that projects within the image, and 0 otherwise. Now the frustum classification labels are generated as $l_i^c = f(\mathbf{P}_i; G, K, H, W)$, where G is known during training. In the Oxford Robotcar and KITTI datasets, we randomly select a pair of image and raw point cloud (I, P_{raw}) , and compute the relative pose from the GPS/INS readings as the ground truth pose G_c^p . We use (I, P_{raw}) with a relative distance within a specified interval in our training data. However, we observe that the rotations in G_c^p are close to zero from the two datasets since the cars used to collect the data are mostly undergoing pure translations. To avoid overfitting to such scenario, we apply randomly generated rotations G_r onto the raw point cloud to get the final point cloud $P = G_r P_{raw}$ in the training data. Furthermore, the ground truth pose is now given by $G = G_c^p G_r^{-1}$. Note that random translation can also be included in G_r , but it does not have any effect on the training since the network is translational equivariant.

Training Procedure. The frustum classification training procedure is summarized as:

1. Select a pair of image and point cloud (I, P_{raw}) with relative pose G_c^p .

2. Generate 3D random transformation G_r , and apply it to get $P = G_r P_{raw}$ and $G = G_c^p G_r^{-1}$.
3. Get the ground truth per-point frustum labels l_i^c according to Eq. 5.
4. Feed (I, P) into the network illustrated in Fig. 2.
5. Frustum prediction $\hat{L}^c = \{\hat{l}_1^c, \dots, \hat{l}_n^c\}$, $\hat{l}_i^c \in \{0, 1\}$.
6. Apply cross entropy loss for the classification tasks to train the network.

5. Pose Optimization

We now formulate an optimization method to get the pose of the camera in the point cloud reference frame with the frustum classification results. Note that we do not use deep learning in this step since the physics and geometry of the camera projection model is already well-established.

Formally, the pose optimization problem is to solve for \hat{G} , given the point cloud P , frustum predictions $\hat{L}^c = \{\hat{l}_1^c, \dots, \hat{l}_N^c\}$, $\hat{l}_i^c \in \{0, 1\}$, and camera intrinsic matrix K . In this section, we describe our inverse camera projection solver to solve for \hat{G} .

5.1. Inverse Camera Projection

The frustum classification of a point, i.e., L^c given G defined in Eq. 5 is based on the forward projection of a camera. The inverse camera projection problem is the other way around, i.e., determine the optimal pose \hat{G} that satisfies a given \hat{L}^c . It can be written more formally as:

$$\hat{G} = \arg \max_{G \in SE(3)} \sum_{i=1}^N (f(\mathbf{P}_i; G, K, H, W) - 0.5) (\hat{l}_i^c - 0.5). \quad (6)$$

Intuitively, we seek to find the optimal pose \hat{G} such that all 3D points with label $\hat{l}_i^c = 1$ from the network are projected into the image, and vice versa. However, a naive search of the optimal pose in the $SE(3)$ space is intractable. To mitigate this problem, we relax the cost as a function of the distance from the projection of a point to the image boundary, i.e., a $H \times W$ rectangle.

Frustum Prediction Equals to 1. Let us consider a point \mathbf{P}_i with the prediction $\hat{l}_i^c = 1$. We define cost function:

$$g(p'_{x_i}; W) = \max(-p'_{x_i}, 0) + \max(p'_{x_i} - W, 0) \quad (7)$$

that penalizes a pose G which causes p'_{x_i} of the projected point $\mathbf{p}'_i = [p'_{x_i}, p'_{y_i}]^\top$ (c.f. Eq. 4) to fall outside the borders of the image width. Specifically, the cost is zero when p'_{x_i} is within the image width, and negatively proportional to the distance to the closest border along the image x-axis otherwise. A cost $g(p'_{y_i}; H)$ can be analogously defined along image y-axis. In addition, cost function $h(\cdot)$ is defined to avoid the ambiguity of \mathbf{P}'_i falling behind the camera:

$$h(z'_i) = \alpha \cdot \max(-z', 0), \quad (8)$$

where α is a hyper-parameter that balances the weighting between $g(\cdot)$ and $h(\cdot)$.

