



# Globally Optimal Point Set Registration by Joint Symmetry Plane Fitting

Lan Hu<sup>1,2</sup> · Laurent Kneip<sup>1</sup>

Received: 19 June 2020 / Accepted: 28 January 2021 / Published online: 26 February 2021  
© The Author(s) 2021

## Abstract

The present work proposes a solution to the challenging problem of registering two partial point sets of the same object with very limited overlap. We leverage the fact that most objects found in man-made environments contain a plane of symmetry. By reflecting the points of each set with respect to the plane of symmetry, we can largely increase the overlap between the sets and therefore boost the registration process. However, prior knowledge about the plane of symmetry is generally unavailable or at least very hard to find, especially with limited partial views. Finding this plane could strongly benefit from a prior alignment of the partial point sets. We solve this chicken-and-egg problem by jointly optimizing the relative pose and symmetry plane parameters. We present a globally optimal solver by employing the branch-and-bound paradigm and thereby demonstrate that joint symmetry plane fitting leads to a great improvement over the current state of the art in globally optimal point set registration for common objects. We conclude with an interesting application of our method to dense 3D reconstruction of scenes with repetitive objects.

**Keywords** Shape registration · Symmetry plane estimation · Global optimization

## 1 Introduction

The alignment of two point sets is a fundamental geometric problem that occurs in many computer vision and robotics applications. In computer vision, the technique is used to stitch together partial 3D reconstructions in order to form a more complete model of an object or environment [30]. In robotics, point set registration is an essential ingredient to simultaneous localization and mapping with affordable consumer depth cameras [29] or powerful 3D laser range scanners [11]. The general approach for aligning two point sets does not require initial correspondences. It is given by the iterative closest point (ICP) algorithm [47], a local search strategy that alternates between geometric correspondence establishment (i.e. by simple nearest neighbour search) and

Procrustes alignment. The iterative procedure depends on a sufficiently accurate initial guess about the relative transformation (e.g. an identity transformation in the case of incremental ego-motion estimation).

The present work is motivated by a common problem that occurs when performing a dense reconstruction of an environment which contains multiple instances of the same object. An example of the latter is given by a room in which the same type of chair occurs more than once. Let us assume that the front end of our reconstruction framework encompasses semantic recognition capabilities which are used to segment out partial point sets of objects of the same class and type [24]. There is a general interest in aligning those partial object point sets towards exploiting their mutual information and completing or even improving the reconstruction of each instance. The difficulty of this partial point set registration problem arises from two factors:

- The relative pose between different objects is arbitrary and unknown.
- Since the objects are observed under an arbitrary pose and with potential occlusions, the measured partial point sets potentially have very little overlap.

---

✉ Lan Hu  
hulan@shanghaitech.edu.cn

Laurent Kneip  
lkneip@shanghaitech.edu.cn

<sup>1</sup> Mobile Perception Lab, ShanghaiTech University, Shanghai, China

<sup>2</sup> Shanghai Institute of Microsystem and Information Technology, Chinese Academy of Sciences, Shanghai, China

Our contribution focuses on the registration of two partial point sets with limited overlap. The plain ICP algorithm is only a local search strategy that depends on a sufficiently accurate initial guess, which is why it may not serve as a valid solution to our problem. A potential remedy is given by the globally optimal ICP algorithm presented by Yang et al. [42]. However, the algorithm still depends on sufficient overlap in the partial point sets, which is not necessarily a given (50% is reported as a requirement for high success rate).

The core idea of the present work consists of exploiting the fact that the majority of commonly observed objects contain a plane of symmetry. By reflecting the points of each partial point set with respect to the plane of symmetry, we may effectively increase the overlap between the two sets and vastly improve the success rate of the registration process. However, given that each point set only observes part of the object, the identification of the plane of symmetry in each individual point set appears to be an equally difficult and ill-posed problem than the partial point set registration itself. It is only after the aligning transformation is found that symmetry plane detection would become a less challenging problem. In conclusion, the solution of each problem strongly depends on a prior solution to the other. We therefore present the following contributions:

- We solve this chicken-and-egg problem by a joint solution of the aligning transformation as well as the symmetry plane parameters. To the best of our knowledge, our work is the first to address those two problems jointly, thus leading to a vast improvement over the existing state-of-the-art in globally optimal point set registration, especially in situations in which the two point sets contain very limited overlap.
- We devise a globally optimal solution to this problem by employing the branch-and-bound optimization paradigm. Our work implicitly provides the first solution to globally optimal symmetry plane estimation in a single point set, or—more generally—symmetry plane detection across two point sets.
- We test our algorithm on synthetic 2d and 3d data, as well as open model and real data. We furthermore show a meaningful application of our algorithm to a dense 3D reconstruction scenario in which multiple instances of the same object occur.

## 2 Related Work

Despite strong mutual dependency, our method is the first to perform joint point set alignment and symmetry plane estimation. Our literature review therefore cites prior art on those two topics individually.

### 2.1 Point-Set Registration

The iterative closest point (ICP) [5,12,47] algorithm is a popular method for aligning two point sets. It does not depend on a prior derivation of point-to-point correspondences and simply aligns the two sets by iteratively alternating between the two steps of finding nearest neighbours (e.g. by evaluating point-to-point distances) and computing the alignment (e.g. using Arun's method [2]). To improve robustness of the algorithm against occlusions and reduced overlap, the method has been extended by outlier rejection [19,47] or data trimming [13] techniques. However, the classical ICP algorithm remains a local search algorithm for which the convergence depends on a good initial guess and sufficient overlap between both point sets.

The approach to global ICP registration is based on the branch-and-bound framework [16,18,21,43]. Those algorithms avoid local minima by searching the entire space of relative transformations. Yang et al. [42,43] propose the Go-ICP algorithm, which applies branch and bound to the ICP problem to find the globally optimal minimum of the sum of L2 distances between nearest neighbours from two aligned point sets. The method is accelerated by using local ICP in the loop. However, missing robustness of the cost-function causes the method to remain sensitive with respect to occlusions and partial overlaps. Campbell et al. [8] finally devise GOGMA, a branch-and-bound variant in which the objective of minimizing point-to-point errors is again replaced by Gaussian mixture alignment.

An entire family of alternative approaches [41,45,48] relies on the idea of feature matching. Some pipelines use point-to-point matches based on local geometric descriptors [20,33,37]. Once candidate correspondences are collected, the alignment is estimated iteratively from sparse subsets of correspondences and then validated on the entire surface. This iterative process is typically based on variants of the RANdom SAMple Consensus (RANSAC) scheme [7,10,17]. However, when the surfaces only partially overlap, existing pipelines often require many iterations to sample a good correspondence set and find a reasonable alignment. In contrast with RANSAC-based algorithms, another line of work combines non-convex optimization algorithms [40,41,48] with data truncation [26] or robust kernels and iterative reweighting [4,19]. Limited overlap ratios nonetheless render the problem ill-conditioned and very challenging to be solved.

Limited overlap ratios cause lots of mismatches in the point-pair establishment. Most existing techniques aim at rejecting such mismatches by robust registration. In contrast, the idea of the present work is to decrease the number of mismatches by increasing the overlap ratio through joint symmetry plane fitting. This leads to a large improvement in partial scan alignment for common symmetric objects. Inspired by recent advances [8,42], we also employ the nested

branch-and-bound strategy integrated with local ICP to find globally optimal alignments and thus verify our formulation.

## 2.2 Symmetry Plane Estimation

In recent years, the exploitation of symmetry in shapes has become a popular topic in the fields of 3D vision and deep learning. It has been used in various works on 3D object completion [32,35,36], shape correspondence prediction [3,28,46] and 3D object detection [1,6,31]. For a comprehensive review of symmetry detection, we kindly refer the reader to Liu et al.'s review [23] and its applications [27]. Here we only focus on the problem of symmetry plane fitting with missing data. The most straightforward solution is given by employing the RANSAC algorithm proposed by Fischler and Bolles [17], a well-known algorithm for robust model fitting for outlier affected data. In the context of shape matching, the basic idea is to extract sparse characteristic points and match them between both sets. We then choose a random subset of correspondences and derive a hypothesis for the global transformation induced by these samples. The alignment quality is finally evaluated by the matching error between the two shapes. The method can be easily applied for detecting a plane of symmetry in a single point set by employing symmetry invariant point descriptors and hypothesizing the plane of symmetry to be orthogonal to the axis connecting a correspondence. For example, Cohen et al. [15] detect symmetries in sparse point clouds by using appearance-based 3D-3D point correspondences in a RANSAC scheme. The detected symmetries are subsequently explored to eliminate noise from the point-clouds. Xu et al. [39] present a voting algorithm to detect the intrinsic reflectional symmetry axis. Using the axis as a hint, a completion algorithm for missing geometry is again shown. Jiang et al. [22] on the other hand propose an algorithm to find intrinsic symmetries in point clouds by using a curve skeleton. A set of filters then produces a candidate set of symmetric correspondences which are finally verified via spectral analysis. Although these works show results on partial data, the amount of missing data is typically small. Inspired by the work of [14] which detects symmetry by registration, we propose to detect the symmetry plane alongside partial point set registration, thus leading to improved performance in situations with limited overlap.

## 3 Preliminaries

We start by introducing the notation used throughout the paper and review the basic formulation of the ICP problem as well as planar reflections.

### 3.1 Notations and Assumptions

Let us denote the two partial object point sets by  $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^M$  and  $\mathcal{Y} = \{\mathbf{y}_i\}_{i=1}^N$  (sometimes called the *model* and *data* point sets, respectively). The goal pursued in this paper is the identification of a Euclidean transformation given by the rotation  $\mathbf{R}$  and the translation  $\mathbf{t}$  that transforms the points of  $\mathcal{Y}$  such that they *align* with the points of  $\mathcal{X}$ . If  $\mathcal{X}$  and  $\mathcal{Y}$  contain points in the 2D plane,  $\mathbf{R}$  and  $\mathbf{t}$  form an element of the group  $SE(2)$ . If  $\mathcal{X}$  and  $\mathcal{Y}$  contain 3D points,  $\mathbf{R}$  and  $\mathbf{t}$  will be an element of  $SE(3)$ . Note that *alignment* denotes a more general idea rather than just the minimization of the sum of distances between each point of  $\mathcal{Y}$  and its closest point within  $\mathcal{X}$ . The point sets have different cardinality and potentially observe very different parts of the object with only very little overlap. This motivates our approach that takes object symmetries into account.

### 3.2 Registration of Two Point Sets

The standard solution to the point set registration problem is given by the ICP algorithm [47], which minimizes the alignment error given by

$$E(\mathbf{R}, \mathbf{t}) = \sum_{i=1}^N e_i^r(\mathbf{R}, \mathbf{t}|\mathbf{y}_i)^2 = \sum_{i=1}^N \|\mathbf{R}\mathbf{y}_i + \mathbf{t} - \mathbf{x}_j\|^2, \quad (1)$$

where  $e_i^r(\mathbf{R}, \mathbf{t}|\mathbf{y}_i)$  is the per-point residual error for  $\mathbf{y}_i$ , and  $\mathbf{x}_j$  is the closest point to  $\mathbf{y}_i$  within  $\mathcal{X}$ , i.e.

$$\mathbf{x}_j = \operatorname{argmin}_{\mathbf{x} \in \mathcal{X}} \|\mathbf{R}\mathbf{y}_i + \mathbf{t} - \mathbf{x}\|. \quad (2)$$

Given an initial transformation  $\mathbf{R}$  and  $\mathbf{t}$ , the ICP algorithm iteratively solves the above minimization problem by alternating between updating the aligning transformation with fixed  $\mathbf{x}_j$  (i.e. using (1)) and updating the closest point matches  $\mathbf{x}_j$  themselves using (2). It is intuitively clear that the ICP algorithm only converges to a local minimum.