Frustum Prediction Equals to 0. We now consider a point \mathbf{P}_i with prediction $\hat{l}_i^c = 0$. The cost defined along the image x-axis is given by:

$$u(p'_{x_i}; W) = \frac{W}{2} - \left| p'_{x_i} - \frac{W}{2} \right|. \quad (9)$$

It is negative when p'_{x_i} falls outside the borders along the image width, and positively proportional to the distance to the closest border along the image x-axis otherwise. Similarly, an analogous cost $u(p'_{y_i}; H)$ along the y-axis can be defined. Furthermore, an indicator function:

$$\mathbb{1}(p'_{x_i}, p'_{y_i}, z'_i; H, W) = \frac{\max(u(p'_{x_i}; W), 0)}{u(p'_{x_i}; W)} \cdot \frac{\max(u(p'_{y_i}; H), 0)}{u(p'_{y_i}; H)} \cdot \frac{\max(z'_i, 0)}{z'_i} \quad (10)$$

is required to achieve the target of zero cost when \mathbf{p}'_i is outside the $H \times W$ image or \mathbf{P}'_i is behind the camera (i.e. $z'_i < 0$).

Cost Function. Finally, the cost function for a single point \mathbf{P}_i is given by:

$$r_i(G; \hat{l}_i^c) = \begin{cases} r_i^0 & : \hat{l}_i^c = 0 \\ r_i^1 & : \hat{l}_i^c = 1 \end{cases}, \quad \text{where} \quad (11)$$

$$\begin{aligned} r_i^0 &= (u(p'_{x_i}; W) + u(p'_{y_i}; H)) \cdot \mathbb{1}(p'_{x_i}, p'_{y_i}, z'_i; H, W), \\ r_i^1 &= g(p'_{x_i}; W) + g(p'_{y_i}; H) + h(z'_i). \end{aligned}$$

p'_{x_i}, p'_{y_i}, z_i are functions of G according to Eq. 2, 3 and 4. Image height H , width W and camera intrinsics K are known. Now the optimization problem in Eq. 6 becomes:

$$\hat{G} = \arg \min_{G \in SE(3)} \sum_{i=1}^n r_i(G; \hat{l}_i^c)^2. \quad (12)$$

This is a typical unconstrained least squares optimization problem. We need proper parameterization of the unknown transformation matrix,

$$G = \begin{bmatrix} R & \mathbf{t} \\ \mathbf{0} & 1 \end{bmatrix}, \quad \text{with } R \in SO(3), \mathbf{t} \in \mathbb{R}^3, \quad (13)$$

where $G \in SE(3)$ is an over-parameterization that can cause problems in the unconstrained continuous optimization. To this end, we use the Lie-algebra representation $\xi \in \mathfrak{se}(3)$ for the minimal parameterization of $G \in SE(3)$. The exponential map $G = \exp_{\mathfrak{se}(3)}(\xi)$ converts $\mathfrak{se}(3) \mapsto SE(3)$, while the log map $\xi = \log_{SE(3)}(G)$ converts $SE(3) \mapsto \mathfrak{se}(3)$. Similar to [10], we define the $\mathfrak{se}(3)$ concatenation operator $\circ : \mathfrak{se}(3) \times \mathfrak{se}(3) \mapsto \mathfrak{se}(3)$ as:

$$\xi_{ki} = \xi_{kj} \circ \xi_{ji} = \log_{SE(3)}(\exp_{\mathfrak{se}(3)}(\xi_{kj}) \cdot \exp_{\mathfrak{se}(3)}(\xi_{ji})), \quad (14)$$

and the cost function in Eq. 12 can be re-written with the proper exponential or log map modifications into:

$$\hat{G} = \exp_{\mathfrak{se}(3)}(\hat{\xi}), \quad \text{where} \quad (15)$$

$$\hat{\xi} = \arg \min_{\xi} \|\mathbf{r}\|^2 = \arg \min_{\xi} \sum_{i=1}^n r_i(\xi; \hat{l}_i)^2.$$

Gauss-Newton Optimization. Eq. 15 is a typical least squares optimization problem that can be solved by the Gauss-Newton method. During iteration i with the current solution $\xi^{(i)}$, the increment $\delta\xi^{(i)}$ is estimated by a Gauss-Newton second-order approximation:

$$\delta\xi^{(i)} = -(J^\top J)^{-1} J^\top \mathbf{r}(\xi^{(i)}), \quad \text{where } J = \left. \frac{\partial \mathbf{r}(\epsilon \circ \xi^{(i)})}{\partial \epsilon} \right|_{\epsilon=0}, \quad (16)$$

and the update is given by $\xi^{(i+1)} = \delta\xi^{(i)} \circ \xi^{(i)}$. Finally the inverse camera projection problem is solved by performing the exponential map $\hat{G} = \exp_{\mathfrak{se}(3)}(\hat{\xi})$. A visualization of our iterative optimization is presented in Fig. 3.

6. Experiments

Our image-to-point cloud registration approach is evaluated with Oxford Robotcar [24] and KITTI [13] dataset.

Oxford dataset. The point clouds are built from the accumulation of the 2D scans from a 2D Lidar. Each point cloud is set at the size of radius 50m, i.e. diameter 100m. The images are captured by the center camera of a Bumblebee tri-camera rig. Similar to 3DFeat-Net [40], 35 traversals are used for training, while 5 traversals are for testing. In training and inference, the image-point cloud pair is selected using the following steps: a) Choose a Lidar point cloud from one of the traversals. b) Randomly select an image from the same traversal and captured within $\pm 10\text{m}$ from the coordinate origin of the point cloud. The relative pose of the camera to the point cloud coordinate frame is G_c^p . c) Apply a random 2D rotation (around the up-axis) and translation (along the x-y plane) G_r to the point cloud. d) The objective is to recover the ground truth transformation $G_{gt} = G_c^p G_r^{-1}$. There are 130,078 point clouds for training and 19,156 for testing.

KITTI Odometry dataset. Point clouds are directly acquired from a 3D Lidar. The image-point cloud pairs are selected in a similar way to that in Oxford dataset, i.e., a pair of image and point cloud is captured within $\pm 10\text{m}$. The images are from both the left and right cameras that are facing the front. We follow the common practice of utilizing the 0-8 sequences for training, and 9-10 for testing. In total there are 20,409 point clouds for training, and 2,792 for testing.

Remarks: Note that a KITTI point cloud is from a single frame 3D Lidar scan, while an Oxford point cloud is an accumulation of 2D Lidar scans over 100m. As a result, point

clouds in KITTI suffers from severe occlusion, sparse measurement at far range, etc. More details of the two dataset configurations are in the supplementary materials.

6.1. Implementation Details

Classification Network. Refer to our supplementary materials for the network implementation details. These includes parameters of the PointNet and SharedMLP modules, number of nodes in $\mathfrak{P}^{(1)}$, $\mathfrak{P}^{(2)}$, number of nearest neighbor k for point cloud interpolation in the decoder, etc.

Inverse Camera Projection. The initial guess $G^{(0)}$ in our proposed inverse camera projection (c.f. Section 5.1) is critical since the solver for Eq. 15 is an iterative approach. To alleviate the initialization problem, we perform the optimization 60 times with randomly generated initialization $G^{(0)}$, and select the solution with the lowest cost. In addition, the 6DoF search space is too large for random initialization. We mitigate this problem by leveraging on the fact that our datasets are from ground vehicles to perform random initialization in 2D instead. Specifically, $R^{(0)}$ is initialized as a random rotation around the up-axis, and $\mathbf{t}^{(0)}$ as a random translation in the x-y horizontal plane. Our algorithm is implemented with Ceres [1].

6.2. Registration Accuracy

The frustum classification accuracy is 98% and 94% on the Oxford and KITTI dataset, respectively. However, these numbers does not translate directly to the registration accuracy. Following the practice of [21, 40], the registration is evaluated with two criteria: average Relative Translational Error (RTE) and average Relative Rotation Error (RRE). The results are shown in Table 1 and Fig. 4.

Grid Cls. + PnP is the result of our ‘‘Grid classification + PnP’’ baseline method (see Supplementary materials for details). The RANSAC PnP algorithm optimizes the full 6-DoF \hat{G} without any constraints. *Frus. Cls. + Inv.Proj.* represents the result of our ‘‘Frustum classification + Inverse camera projection’’ method. The difference between *Frus. Cls. + Inv.Proj. 3D* and *Frus. Cls. + Inv.Proj. 2D* is that the former is optimizing the full 6-DoF \hat{G} , while the latter constrains \hat{G} to be 3-DOF, i.e., translation on x-y horizontal plane and rotation around the up-axis.

Due to the lack of existing approaches in solving the image-to-point cloud registration problem under the same setting, we further compare our DeepI2P with 4 other approaches that may contain unfair advantages over ours in their input data modality or configurations.