### 3.3 Modelling and Identifying Symmetry

Symmetry is modelled by a reflection by the symmetry plane. Let us define the symmetry plane by the normal  $\mathbf{n}$  and depth of plane  $d$  such that any point  $\mathbf{x}$  on the plane satisfies the relation  $\mathbf{x}^T \mathbf{n} + d = 0$ . Let  $\mathbf{x}$  now be a (2D or 3D) point from  $\mathcal{X}$ . The reflection plane reflects  $\mathbf{x}$  to a single reflected point  $\hat{\mathbf{x}}$  given by

$$\hat{\mathbf{x}} = \mathbf{x} - 2\mathbf{n}(\mathbf{x}^T \mathbf{n} + d). \quad (3)$$

The term in parentheses is the signed distance between  $\mathbf{x}$  and the reflection plane. The subtraction of  $2\mathbf{n}$  times this

distance reflects the point to the other side of the plane. The problem of symmetry identification may now be formulated as a minimization of the symmetry distance defined by

$$E(\mathbf{n}, d) = \sum_{i=1}^M e_i^s(\mathbf{n}, d|\mathbf{x}_i)^2 = \sum_{i=1}^M \|\hat{\mathbf{x}}_i - \mathbf{x}_j\|^2, \tag{4}$$

where  $e_i^s(\mathbf{n}, d|\mathbf{x}_i)$  is the per-point residual error for  $\mathbf{x}_i$ , and  $\mathbf{x}_j$  is the nearest point to  $\hat{\mathbf{x}}_i$  in  $\mathcal{X}$ , i.e.

$$\mathbf{x}_j = \underset{\mathbf{x} \in \mathcal{X}}{\operatorname{argmin}} \|\hat{\mathbf{x}}_i - \mathbf{x}\|. \tag{5}$$

It is intuitively clear that the symmetry plane fitting problem may also be solved via ICP, the only difference being the parameters over which the problem is solved (i.e.  $\mathbf{n}$  and  $d$ ).

### 3.4 Transformed Symmetry Plane Parameters

Let us still assume that  $\mathbf{n}$  and  $d$  are the symmetry plane parameters of a point set  $\mathcal{X}$ , and  $\mathbf{R}$  and  $\mathbf{t}$  are the parameters that align a point set  $\mathcal{Y}$  with  $\mathcal{X}$ . Each transformed point  $\mathbf{R}\mathbf{y} + \mathbf{t}$  and its transformed, symmetric equivalent  $\mathbf{R}\hat{\mathbf{y}} + \mathbf{t}$  must still fulfill the original reflection equation (3):

$$\mathbf{R}\hat{\mathbf{y}} + \mathbf{t} = \mathbf{R}\mathbf{y} + \mathbf{t} - 2\mathbf{n}((\mathbf{R}\mathbf{y} + \mathbf{t})^T \mathbf{n} + d). \tag{6}$$

Cancelling  $\mathbf{t}$  and multiplying by  $\mathbf{R}^T$  on either side, we easily obtain

$$\hat{\mathbf{y}} = \mathbf{y} - 2\mathbf{R}^T \mathbf{n}(\mathbf{y}^T \mathbf{R}^T \mathbf{n} + \mathbf{t}^T \mathbf{n} + d). \tag{7}$$

By comparing to (3), it is obvious that  $\hat{\mathbf{n}} = \mathbf{R}^T \mathbf{n}$  and  $\hat{d} = \mathbf{t}^T \mathbf{n} + d$  must represent the symmetry plane parameters for the original, untransformed set  $\mathcal{Y}$ .

## 4 Alignment and Symmetry as a Joint Optimization Problem

We now introduce our novel optimization objective which jointly optimizes an aligning point set transformation as well as the plane of symmetry. The objective is then solved in a branch-and-bound optimization paradigm, for which we introduce both the domain parameterization and the derivation of upper and lower bounds.

### 4.1 Objective Function

We still assume that our two partial object point sets are given by  $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^M$  and  $\mathcal{Y} = \{\mathbf{y}_i\}_{i=1}^N$ , and that the symmetry plane is represented by the normal  $\mathbf{n}$  and depth  $d$ . We define  $\mathcal{X}^s = \{\mathbf{x}_i^s | \mathbf{x}_i^s = \mathbf{x}_i - 2\mathbf{n}(\mathbf{x}_i^T \mathbf{n} + d), i = 1, \dots, M\}$  to be

the corresponding reflected point set of  $\mathcal{X}$ . We, furthermore, define  $\mathcal{X}^r = \{\mathbf{x}_i^r | \mathbf{x}_i^r = \mathbf{R}^T(\mathbf{x}_i - \mathbf{t}), i = 1, \dots, M\}$  and  $\mathcal{Y}^r = \{\mathbf{y}_i^r | \mathbf{y}_i^r = \mathbf{R}\mathbf{y}_i + \mathbf{t}, i = 1, \dots, N\}$  to be the aligned sets in either direction. The symmetry fitting objective of  $\mathcal{X}$  employs

$$e_i^s(\mathbf{n}, d|\mathbf{x}_i) = \|\mathbf{x}_i - 2\mathbf{n}(\mathbf{x}_i^T \mathbf{n} + d) - \mathbf{x}_j\|, \tag{8}$$

where the difference to (4) is given by the fact that  $\mathbf{x}_j$  is now the nearest neighbour from the set  $\mathcal{X} \cup \mathcal{Y}^r$ . Similarly, using Eq. (7), the symmetry objective function for  $\mathcal{Y}$  employs

$$e_i^s(\hat{\mathbf{n}}, \hat{d}|\mathbf{y}_i) = \|\mathbf{y}_i - 2\hat{\mathbf{n}}(\mathbf{y}_i^T \hat{\mathbf{n}} + \hat{d}) - \mathbf{y}_j\|, \tag{9}$$

where  $\mathbf{y}_j$  is now chosen as the nearest neighbour from the set  $\mathcal{X}^r \cup \mathcal{Y}$ .

The registration error adopts the formulation

$$e_i^r(\mathbf{R}, \mathbf{t}|\mathbf{y}_i) = \begin{cases} \|\mathbf{R}\mathbf{y}_i + \mathbf{t} - \mathbf{x}_j\| & \text{if } \|\mathbf{R}\mathbf{y}_i + \mathbf{t} - \mathbf{x}_j\| < c \\ w_1 \|\mathbf{R}\mathbf{y}_i + \mathbf{t} - \mathbf{x}_j^s\| & \text{else} \end{cases} \tag{10}$$

where  $\mathbf{x}_j$  and  $\mathbf{x}_j^s$  are chosen as the nearest neighbours from the sets  $\mathcal{X}$  and  $\mathcal{X}^s$ , respectively. We give priority to align points to the directly registered points from  $\mathcal{X}$  and transfer to the alignment with a point from  $\mathcal{X}^s$  if the former is larger than a chosen design parameter  $c$ . Lastly, residual errors that are very large are simply truncated from the problem.

The overall objective function becomes

$$E(\mathbf{R}, \mathbf{t}, \mathbf{n}, d) = \sum_{i=1}^N e_i^r(\mathbf{R}, \mathbf{t}|\mathbf{y}_i)^2 + w_2 \left\{ \sum_{i=1}^M e_i^s(\mathbf{n}, d|\mathbf{x}_i)^2 + \sum_{i=1}^N e_i^s(\hat{\mathbf{n}}, \hat{d}|\mathbf{y}_i)^2 \right\}, \tag{11}$$

where  $w_2$  balances the registration and symmetry objectives. Direct optimization over  $\mathbf{R}$ ,  $\mathbf{t}$ ,  $\mathbf{n}$ , and  $d$  using traditional ICP would easily get trapped in the nearest local minimum. We therefore propose to minimize the energy objective using the globally optimal branch-and-bound paradigm, an exhaustive search strategy that branches over the entire parameter space. In order to speed up the execution, the method derives upper and lower bounds for the minimal energy on each branch (i.e. sub-volume of the optimization space) and discards branches for which the lower bound remains higher than the upper bound in another branch. We use the axis-angle vector parametrization given by  $\mathbf{r}$  and  $\alpha$  to represent the relative rotation  $\mathbf{R}$  and the symmetry normal  $\mathbf{n}$ , respectively. In particular,  $\mathbf{n}(\alpha) = \mathbf{R}_\alpha [1 \ 0 \ 0]^T$ . In the remainder of this section, we will discuss how to find concrete values for the upper and lower bounds.

### 4.2 Derivation of the Upper and Lower Bounds

The basic idea of BnB is to partition the feasible set into convex sets on which we derive upper and lower bounds for the objective  $L_2$ -energy. Given that our objective (11) consists of a sum of many squared and positive sub-energies, lower and upper bounds on the overall objective energy can be derived by calculating individual upper and lower bounds on the alignment and symmetry residuals. Using the above introduced parametrizations, the upper bound  $\bar{E}$  and the lower bound  $\underline{E}$  of the optimal, joint  $L_2$  registration and symmetry cost  $E^*$  on a given interval of variables are therefore given as

$$\begin{aligned} \bar{E} &\doteq \sum_{i=1}^N \overline{e_i^r}(\mathbf{r}, \mathbf{t}|\mathbf{y}_i)^2 + w_2 \left\{ \sum_{i=1}^M \overline{e_i^s}(\boldsymbol{\alpha}, d|\mathbf{x}_i)^2 \right. \\ &\quad \left. + \sum_{i=1}^N \overline{e_i^s}(\mathbf{r}, \mathbf{t}, \boldsymbol{\alpha}, d|\mathbf{y}_i)^2 \right\} \\ \underline{E} &\doteq \sum_{i=1}^N \underline{e_i^r}(\mathbf{r}, \mathbf{t}|\mathbf{y}_i)^2 + w_2 \left\{ \sum_{i=1}^M \underline{e_i^s}(\boldsymbol{\alpha}, d|\mathbf{x}_i)^2 \right. \\ &\quad \left. + \sum_{i=1}^N \underline{e_i^s}(\mathbf{r}, \mathbf{t}, \boldsymbol{\alpha}, d|\mathbf{y}_i)^2 \right\}. \end{aligned}$$

Upper bounds on an interval are easily given by the energy of an arbitrary point within the interval. The remainder of this section discusses the derivation of the lower bounds.

*Lower bound for the alignment error  $e_i^r(\mathbf{r}, \mathbf{t}|\mathbf{y}_i)$ :* For a rotation interval of half-length  $\sigma_r$  with centre  $\mathbf{r}_0$ , we have

$$\|\mathbf{R}_r \mathbf{y} - \mathbf{R}_{r_0} \mathbf{y}\| \leq 2\sin(\min(\sigma_r/2, \pi/2))\|\mathbf{y}\| \doteq \gamma_r \|\mathbf{y}\|. \quad (12)$$

$\gamma_r$  is also called the rotation uncertainty radius.