1) Direct Regression uses a deep network to directly regress the relative poses. It consists of the Point Cloud Encoder and Image Encoder in Section 4. The global point cloud feature and global image feature are concatenated into

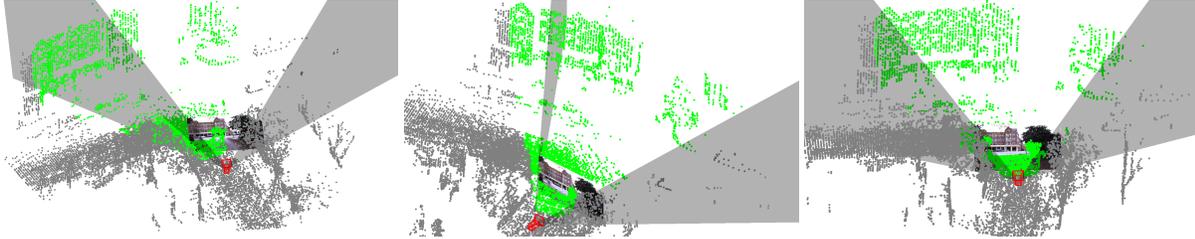


Figure 3. Visualizations of the Gauss-Newton at iteration 0 / 40 / 80 from left to right. Green points are classified as inside image FoV.

Table 1. Registration accuracy on the Oxford and KITTI datasets.

	Oxford		KITTI	
	RTE (m)	RRE ($^{\circ}$)	RTE (m)	RRE ($^{\circ}$)
Direct Regression	5.02 ± 2.89	10.45 ± 16.03	4.94 ± 2.87	21.98 ± 31.97
MonoDepth2 [14] + USIP [21]	33.2 ± 46.1	142.5 ± 139.5	30.4 ± 42.9	140.6 ± 157.8
MonoDepth2 [14] + GT-ICP	1.3 ± 1.5	6.4 ± 7.2	2.9 ± 2.5	12.4 ± 10.3
2D3D-MatchNet [11] (No Rot [§])	1.41	6.40	NA	NA
Grid Cls. + PnP	1.91 ± 1.56	5.94 ± 10.72	3.22 ± 3.58	10.15 ± 13.74
Frus. Cls. + Inv.Proj. 3D	2.27 ± 2.19	15.00 ± 13.64	3.17 ± 3.22	15.52 ± 12.73
Frus. Cls. + Inv.Proj. 2D	1.65 ± 1.36	4.14 ± 4.90	3.28 ± 3.09	7.56 ± 7.63

[§]Point clouds are not randomly rotated in the experiment setting of 2D3D-MatchNet [11].

a single vector and processed by a MLP that directly regresses \hat{G} . See the supplementary materials for more details of this method. Table 1 shows that our DeepI2P significantly outperforms the simple regression method.

2) Monodepth2+USIP converts the cross-modality registration problem into point cloud-based registration by using Monodepth2 [14] to estimate a depth map from a single image. The Lidar point cloud is used to calibrate the scale of depth map from Monodepth2, i.e. the scale of the depth map is perfect. Subsequently, the poses between the depth map and point cloud are estimated with USIP [21]. This is akin to same modality point cloud-to-point cloud registration. Nonetheless, Table 1 shows that this approach underperforms. This is probably because the depth map is inaccurate and USIP does not generalize well on depth maps.

3) Monodepth2+GT-ICP acquires a depth map with absolute scale in the same way as Monodepth2+USIP. However, it uses Iterative Closest Point (ICP) [2, 5] to estimate the pose between the depth map and point cloud. Note that ICP fails without proper initialization, and thus we use the ground truth (GT) relative pose for initialization. Table 1 shows that our DeepI2P achieves similar RTE and better RRE compared to Monodepth2+GT-ICP, despite the latter has the unfair advantages of ground truth initialization and the depth map is perfectly calibrated.