**Proof** Using Lemmas 3.1 and 3.2 of [21], we have

$$\angle(\mathbf{R}_r \mathbf{y}, \mathbf{R}_{r_0} \mathbf{y}) \leq d_{\angle}(\mathbf{R}_r, \mathbf{R}_{r_0}) \leq \|\mathbf{r} - \mathbf{r}_0\| \quad (13)$$

and

$$\begin{aligned} \|\mathbf{R}_r \mathbf{y} - \mathbf{R}_{r_0} \mathbf{y}\| &= 2\sin(\angle(\mathbf{R}_r \mathbf{y}, \mathbf{R}_{r_0} \mathbf{y})/2)\|\mathbf{y}\| \\ &\leq 2\sin(\min(d_{\angle}(\mathbf{R}_r, \mathbf{R}_{r_0}), \pi)/2)\|\mathbf{y}\| \\ &\leq 2\sin(\min(\|\mathbf{r} - \mathbf{r}_0\|, \pi)/2)\|\mathbf{y}\| \\ &\leq 2\sin(\min(\sigma_r/2, \pi/2))\|\mathbf{y}\|. \text{ (end of proof)} \end{aligned} \quad (14)$$

We can similarly derive a translation uncertainty radius  $\gamma_t$ . For a translation volume with half side length  $\sigma_t$  centred at  $\mathbf{t}_0$ , we have

$$\|(\mathbf{x} + \mathbf{t}) - (\mathbf{x} + \mathbf{t}_0)\| = \|\mathbf{t} - \mathbf{t}_0\| \leq \|\sigma_t\| \doteq \gamma_t. \quad (15)$$

The registration term in Eq. (10) becomes

$$e_i^r(\mathbf{r}, \mathbf{t}|\mathbf{y}_i) \geq e_i^r(\mathbf{r}_0, \mathbf{t}_0|\mathbf{y}_i) - w(x)(\gamma_r \|\mathbf{y}_i\| + \gamma_t),$$

where  $w(\mathbf{x}) = 1$  if correspondence  $\mathbf{x} \in \mathcal{X}$ , or  $w(\mathbf{x}) = w_2$  if  $\mathbf{x} \in \mathcal{X}^s$ . The final lower bound of the registration term is

$$\underline{e}_i^r(\mathbf{r}, \mathbf{t}|\mathbf{y}_i) \doteq \max\{e_i^r(\mathbf{r}_0, \mathbf{t}_0|\mathbf{y}_i) - (\gamma_r \|\mathbf{y}_i\| + \gamma_t), 0\} \quad (16)$$

For more details, we kindly refer the reader to [21,42].  $\square$

*Lower bound of symmetry term  $e_i^s(\boldsymbol{\alpha}, d|\mathbf{x}_i)$ :* Assuming that the normal is defined by an  $\boldsymbol{\alpha}$ -interval of half-length  $\sigma_\alpha$  and with centre  $\boldsymbol{\alpha}_0$ , we have

$$\begin{aligned} \|\mathbf{x}^T(\mathbf{n} - \mathbf{n}_0)\| &\leq \|\mathbf{n} - \mathbf{n}_0\| \|\mathbf{x}\| = \|(\mathbf{R}_\alpha - \mathbf{R}_{\alpha_0})[1 \ 0 \ 0]^T\| \|\mathbf{x}\| \\ &\leq 2\sin(\min(\sigma_\alpha/2, \pi/2))\|\mathbf{x}\| \leq \gamma_\alpha \|\mathbf{x}\|. \end{aligned} \quad (17)$$

For the depth  $d \in [d_0 - \sigma_d, d_0 + \sigma_d]$ , we simply have

$$|d - d_0| \leq \sigma_d \doteq \gamma_d. \quad (18)$$

Now let  $\mathbf{x}_j \in \mathcal{X} \cup \mathcal{Y}^r$  be the closest point to  $(\mathbf{x}_i - 2\mathbf{n}(\mathbf{x}_i^T \mathbf{n} + d))$ , and let  $\mathbf{x}_{j_0} \in \mathcal{X} \cup \mathcal{Y}^r$  be the closest point to  $(\mathbf{x}_i - 2\mathbf{n}_0(\mathbf{x}_i^T \mathbf{n}_0 + d_0))$ . The lower bound is derived as follows:

$$\begin{aligned} e_i^s(\boldsymbol{\alpha}, d|\mathbf{x}_i) &= \|\mathbf{x}_i - 2\mathbf{n}(\mathbf{x}_i^T \mathbf{n} + d) - \mathbf{x}_j\| \\ &= \|\mathbf{x}_i - 2\mathbf{n}_0(\mathbf{x}_i^T \mathbf{n}_0 + d_0) - \mathbf{x}_j - 2(\mathbf{n}(\mathbf{x}_i^T \mathbf{n} + d) \\ &\quad - \mathbf{n}_0(\mathbf{x}_i^T \mathbf{n}_0 + d_0))\| \\ &\geq \|\mathbf{x}_i - 2\mathbf{n}_0(\mathbf{x}_i^T \mathbf{n}_0 + d_0) - \mathbf{x}_{j_0}\| \\ &\quad - 2\|\mathbf{n}(\mathbf{x}_i^T \mathbf{n} + d) - \mathbf{n}_0(\mathbf{x}_i^T \mathbf{n}_0 + d_0)\| \\ &\geq e_i^s(\boldsymbol{\alpha}_0, d_0|\mathbf{x}_i) - 2\|\mathbf{x}_i^T \mathbf{n} \mathbf{n} - \mathbf{x}_i^T \mathbf{n}_0 \mathbf{n}_0\| - 2\|\mathbf{n}d - \mathbf{n}_0d_0\| \end{aligned} \quad (19)$$

We, furthermore, have

$$\begin{aligned} \|\mathbf{x}_i^T \mathbf{n} \mathbf{n} - \mathbf{x}_i^T \mathbf{n}_0 \mathbf{n}_0\| &= \|\mathbf{x}_i^T \mathbf{n} \mathbf{n} - \mathbf{x}_i^T \mathbf{n}_0 \mathbf{n} + \mathbf{x}_i^T \mathbf{n}_0 \mathbf{n} - \mathbf{x}_i^T \mathbf{n}_0 \mathbf{n}_0\| \\ &\leq |\mathbf{x}_i^T \mathbf{n} - \mathbf{x}_i^T \mathbf{n}_0| \cdot \|\mathbf{n}\| + |\mathbf{x}_i^T \mathbf{n}_0| \cdot \|\mathbf{n} - \mathbf{n}_0\| \\ &\leq 2\|\mathbf{x}_i\| \cdot \|\mathbf{n} - \mathbf{n}_0\| = 2\gamma_\alpha \|\mathbf{x}_i\| \end{aligned} \quad (20)$$

and

$$\begin{aligned} \|\mathbf{n}d - \mathbf{n}_0d_0\| &= \|\mathbf{n}d - \mathbf{n}d_0 + \mathbf{n}d_0 - \mathbf{n}_0d_0\| \\ &\leq |d - d_0| + |d_0| \|\mathbf{n} - \mathbf{n}_0\| \leq \gamma_d + |d_0| \gamma_\alpha. \end{aligned} \quad (21)$$

Substituting (20) and (21) in (19), we finally obtain

$$\begin{aligned} \underline{e}_i^s(\boldsymbol{\alpha}, d|\mathbf{x}_i) &\doteq \max\{e_i^s(\boldsymbol{\alpha}_0, d_0|\mathbf{x}_i) - 2(2\gamma_\alpha \|\mathbf{x}_i\| \\ &\quad + \gamma_d + |d_0| \gamma_\alpha), 0\}. \end{aligned} \quad (22)$$

**Lower Bound of symmetry term**  $e_i^s(\mathbf{r}, \mathbf{t}, \boldsymbol{\alpha}, d|y_i)$ : By using  $\hat{\mathbf{n}} = \mathbf{R}^T \mathbf{n}$  and  $\hat{d} = \mathbf{t}^T \mathbf{n} + d$ , we analogously derive

$$e_i^s(\mathbf{r}, \mathbf{t}, \boldsymbol{\alpha}, d|y_i) \geq e_i^s(\mathbf{r}_0, \mathbf{t}_0, \boldsymbol{\alpha}_0, d_0|y_i) - 2\|\mathbf{y}_i^T \hat{\mathbf{n}}\hat{\mathbf{n}} - \mathbf{y}_i^T \hat{\mathbf{n}}_0 \hat{\mathbf{n}}_0\| - 2\|\hat{\mathbf{n}}\hat{d} - \hat{\mathbf{n}}_0 \hat{d}_0\| \tag{23}$$

By substituting  $\hat{\mathbf{n}} = \mathbf{R}^T \mathbf{n}$ , similar to (20), the first term gives

$$\begin{aligned} & \|\mathbf{y}_i^T \mathbf{R}^T \mathbf{n} \mathbf{R}^T \mathbf{n} - \mathbf{y}_i^T \mathbf{R}_0^T \mathbf{n}_0 \mathbf{R}_0^T \mathbf{n}_0\| \\ & \leq 2\|\mathbf{y}_i\| \|\mathbf{R}_0^T \mathbf{n}_0 - \mathbf{R}^T \mathbf{n}\| \\ & = 2\|\mathbf{y}_i\| \|\mathbf{R}_0^T \mathbf{n}_0 - \mathbf{R}^T \mathbf{n}_0 + \mathbf{R}^T \mathbf{n}_0 - \mathbf{R}^T \mathbf{n}\| \\ & \leq 2\|\mathbf{y}_i\| (\gamma_\alpha + \gamma_r). \end{aligned} \tag{24}$$

By also substituting  $\hat{d} = \mathbf{t}^T \mathbf{n} + d$ , the second term gives

$$\begin{aligned} & \|(\mathbf{t}_0^T \mathbf{n}_0 + d_0) \mathbf{R}_0^T \mathbf{n}_0 - (\mathbf{t}^T \mathbf{n} + d) \mathbf{R}^T \mathbf{n}\| \\ & \leq |\mathbf{t}_0^T \mathbf{n}_0 + d_0| \cdot \|\mathbf{R}_0^T \mathbf{n}_0 - \mathbf{R}^T \mathbf{n}\| \\ & \quad + \|\mathbf{R}^T \mathbf{n}\| \|\mathbf{t}_0^T \mathbf{n}_0 - \mathbf{t}^T \mathbf{n}\| + \|\mathbf{R}^T \mathbf{n}\| |d_0 - d| \\ & \leq (\gamma_\alpha + \gamma_r) |\mathbf{t}_0^T \mathbf{n}_0 + d_0| + \|\mathbf{t}_0\| \gamma_\alpha + \gamma_t + \gamma_d, \end{aligned} \tag{25}$$

where we have used  $\|\mathbf{t}_0^T \mathbf{n}_0 - \mathbf{t}^T \mathbf{n}\| = \|\mathbf{t}_0^T \mathbf{n}_0 - \mathbf{t}_0^T \mathbf{n} + \mathbf{t}_0^T \mathbf{n} - \mathbf{t}^T \mathbf{n}\| \leq \|\mathbf{t}_0\| \|\mathbf{n}_0 - \mathbf{n}\| + \|\mathbf{n}\| \|\mathbf{t}_0 - \mathbf{t}\|$ . Substituting (24) and (25) in (23), we finally obtain

$$\begin{aligned} e_i^s(\mathbf{r}, \mathbf{t}, \boldsymbol{\alpha}, d|y_i) & \doteq \max\{e_i^s(\mathbf{r}_0, \mathbf{t}_0, \boldsymbol{\alpha}_0, d_0|y_i) - 2(\gamma_t + \gamma_d \\ & \quad + (\gamma_\alpha + \gamma_r) (2\|\mathbf{y}_i\| + |\mathbf{t}_0^T \mathbf{n}_0 + d_0|) + \|\mathbf{t}_0\| \gamma_\alpha), 0\}. \end{aligned} \tag{26}$$

Here, the lower bound has been obtained by inserting (10), (22) and (26) into (12).

### 5 Implementation

Similar to prior art [42], we improve the algorithm’s ability to handle the dimensionality of the problem by installing a nested BnB paradigm.