4) 2D3D-MatchNet [11] is the only prior work for cross-modal image-to-point cloud registration to our best knowledge. However, the rotation between camera and Lidar is almost zero in their experiment setting. This is because the images and point clouds are taken from temporally consecutive timestamps without additional augmentation. In con-

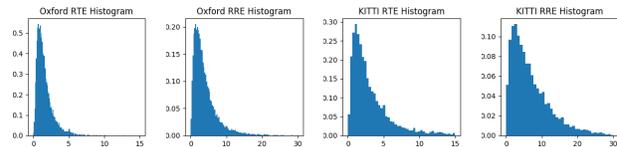


Figure 4. Histograms of image-point cloud registration RTE and RRE on the Oxford and KITTI datasets. x-axis is RTE (m) and RRE ($^{\circ}$), and y-axis is the percentage.

trast, the point clouds in our experiments are always randomly rotated. This means 2D3D-MatchNet is solving a much easier problem, but their results are worse than ours.

Distribution of Errors. The distribution of the registration RTE (m) and RRE ($^{\circ}$) on the Oxford and KITTI dataset are shown in Fig. 4. It can be seen that our performance is better on Oxford than KITTI. Specifically, the mode of the translational/rotational errors are $\sim 1.5\text{m}/3^{\circ}$ on Oxford and $\sim 2\text{m}/5^{\circ}$ on KITTI. The translational/rotational error variances are also smaller on Oxford.

Acceptance of RTE/RRE. There are other same-modality methods that solves registration on Oxford and KITTI dataset, e.g. USIP [21] and 3DFeatNet [40] that gives much better RTE and RRE. These methods work only on point cloud-to-point cloud instead of image-to-point cloud data, and thus are not directly comparable to our method. Furthermore, we note that the reported accuracy of our DeepI2P in Table 1 is sufficient for non life-critical applications such as frustum localization of mobile devices in both indoor and outdoor environments.

Oxford vs KITTI. Our performance on Oxford is better than on KITTI for several reasons: 1) The point clouds in

Oxford are built from 2D scans accumulated over 100m from a Lidar scanner, while KITTI point cloud is a single scan from a 3D Lidar as shown in Fig. 6. Consequently, the occlusion effect in KITTI is severe. For example, the image captured in timestamp t_j is seeing things that are mostly occluded/unobserved from the point cloud of timestamp t_i . Given the limited Field-of-View (FoV) of the camera, the cross-modality registration becomes extremely challenging since two modalities are observing different scene contents. 2) The vertical field-of-view of the point clouds in KITTI is very limited, which leads to lack of distinctive vertical structures for cross-modality matching. Most of the points on the ground are structureless and thus not useful for matching across modalities. 3) “KITTI Odometry” is a small dataset that contains only 20,409 point clouds for training, while Oxford is $\sim 6.4\times$ larger with 130,078 point clouds. As a result, we observe severe network overfitting in the KITTI dataset but not in the Oxford dataset.

6.3. Ablation Study

Initialization of our Gaussian-Newton and Monodepth2+ICP. In the 2D registration setting, there are 3 unknown parameters - rotation θ , translation t_x, t_y . For our method, the initial θ is easily obtained by aligning the average yaw-angles of the predicted in-frustum points with the camera principle axis. Therefore, our 60-fold initialization is only for two-dimensional search for t_x, t_y . In contrast, Monodepth2+ICP requires three-dimensional search for θ, t_x, t_y . As shown in Tab. 2, our DeepI2P is robust to initialization, while Monodepth2+ICP performs a lot worse with 60-fold random initialization.

In-image occlusion. In Oxford/Kitti, occlusion effect can be significant when the translation is large, e.g. $> 5m$. Tab. 2 shows that the registration improves when the maximum translation limit decrease ($15m \rightarrow 10m \rightarrow 5m$).

Cross-attention module. Tab. 2 shows the significant drop in registration accuracy w/o cross-attention module. Additionally, we also see the coarse classification accuracy drops significantly from 98% to 80%.

3D point density. As shown in Tab. 2, our registration accuracy decreases with reducing point density. Nonetheless, the performance drop is reasonable even when the point density drops to 1/4 ($20480 \rightarrow 5120$).