#### 5.1 Nested BnB

Detailed descriptions are given in Algorithm 1 (the Outer Algorithm) and Algorithm 2 (the Inner BnB). We install a nested BnB scheme in which the outer layer searches through the space  $C_{r\alpha}$  of all angular parameters (Algorithm 1), the inner layer optimizes over the space  $C_{td}$  of translation and depth (Algorithm 2). While finding the bounds in a sub-volume of the angle space, the algorithm calls the inner BnB algorithm to identify the optimal translation and depth. One

**Algorithm 1** Main Algorithm: BnB search for optimal registration and symmetry parameters

```

Input: Data and model point set; threshold  $\tau_1$ ; initial intervals  $C_{r\alpha}$  and  $C_{td}$ .
Output: Globally minimal error  $E^*$  and the optimal  $\mathbf{r}^*, \mathbf{t}^*, \boldsymbol{\alpha}^*, d^*$ 
Put  $C_{r\alpha}$  into priority queue  $Q_{r\alpha}$ . Set  $E^* = +\infty$ .
loop
  Read out interval with lowest  $\underline{E}_{r\alpha}$  from  $Q_{r\alpha}$ .
  Quit the loop if  $E^* - \underline{E}_{r\alpha} < \tau_1$ .
  Divide interval into sub-intervals.
  for each sub-interval  $C_{r\alpha}$  do
    Compute the corresponding optimal  $\mathbf{t}'$  and  $d'$  by calling Algorithm 2 with  $\mathbf{R}_{r_0}$  and  $\mathbf{n}_0(\boldsymbol{\alpha}_0)$ .
    Compute  $\bar{E}_{r\alpha}$  and  $\underline{E}_{r\alpha}$  for  $C_{r\alpha}$  with the estimated  $\mathbf{t}', d'$  (zero translation and depth uncertainty).
    if  $\bar{E}_{r\alpha} < E^*$  then
      Run ICP with initialization of  $(\mathbf{r}_0, \mathbf{t}', \boldsymbol{\alpha}_0, d')$ 
      Update  $E^*, \mathbf{r}^*, \boldsymbol{\alpha}^*, \mathbf{t}^*,$  and  $d^*$  with the results of ICP.
    end if
    Discard  $C_{r\alpha}$  if  $\underline{E}_{r\alpha} \geq E^*$ ; otherwise put it into  $Q_{r\alpha}$ 
  end for
end loop

```

**Algorithm 2** BnB search for optimal translation and depth given rotation and normal

```

Input: Data and model point set; threshold  $\tau_2$ ; initial intervals  $C_{td}$ ; Currently lowest error  $E^*$ 
Output: Minimal error  $E^*$  and the optimal  $\mathbf{t}^*, d^*$ 
Put  $C_{td}$  into priority queue  $Q_{td}$ .
loop
  Read out interval with lowest  $\underline{E}_{td}$  from  $Q_{td}$ .
  Quit the loop if  $E^* - \underline{E}_{td} < \tau_2$ .
  Divide interval into sub-intervals.
  for each sub-interval  $C_{td}$  do
    Compute  $\bar{E}_{td}$  and  $\underline{E}_{td}$  for  $C_{td}$  with the  $\mathbf{r}_0, \boldsymbol{\alpha}_0$  (zero rotation and normal uncertainty).
    if  $\bar{E}_{td} < E^*$  then
      Update  $E^*$  and  $\mathbf{t}^*, d^*$ .
    end if
    Discard  $C_{td}$  if  $\underline{E}_{td} \geq E^*$ ; otherwise put it into  $Q_{td}$ 
  end for
end loop

```

important approximation that accelerates the execution is that when we estimate the bounds in the outer BnB search, the uncertainties  $\gamma_t$  and  $\gamma_d$  are set to zero. When we estimate the bounds in the inner BnB, the uncertainties  $\gamma_\alpha$  and  $\gamma_r$  are set to zero. Nested BnB implementations are commonly used to speed up the execution. For more details, please refer to [42].

#### 5.2 Integration with Local ICP

Within the outer layer, whenever BnB identifies an interval  $C_{r\alpha}$  with an improved upper bound, we will execute a conventional local ICP algorithm starting from the centre of the  $C_{r\alpha}$  and taking  $\mathbf{t}^*$  and  $d^*$  as an initial value. Once ICP converges to a local minimum with a lower function value, the new value is used to further reduce the upper bound. The technique is inspired by Yang et al. [43].

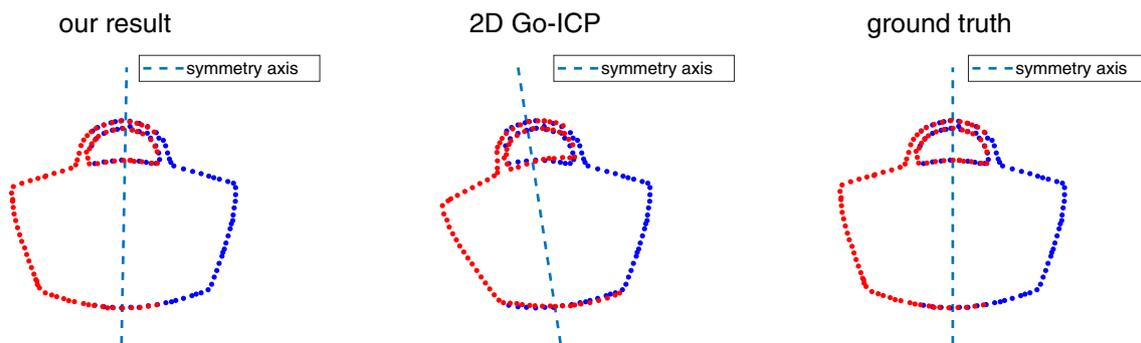


Fig. 1 Example of 2D point set registration with an overlap ratio of 0.3302

### 5.3 Trimming for Outlier Handling

A general problem with partially overlapping point sets is that—even at the global optimum—some points may simply not have a correspondence, and should hence be treated as outliers. Although the addition of symmetry and point reflections already greatly alleviates this problem, we still add the strategy proposed in Trimmed ICP [13] for robust point-set registration. More specifically, in each iteration, only a subset of the matched data points with smallest point-to-point distances are used for motion computation. In this work, we choose a 90%-subset for both symmetry and registration residuals.

## 6 Experiments

We now report our experimental results on both synthetic and real data. In all experiments, we pre-normalize the pointsets such that all points are within the volume  $[-1, 1]^3$ . The space of possible axis-angle rotation vectors  $\mathbf{r}_i$  and translation vectors  $\mathbf{t}_i$  can therefore be limited to the intervals  $[-\pi, +\pi]$  and  $[-0.5, 0.5]$  in each dimension, respectively. For the symmetry plane, the angles of the normal vector and the scale  $d$  can be constrained to the intervals  $[-\frac{\pi}{2}, +\frac{\pi}{2}]$  and  $[-0.5, +0.5]$ , respectively. The threshold  $c$  and penalty parameter  $w_1$  in (10) are set to 0.02 and 2. The weight  $w_2$  balancing the registration and symmetry costs in (11) is 0.1. We run experiments on both 2D and 3D data. For the stopping criteria, parameter  $\tau_1$  in the outer-BnB of both our and reference implementations are set to 0.01, and the threshold  $\tau_2$  of the inner-BnB is set to 0.03. Note that these values are too small to have any noticeable impact on the convergence of the algorithm. Theoretically, making  $\tau_1$  and  $\tau_2$  smaller would merely make the results more accurate by trading computational efficiency versus higher solution space resolution and thus accuracy. The parameters used in our experiments simply strike a good balance between runtime and accuracy.

### 6.1 Performance on 2D Synthetic Data

We compare our algorithm against a sequential combination of the 2D version of Go-ICP [42] for registration and the ransac-based algorithm presented in [15] for symmetry detection over the aligned point sets. In the case of 2D data, there are 5 degrees of freedom (DoF) to estimate. One is given by the relative rotation angle, two by the translational displacement, one by the symmetry normal angle, and one by the corresponding scale. Note that in this case, the axis-angle representations of normal  $\alpha$  and rotation  $\mathbf{r}$  become a 1-dimensional vector (i.e. a scalar).

**Description of experiments** Each experiment is generated by taking an image that contains a symmetrical object and using the Sobel edge detector to extract the object's contour points. To evaluate the performance, we randomly divide the contour points into two subsets with a defined and structured overlap. To conclude,  $\mathcal{Y}$  is transformed by a random rotation and translation drawn from the intervals  $\pm 180$  degrees and  $\pm 0.5$ .

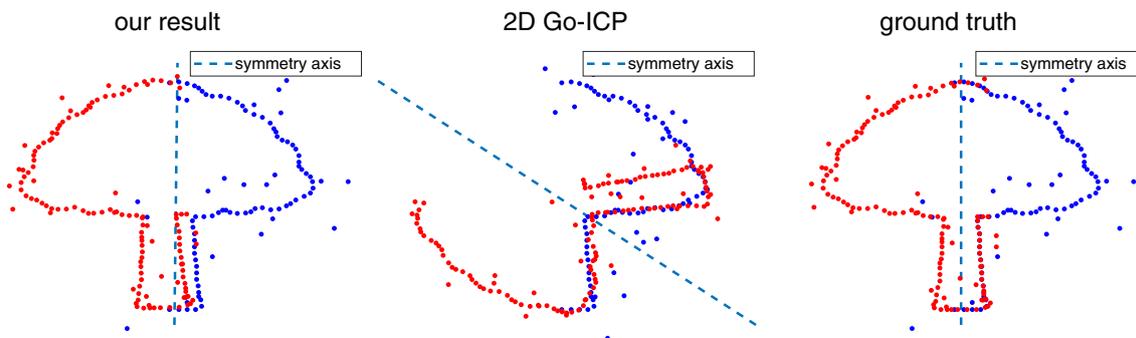
**Handling of limited overlap** The overlap between both point sets is varied from 10 to 83%. For each overlap ratio, we repeat 50 experiments each time choosing a random object, rotation and translation. Figure 1 shows an example result, where the left is the result of our proposed method, the centre shows the result of the 2D version of Go-ICP, and the right one shows the ground truth alignment. Table 1 shows the rotation and symmetry plane normal errors. As demonstrated by vastly reduced mean errors, our method has a substantially better ability to deal with limited partial overlaps compared to Go-ICP. Note that results for translation and depth behave analogously.

**Outlier handling** To test resilience against outliers, we repeat the same experiment but add up to 30% outliers with a zero mean and standard deviation 0.05 to both  $\mathcal{X}$  and  $\mathcal{Y}$ . Figure 2 indicates an example result, and Table 2 again illustrates the evaluated errors over all experiments. While the registration error starts to increase earlier (starting from an overlap ratio within the interval  $[0.4, 0.6]$ ) and average angular errors

**Table 1** Rotation symmetry normal errors on 2D synthetic data with no outliers

Overlap ratio	[0, 0.2]	[0.2, 0.4]	[0.4, 0.6]	[0.6, 0.8]	[0.8, 1]
Rotation (ours)	<b>21.78/4.64</b>	<b>15.69/12.96</b>	<b>1.86/0.40</b>	0.70/0.31	0.35/0.23
Rotation (Go-ICP)	70.32/51.73	42.93/13.75	3.57/0.45	<b>0.69/0.20</b>	<b>0.34/0.14</b>
Symmetry normal (ours)	<b>26.84/20.29</b>	<b>12.62/11.70</b>	<b>1.82/0.86</b>	<b>0.90/0.49</b>	0.54/0.34
Symmetry normal (Go-ICP)	60.81/78.46	30.62/ <b>2.00</b>	5.02/ <b>0.37</b>	1.79/ <b>0.37</b>	<b>0.41/0.22</b>

Indicated errors are the mean/median in unit degrees



**Fig. 2** Example of 2D point set registration with an overlap ratio of 0.3288 and 30% outliers

**Table 2** Rotation symmetry normal errors on 2D synthetic data with 30% outliers

Overlap ratio	[0, 0.2]	[0.2, 0.4]	[0.4, 0.6]	[0.6, 0.8]	[0.8, 1]
Rotation error (ours)	<b>17.33/13.33</b>	<b>23.61/19.73</b>	<b>5.37/1.07</b>	1.90/0.65	1.32/0.43
Rotation error (Go-ICP)	82.72/61.39	64.04/42.70	21.63/ <b>0.76</b>	<b>1.49/0.17</b>	<b>0.54/0.17</b>
Symmetry normal (ours)	<b>37.91/23.73</b>	<b>15.46/12.89</b>	<b>4.07/1.11</b>	<b>2.26/0.85</b>	1.31/0.70
Symmetry normal (Go-ICP)	62.38/78.48	35.43/ <b>7.62</b>	21.17/1.39	5.14/ <b>0.48</b>	<b>0.83/0.32</b>

Indicated errors are the mean/median in unit degrees

tend to be higher, it can still be concluded that our method significantly outperforms the combination of Go-ICP and the symmetry plane detection method in [15].