6.4. Visualizations

Fig. 5 shows examples of the results from our frustum classification network and the baseline grid classification network (see supplementary materials). The point clouds are projected into the images using the ground truth pose G . The colors of the points represent the correctness of the frustum or grid predictions as described in the caption. The

Table 2. Registration accuracy on Oxford

	# points	# init.	t limit	RTE	RRE
DeepI2P	20480	60	10	1.65	4.14
DeepI2P	20480	30	10	1.81	4.37
DeepI2P	20480	10	10	2.00	4.56
DeepI2P	20480	1	10	3.52	5.34
DeepI2P	20480	60	5	1.52	3.30
DeepI2P	20480	60	15	1.96	4.74
DeepI2P	10240	60	10	1.80	4.35
DeepI2P	5120	60	10	1.94	4.63
DeepI2P noAtten.	20480	60	10	6.88	20.93
MonoDepth2+ICP	20480	60	10	8.45	75.54

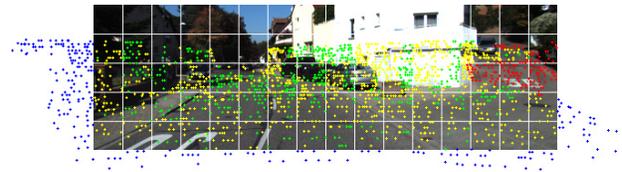


Figure 5. Visualization of the frustum and grid classification results projected onto the images. Green - both frustum and grid predictions are correct. Yellow - frustum prediction is correct but grid prediction is wrong. Red - frustum prediction is outside image FoV, but ground truth label is inside FoV. Blue - frustum prediction is inside image FoV, but ground truth label is outside FoV. Best view in color and zoom-in.

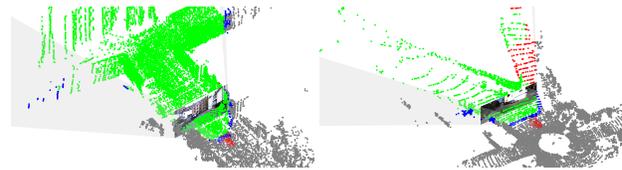


Figure 6. 3D Visualization of the frustum classification and inverse camera projection on the Oxford (Left) and KITTI (Right).

accuracy of the frustum and grid classifications are around 98% and 51% in Oxford, and 94% and 39% in KITTI. The low classification accuracy in KITTI leads to larger RTE and RRE during cross-modality registration. 3D visualization of the frustum classification and the inverse camera projection problem is illustrated in Fig. 6. It illustrates the intuition that a camera pose can be found by aligning the camera frustum to the classification result.

7. Conclusions

The paper proposes an approach for cross-modality registration between images and point clouds. The challenging registration problem is converted to a classification problem solved by deep networks and an inverse camera projection problem solved by least squares optimization. The feasibility of our proposed classification-optimization framework is verified with Oxford and KITTI dataset.

Acknowledgment. This work is supported in part by the Singapore MOE Tier 1 grant R-252-000-A65-114.

References

- [1] Sameer Agarwal, Keir Mierle, and Others. Ceres solver. <http://ceres-solver.org>.
- [2] 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.
- [3] Peter Biber and Wolfgang Straßer. The normal distributions transform: A new approach to laser scan matching. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, volume 3, pages 2743–2748. IEEE, 2003.
- [4] Daniele Cattaneo, Matteo Vaghi, Simone Fontana, Augusto Luis Ballardini, and Domenico Giorgio Sorrenti. Global visual localization in lidar-maps through shared 2d-3d embedding space. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4365–4371. IEEE, 2020.
- [5] Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. *Image and vision computing*, 10(3):145–155, 1992.
- [6] Haowen Deng, Tolga Birdal, and Slobodan Ilic. Ppf-foldnet: Unsupervised learning of rotation invariant 3d local descriptors. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 602–618, 2018.
- [7] Haowen Deng, Tolga Birdal, and Slobodan Ilic. Ppfnet: Global context aware local features for robust 3d point matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 195–205, 2018.
- [8] Chitra Dorai and Anil K. Jain. Cosmos-a representation scheme for 3d free-form objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(10):1115–1130, 1997.
- [9] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence*, 40(3):611–625, 2017.
- [10] Jakob Engel, Thomas Schöps, and Daniel Cremers. Lsdslam: Large-scale direct monocular slam. In *European conference on computer vision*, pages 834–849. Springer, 2014.
- [11] Mengdan Feng, Sixing Hu, Marcelo H Ang, and Gim Hee Lee. 2d3d-matchnet: Learning to match keypoints across 2d image and 3d point cloud. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 4790–4796. IEEE, 2019.
- [12] 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.
- [13] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- [14] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 3828–3838, 2019.
- [15] Zan Gojcic, Caifa Zhou, Jan D Wegner, and Andreas Wieser. The perfect match: 3d point cloud matching with smoothed densities. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5545–5554, 2019.
- [16] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [18] Wolfgang Hess, Damon Kohler, Holger Rapp, and Daniel Andor. Real-time loop closure in 2d lidar slam. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1271–1278. IEEE, 2016.
- [19] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnnp: An accurate o(n) solution to the pnp problem. *International journal of computer vision*, 81(2):155, 2009.
- [20] Jiaxin Li, Ben M Chen, and Gim Hee Lee. So-net: Self-organizing network for point cloud analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9397–9406, 2018.
- [21] Jiaxin Li and Gim Hee Lee. Usip: Unsupervised stable interest point detection from 3d point clouds. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 361–370, 2019.
- [22] David G Lowe. Object recognition from local scale-invariant features. In *Proceedings of the seventh IEEE international conference on computer vision*, volume 2, pages 1150–1157. Ieee, 1999.
- [23] Weixin Lu, Guowei Wan, Yao Zhou, Xiangyu Fu, Pengfei Yuan, and Shiyu Song. Deepicp: An end-to-end deep neural network for 3d point cloud registration. *arXiv preprint arXiv:1905.04153*, 2019.
- [24] Will Maddern, Geoffrey Pascoe, Chris Linegar, and Paul Newman. 1 year, 1000 km: The oxford robotcar dataset. *The International Journal of Robotics Research*, 36(1):3–15, 2017.
- [25] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics*, 31(5):1147–1163, 2015.
- [26] Quang-Hieu Pham, Mikaela Angelina Uy, Binh-Son Hua, Duc Thanh Nguyen, Gemma Roig, and Sai-Kit Yeung. Lcd: learned cross-domain descriptors for 2d-3d matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11856–11864, 2020.
- [27] 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.
- [28] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in neural information processing systems*, pages 5099–5108, 2017.