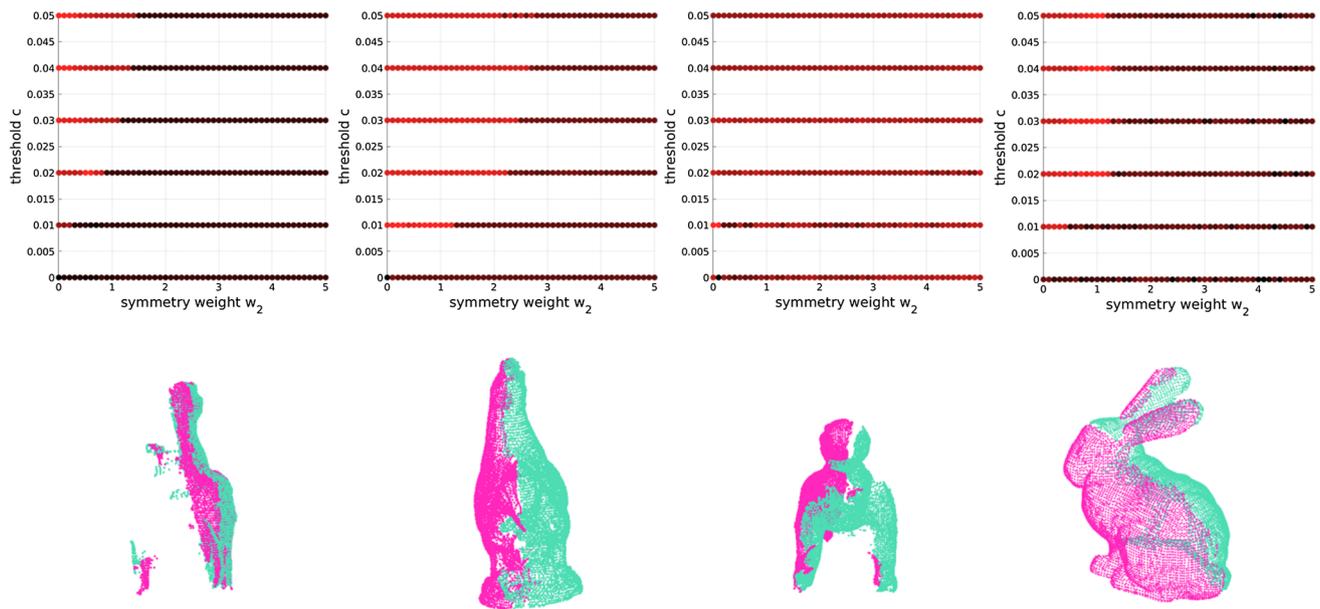
### 6.2 Performance on Open 3D Shape Data

The disadvantage of BnB is that its complexity grows exponentially in the dimensionality of the problem. We therefore take prior information about the 3D point sets into account that helps to reduce the dimensionality. We make the assumption that most objects are standing upright on the ground plane. The rotation between different partial point sets is therefore still constrained to be a 1D rotation about the vertical axis. We, furthermore, make the assumption that the symmetry plane is vertical; thus, the normal vector  $\mathbf{n}$  remains a 1D variable and in the horizontal plane. Finally, given zero vertical displacements, the translation remains a 2D vector over the interval  $[-0.5, +0.5]^2$ . In summary, the 3D case is still solved as a 5-dimensional estimation problem by using the assumption of upright placement on the ground plane. Hence, the axis-angle representations of normal  $\alpha$  and rotation  $\mathbf{r}$  are also reduced to 1-degree of freedom variables. In all object registration visualizations here, the red points are

the *model* (or *target*) points, and the green points are the *data* (or *source*) points. The symmetry plane is indicated in yellow. Note that the assumption that the ground plane is known is a very common assumption in object-level SLAM systems such as SLAM++ [34] and CubeSLAM [44]. The ground plane could be detected algorithmically, and its extraction could be supported by an IMU sensor.

*Description of experiments* We use open 3D shape data to test our algorithm. The point-sets are generated by different scanners with differing noise levels. We use the *bunny*, *dragon* and *Buddha* models from the Stanford 3D dataset. The *chef*, *dinosaur* and *chicken* models are from [25]. While the latter three models are practically symmetric, the bunny, dragon and Buddha models are only partially symmetric shapes. Each model contains multiple scans, and we do the alignment for each pair of two scans.

*Parameter Analysis* Before illustrating further comparative results, we first analyse the influence of some of the parameters used in the objective (11). [8,43] already analyse many details about branch and bound for point set registration. Here we focus on further analysing how the quality depends on the adjustment of the weighting factor  $w_2$  and the threshold parameter  $c$ . Figure 3 indicates the intersection-over-union



**Fig. 3** Parameter analysis: Each column shows one individual example registered using different choices for  $w_2$  and  $c$ . Results in red indicate a high IoU with ground truth, and darker colours inferior results

(IoU) with the ground truth bounding box resulting from different parameter combinations, a measure that encapsulates both the quality of the registration and the accuracy of the fitted symmetry plane. The IoU is each time indicated relatively through a colour code. Red means a high IoU, whereas darker colours down to black indicate decreasing IoU results. The second row gives a qualitative impression of the correctly registered point-sets of each corresponding model. Each column in the figure illustrates a different case:

- The first column represents the most common situation in which the two point sets have a large overlap ratio and are distributed mostly on one side of the symmetry plane. The symmetry parameter  $w_2$  disturbs the accuracy of the registration if it is set to a larger value. In order to maintain correct functionality for non-symmetric or at least largely overlapping examples, the parameter  $w_2$  should be chosen smaller than 0.15.
- The second column illustrates a constellation in which the point sets are distributed on different sides of the symmetry plane. It is the case our algorithm is designed for; thus, the sensitivity with respect to a proper choice of the weight  $w_2$  is lower compared to the first example.
- The third example is a case in which one point set is symmetric in itself, whereas the other point set contains points from one side of the symmetry plane, only. The registration in this case is generally less stable, irrespective of the choice of the weight parameter  $w_2$ . The case is affected by an increase in mismatched correspondences

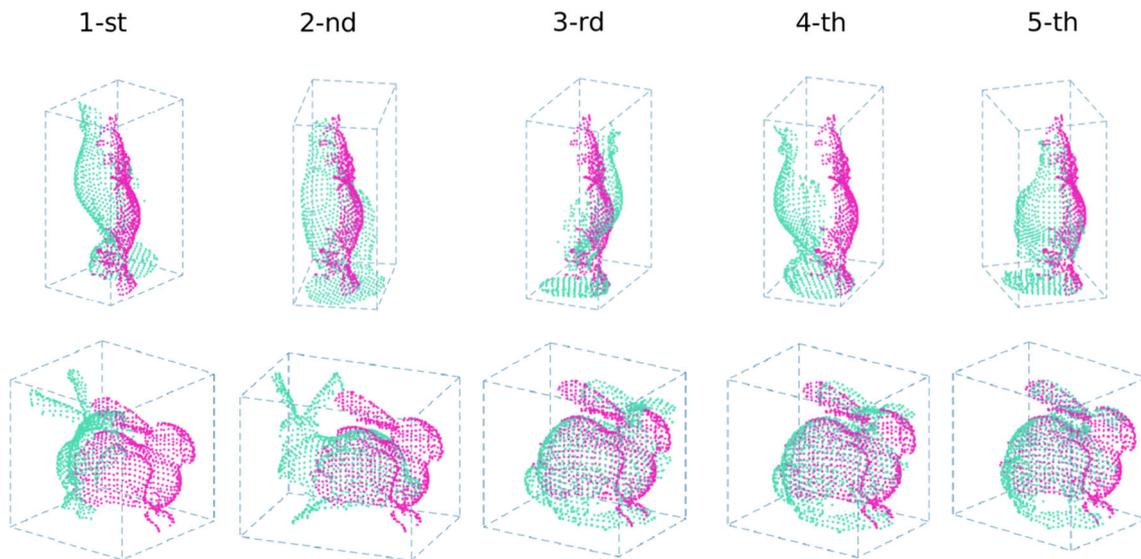
since even the correctly estimated symmetry plane will not increase the overlap ratio.

- The fourth column illustrates an example of registering two scans from either side of a partially symmetric shape. Estimating a plane of symmetry is still helpful in increasing the overlap ratio. Constraints on the weight  $w_2$  are similar than in the first situation.

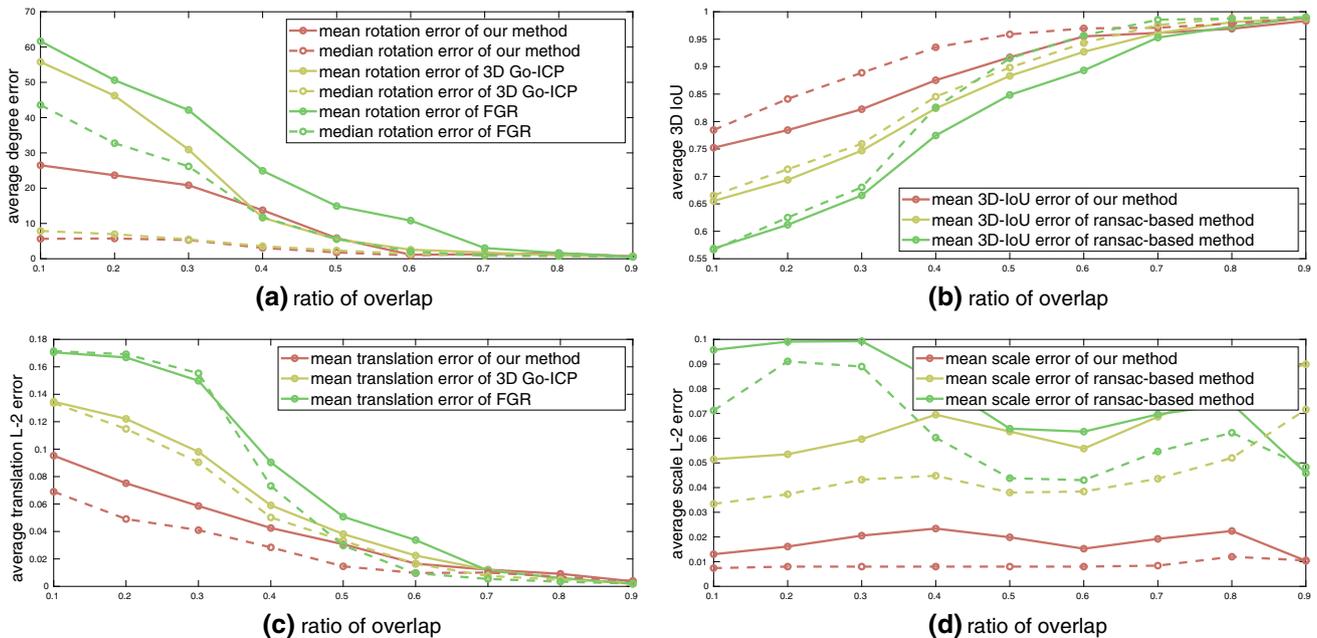
In summary,  $w_2$  should not be chosen too large and the sensitivity with respect to the threshold parameter  $c$  is generally much lower. Note, however, that the intervals for successful parameter tuning are sufficiently large and no further parameter tuning for individual experiments is required once acceptable values have been set. During all experiments in this paper, the parameters have indeed never been adjusted.

**Visualization of iterations** Figure 4 illustrates the iterations of two different registration examples. Note that the orientation and location of each bounding box is constrained to be symmetric with respect to the current symmetry plane location. In each example, we illustrate the first five instances of the currently best registration parameters  $E^*$ ,  $r^*$ ,  $\alpha^*$ ,  $t^*$  and  $d^*$ . As can be observed, our proposed algorithm gradually finds the symmetry plane and leads to a good registration result.

**Error Evaluation for 3D data** We compare our method against Go-ICP [42] as well as the feature-based alignment method fast global registration (FGR) [48]. For the latter approach, the parameters for point downsampling and feature extraction are adjusted to maximize the quality of the registration results. The maximum number of iterations is



**Fig. 4** Iterative registration updates for two different examples



**Fig. 5** Errors of 3D registration on symmetric data compared against 3D Go-ICP [43] and FGR [48], both followed by ransac-based symmetry detection [15]

set to 500, which is significantly larger than the default value of 200. For FGR, the symmetry plane is again fitted in a post-processing step taken from [15]. The stopping threshold  $\tau_1$  of the outer BnB of both our algorithm and Go-ICP is set to 0.01, while the threshold for the inner BnB is set to 0.03 to reduce the execution time. Figures 5 and 6 show quantitative results for perfectly symmetric and partially symmetric shapes, respectively. In order to also evaluate the quality of the fitted symmetry plane, we continue to evaluate the (3D) IoU between the recovered and the ground truth bound-

ing box. Note, however, that, as indicated in subfigures (d), the measure of the symmetry plane scale accuracy does not depend much on the actual overlap ratio in the partial point clouds. If the point sets have a large overlap they may cover only a part of the object, thus not constraining the symmetry plane. On the other hand, if the overlap is small, our method may actually achieve slightly better accuracy.

If we compare Figs. 5 and 6, we observe that our algorithm performs better than Go-ICP and FGR, for both symmetric and partially symmetric shapes. The trend of the average

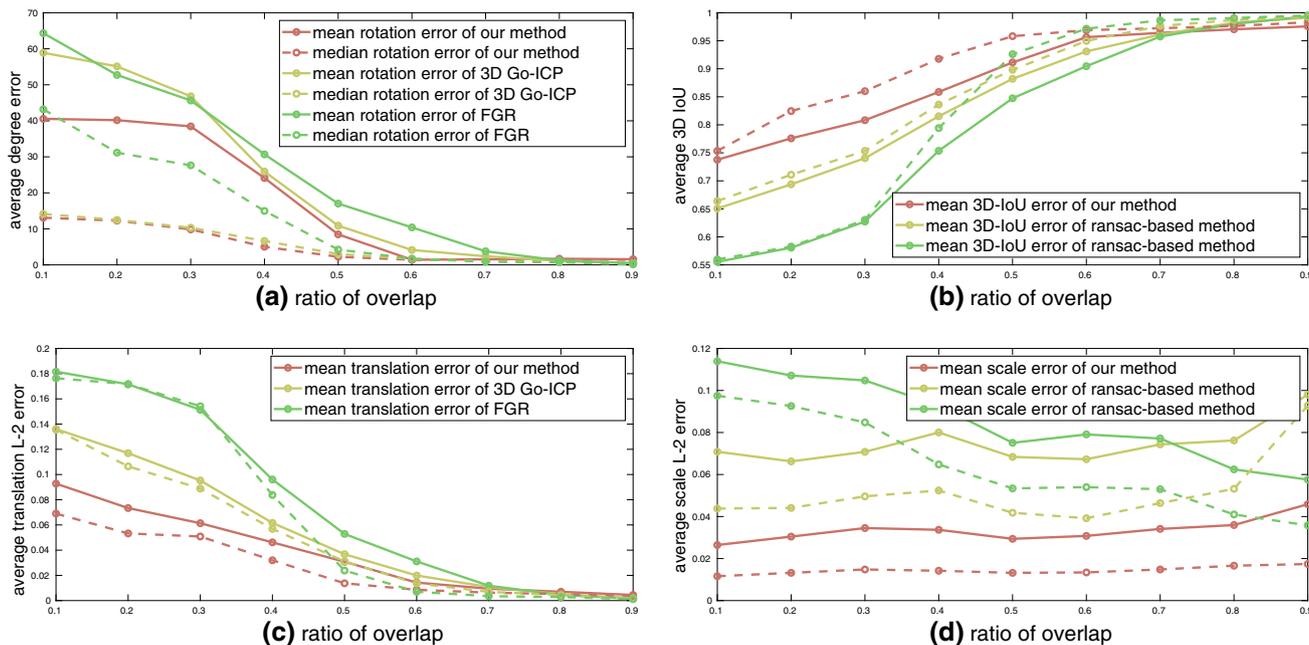


Fig. 6 Errors of 3D registration on partially symmetric data compared against 3D Go-ICP and FGR, both followed by ransac-based symmetry detection [15]

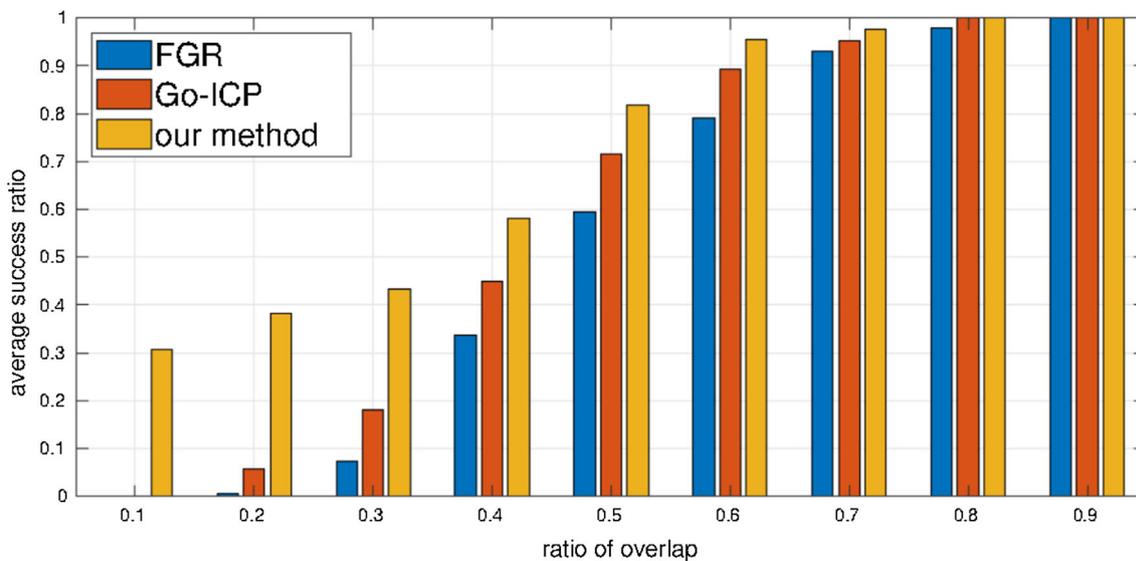
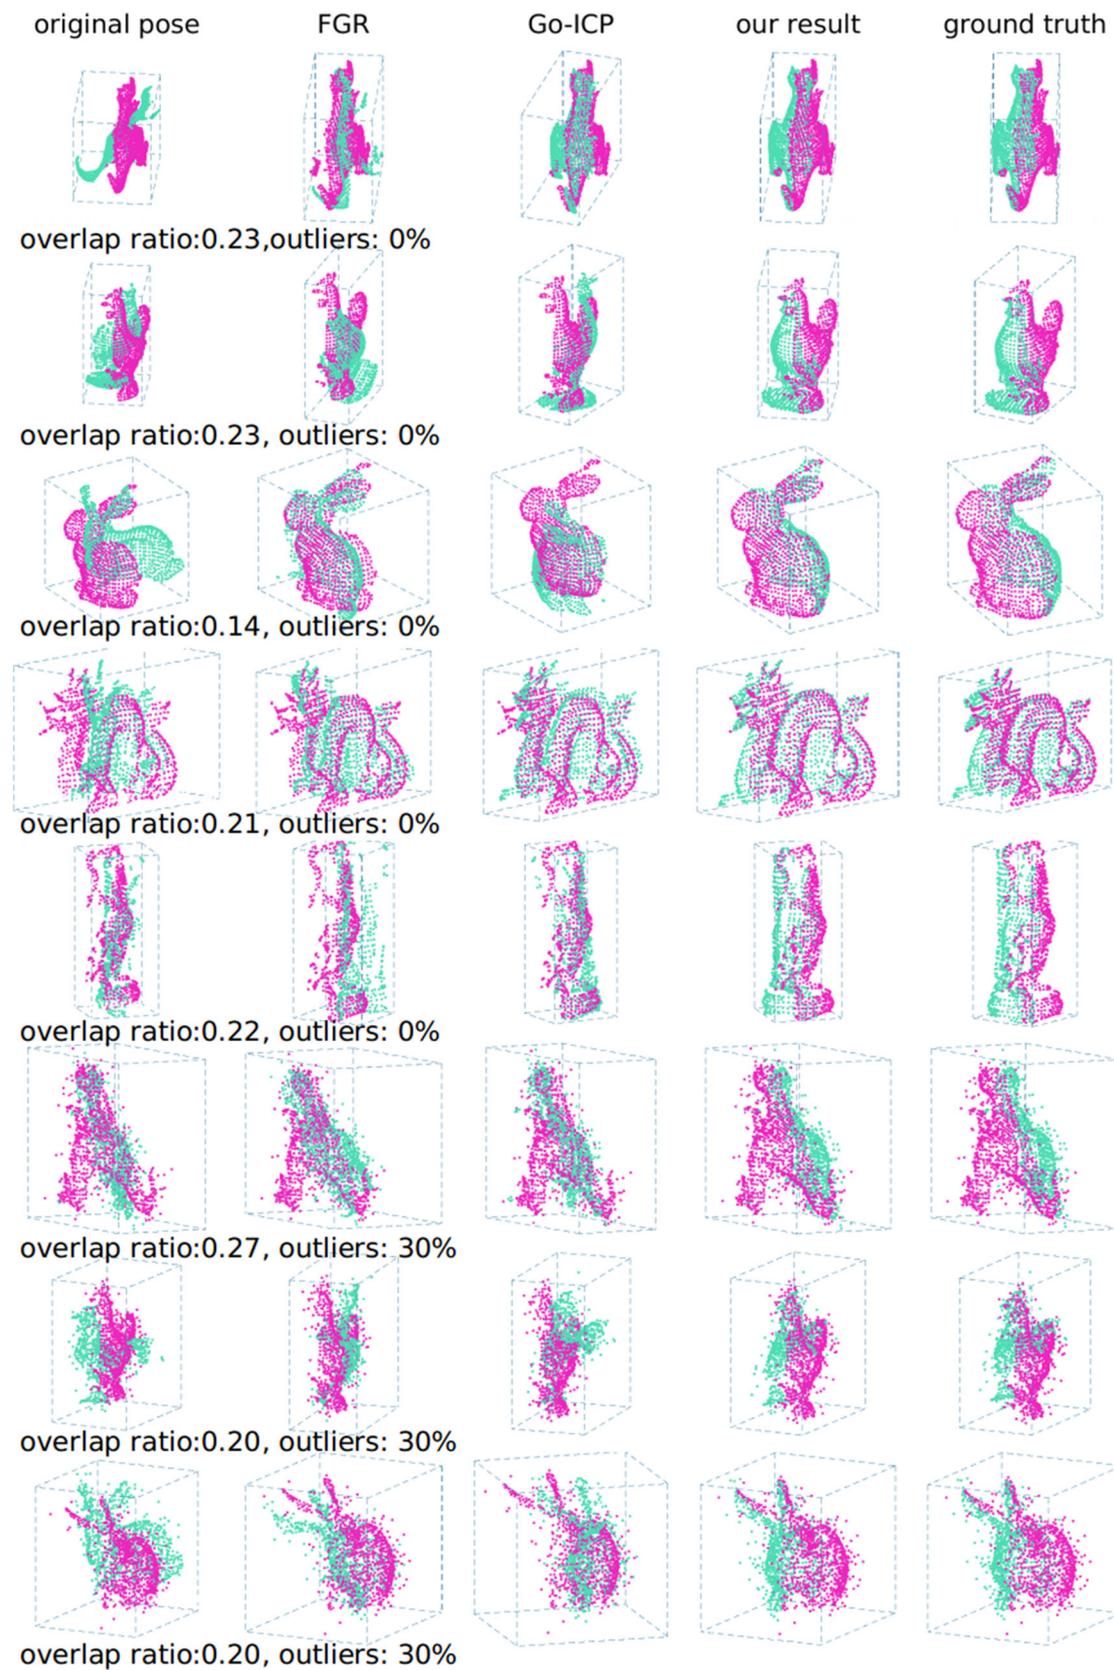


Fig. 7 Average success ratio over 1522 pair-points from the 3D shape data

success ratio we evaluated for all 3D shape registrations in Fig. 7 is similar. We define a registration success when the rotation error is less than 5 degrees and the translation error is less than 0.05. As can be observed, when the overlap ratio is less than 0.4, Go-ICP and FGR have a low success ratio, whereas our algorithm still obtains more than 30% success ratio. Figure 8 illustrates more visual examples on registered point clouds. Noise and overlap ratio are indicated at the bottom left of each example.