- [29] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *2011 International conference on computer vision*, pages 2564–2571. Ieee, 2011.
- [30] Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast point feature histograms (fpfh) for 3d registration. In *2009 IEEE international conference on robotics and automation*, pages 3212–3217. IEEE, 2009.
- [31] Radu Bogdan Rusu and Steve Cousins. 3d is here: Point cloud library (pcl). In *2011 IEEE international conference on robotics and automation*, pages 1–4. IEEE, 2011.
- [32] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Improving image-based localization by active correspondence search. In *European conference on computer vision*, pages 752–765. Springer, 2012.
- [33] Yoli Shavit and Ron Ferens. Introduction to camera pose estimation with deep learning. *arXiv preprint arXiv:1907.05272*, 2019.
- [34] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In *European conference on computer vision*, pages 356–369. Springer, 2010.
- [35] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Performance evaluation of 3d keypoint detectors. *International Journal of Computer Vision*, 102(1-3):198–220, 2013.
- [36] Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. Bundle adjustment—a modern synthesis. In *International workshop on vision algorithms*, pages 298–372. Springer, 1999.
- [37] Shimon Ullman. The interpretation of structure from motion. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 203(1153):405–426, 1979.
- [38] Yue Wang and Justin M Solomon. Deep closest point: Learning representations for point cloud registration. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3523–3532, 2019.
- [39] Jiaolong Yang, Hongdong Li, and Yunde Jia. Go-icp: Solving 3d registration efficiently and globally optimally. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1457–1464, 2013.
- [40] Zi Jian Yew and Gim Hee Lee. 3dfeat-net: Weakly supervised local 3d features for point cloud registration. In *European Conference on Computer Vision*, pages 630–646. Springer, 2018.
- [41] Zi Jian Yew and Gim Hee Lee. Rpm-net: Robust point matching using learned features. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [42] Huai Yu, Weikun Zhen, Wen Yang, Ji Zhang, and Sebastian Scherer. Monocular camera localization in prior lidar maps with 2d-3d line correspondences. *arXiv preprint arXiv:2004.00740*, 2020.
- [43] Andy Zeng, Shuran Song, Matthias Nießner, Matthew Fisher, Jianxiong Xiao, and Thomas Funkhouser. 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1802–1811, 2017.
- [44] Ji Zhang and Sanjiv Singh. Loam: Lidar odometry and mapping in real-time. In *Robotics: Science and Systems*, volume 2, 2014.
- [45] Yu Zhong. Intrinsic shape signatures: A shape descriptor for 3d object recognition. In *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*, pages 689–696. IEEE, 2009.