**Asymmetric objects and objects with multiple symmetry planes** Our proposed algorithm assumes that the object has at least one symmetry plane. We therefore evaluate the performance of our algorithm on asymmetric objects and objects with multiple symmetry planes. The objects tested here are from ModelNet40 [38]. The reader is referred to the deep shape registration work of [45] to look up how the test cases are generated. Figure 9 shows three objects: a piano, a stair and a sofa (the latter one still being partially symmetric). The detected symmetry varies and is likely to be a hyper-



**Fig. 8** 3D alignment results for concrete experiments. Overlap ratio and added outliers are indicated for each pair

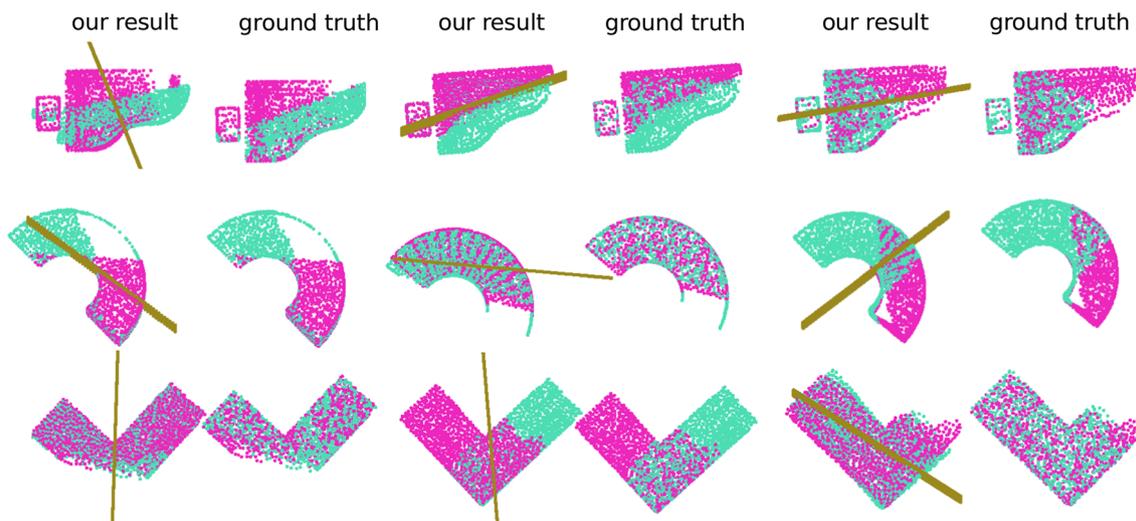


Fig. 9 Example on asymmetric objects, the three row objects are *Piano*, *Stair*, *Sofa*

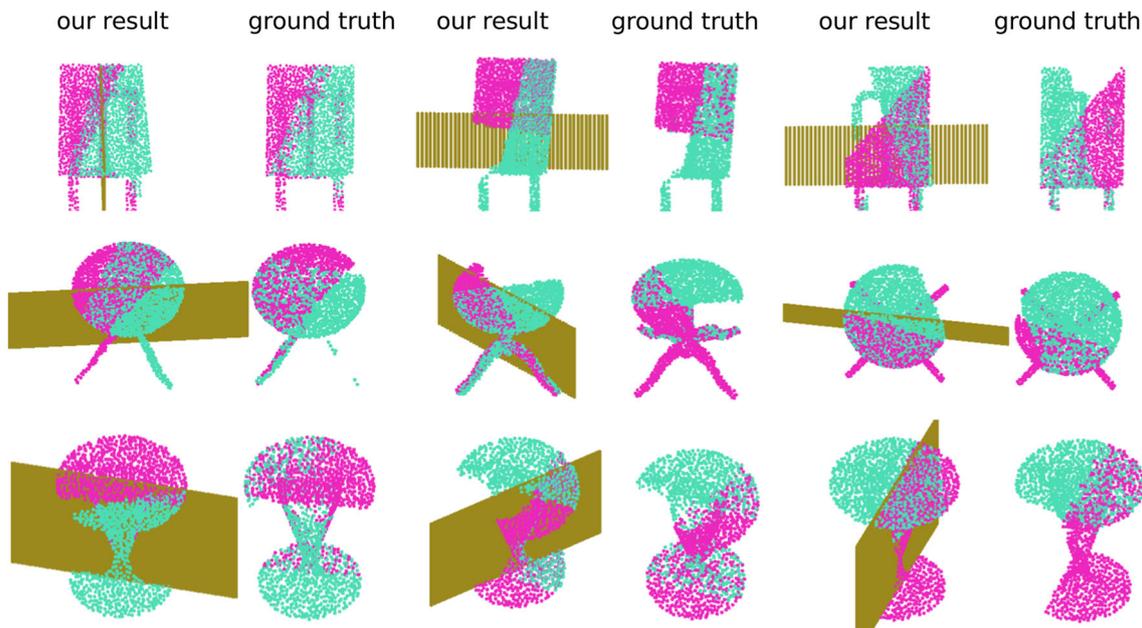


Fig. 10 Visualization of objects with multiple symmetry planes

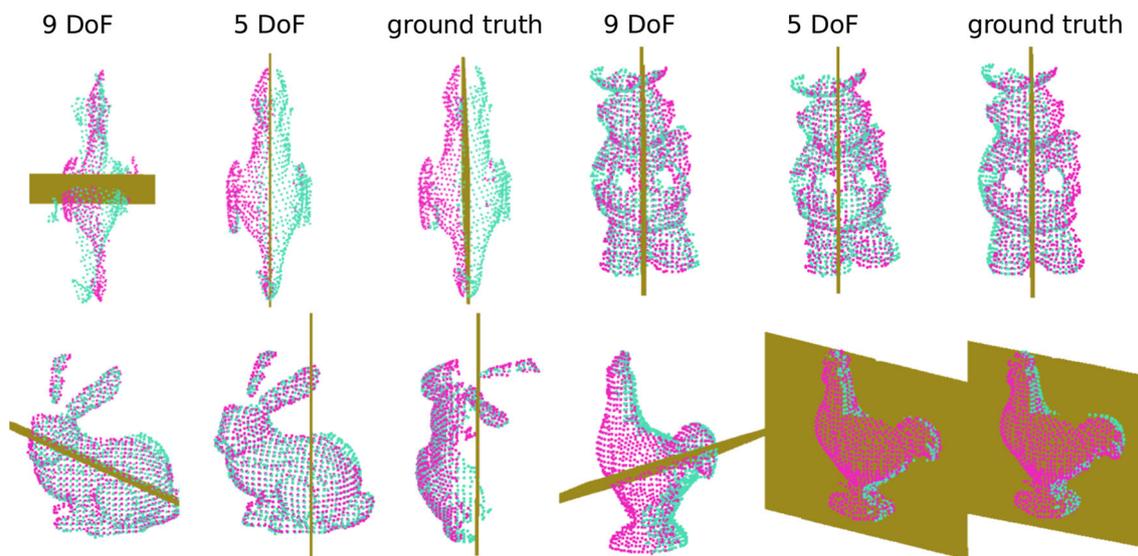
plane that divides the points into two similarly large point sets that—if reflected—maximize the symmetry fitting objective for the points. The situation for objects with multiple symmetry planes is slightly different. Figure 10 shows results for three different tables with 2, 8 and infinitely many symmetry planes. When an object has more than one symmetry plane, there will simply be more than one possible registration solution. While some of the aligned results are different from the ground truth alignment, they are still valid solutions in the sense of the fitted objective. Which solution will be picked is hard to predict and may depend on properties such as noise and point density.

**Time cost**

The branch-and-bound algorithm typically searches the whole parameter space and therefore has an exponential worst-case time complexity. If we do not make the assumption that objects are placed upright with known gravity direction, the total number of parameters to estimate in the 3D case will be 9. In the branch step, we need to search  $2^9 = 512$  sub-intervals resulting in 16 times more sub-intervals than in the 5 parameter case. Table 3 gives the time cost of 9-parameter optimization examples compared against the original 5 degree of freedom case. The time cost for 9 degrees of freedom is around 16 times higher, which agrees

**Table 3** Runtime comparison between 5 and 9 parameter optimizations

Time(s)	Chef	Dinosaur	Chicken	Bunny	Dragon	Buddha
5 DoF	36.8417	26.2592	25.6426	22.0002	19.6215	22.3936
9 DoF	1115.7622	1264.3211	1789.4401	1668.7207	1092.5371	1595.2189

**Fig. 11** Visualization of 5 DoF and 9 DoF optimization

with the initial intuition. Figure 11 illustrates four examples comparing 5 and 9 parameter scenarios. For 9 degrees of freedom, the number of potential solutions for registration and symmetry detection is potentially higher, which further increases the difficulty to identify the ground truth solution.

Compared against Go-ICP, the improved accuracy of our algorithm comes at the expense of increased time cost. There are two reasons. First, our algorithm branches over two additional variables, which are normal and scale of the symmetry plane. Second, our algorithm finds the nearest points in a union of sets of points which is incrementally updated as the algorithm proceeds. As indicated in Fig. 12, our algorithm is slower than the original Go-ICP. While it is not suitable for real-time shape matching, it is still instrumental in demonstrating the value of joint symmetry plane fitting, and we believe that the algorithm can be employed as part of a latency-tolerant back-end thread that for example performs shape completion.

**Failure cases** Figure 13 illustrates situations in which our algorithm fails. When the missing part of the source point set does not have corresponding points in the target set or its reflection (points indicated by the red dotted circles in Fig. 13), registration quality decreases. While this result is common and intuitively clear, we would like to stress that the achievable accuracy or registration quality by our algorithm would still be comparable to Go-ICP.

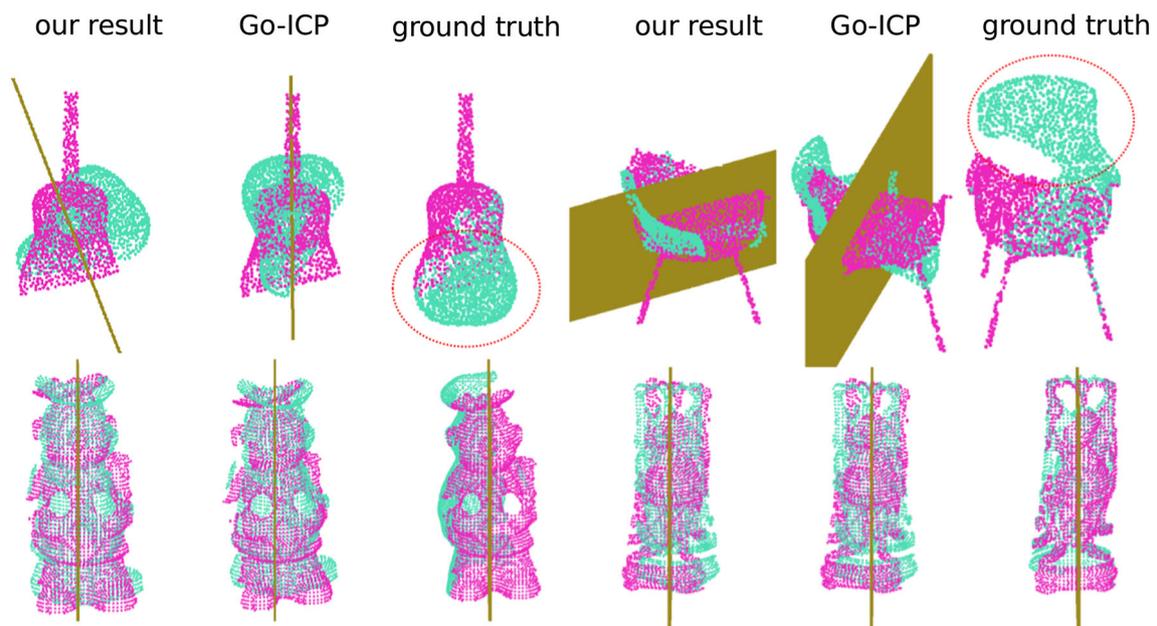
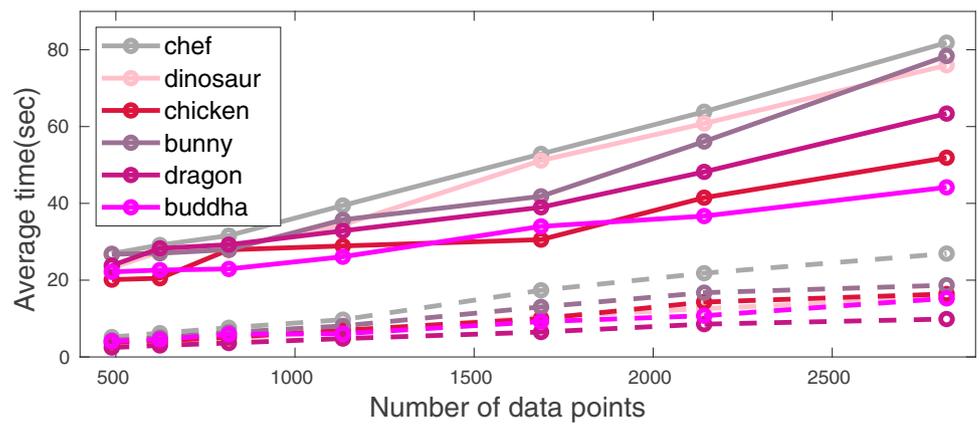
**Performance on further 3D synthetic scans** To evaluate the performance on more data, we choose 24 symmetric CAD models from ShapeNet [9] containing 8 individual types of 3 different classes. For each example, we generate 16 depth images (with occlusions) from random views around the object (rejecting very similar views). We sample random pairs of scans from each set of 16 depth images and add random transformations, thus leading to a total of about 6000 point-set registration experiments. Figure 14 illustrates some of the scans aligned by our algorithm. Figure 15 again indicates mean and median errors of all estimated quantities as a function of the overlap between the sets.

FGR suffers from its reliance on the 3D feature FPFH, for which the varying density in the synthetic data influences accuracy. This can be observed from the green lines, indicating angular errors that tend towards  $90^\circ$ , the centre of the interval  $[0; 180]$ .

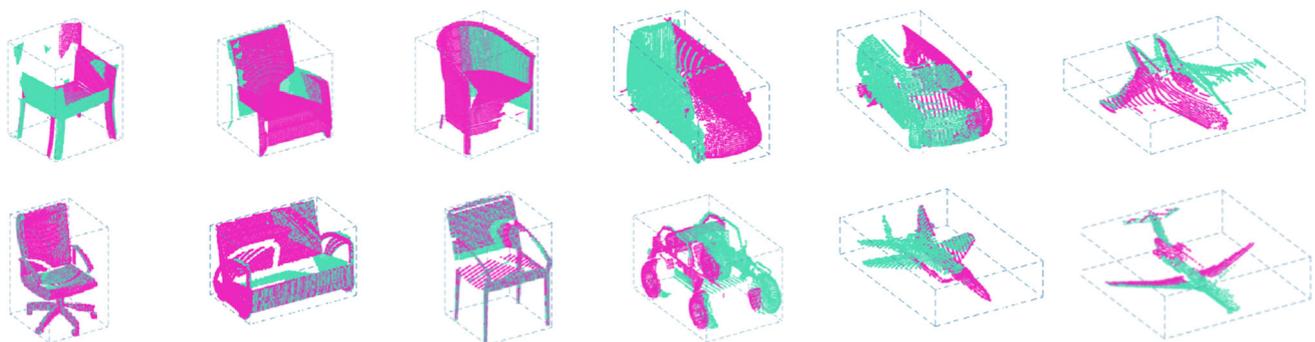
### 6.3 Experiments on Real Scene

Our last experiment is an exciting application to real data that goes back to the initial motivation in the introduction. Figure 16 shows depth images captured by a Kinect camera, each one containing three instances of the same object under different orientations. By pairwise alignment of partial object scans, the mutual information is transferred thus leading to more complete perception of each individual object.

**Fig. 12** Execution time for scan registration on the shape dataset with different point densities



**Fig. 13** Failure cases for examples where the source point set does not have corresponding points in either target set or its reflection



**Fig. 14** Alignment results with partial scans generated from CAD models from ShapeNet [9]

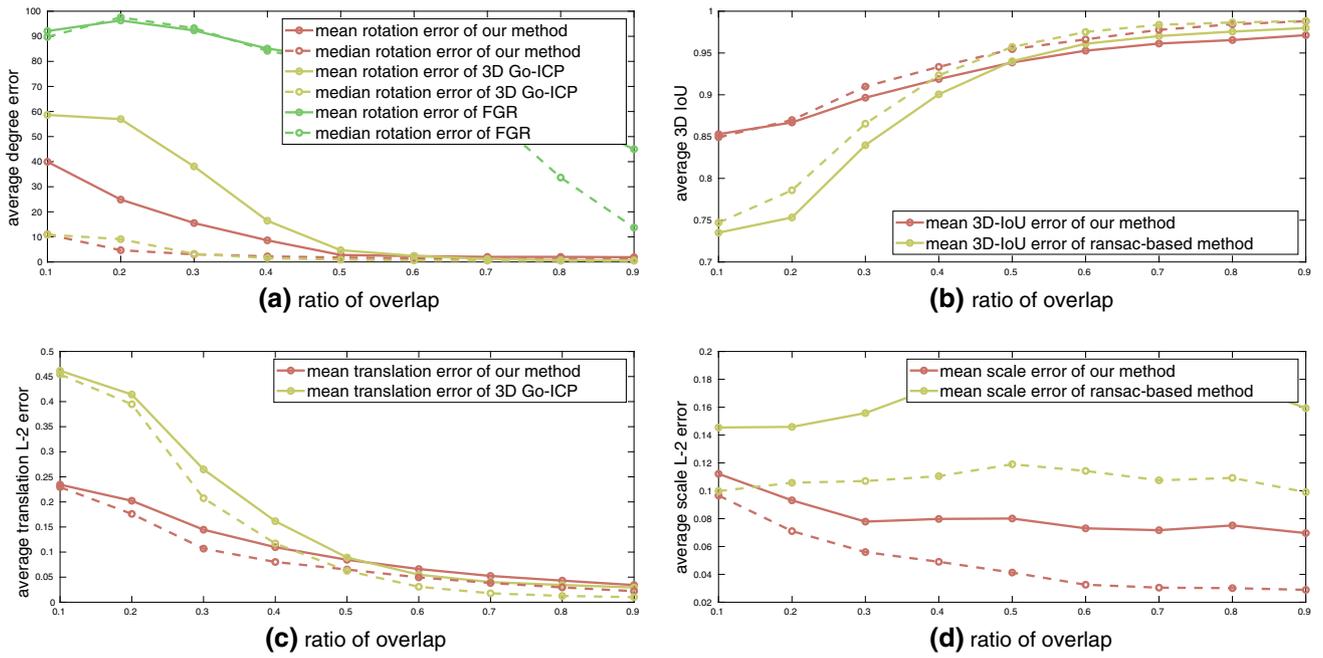


Fig. 15 3D registration errors averaged over 6000 experiments generated from CAD models from ShapeNet [9]

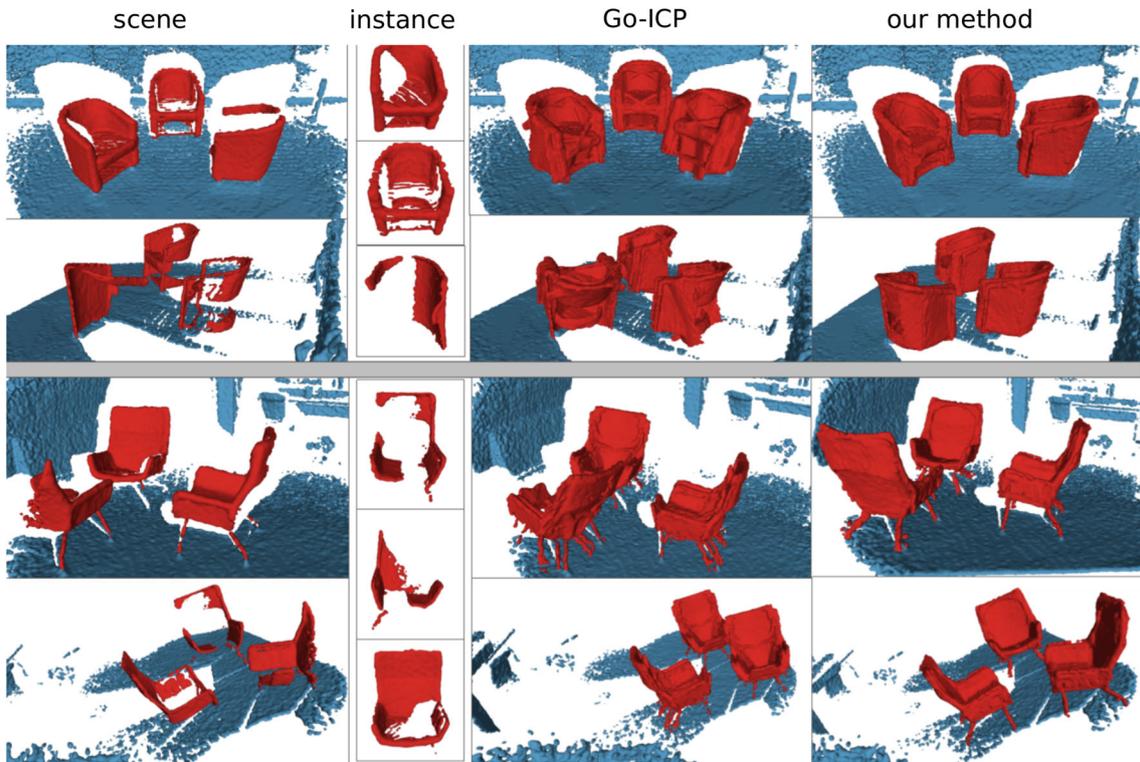
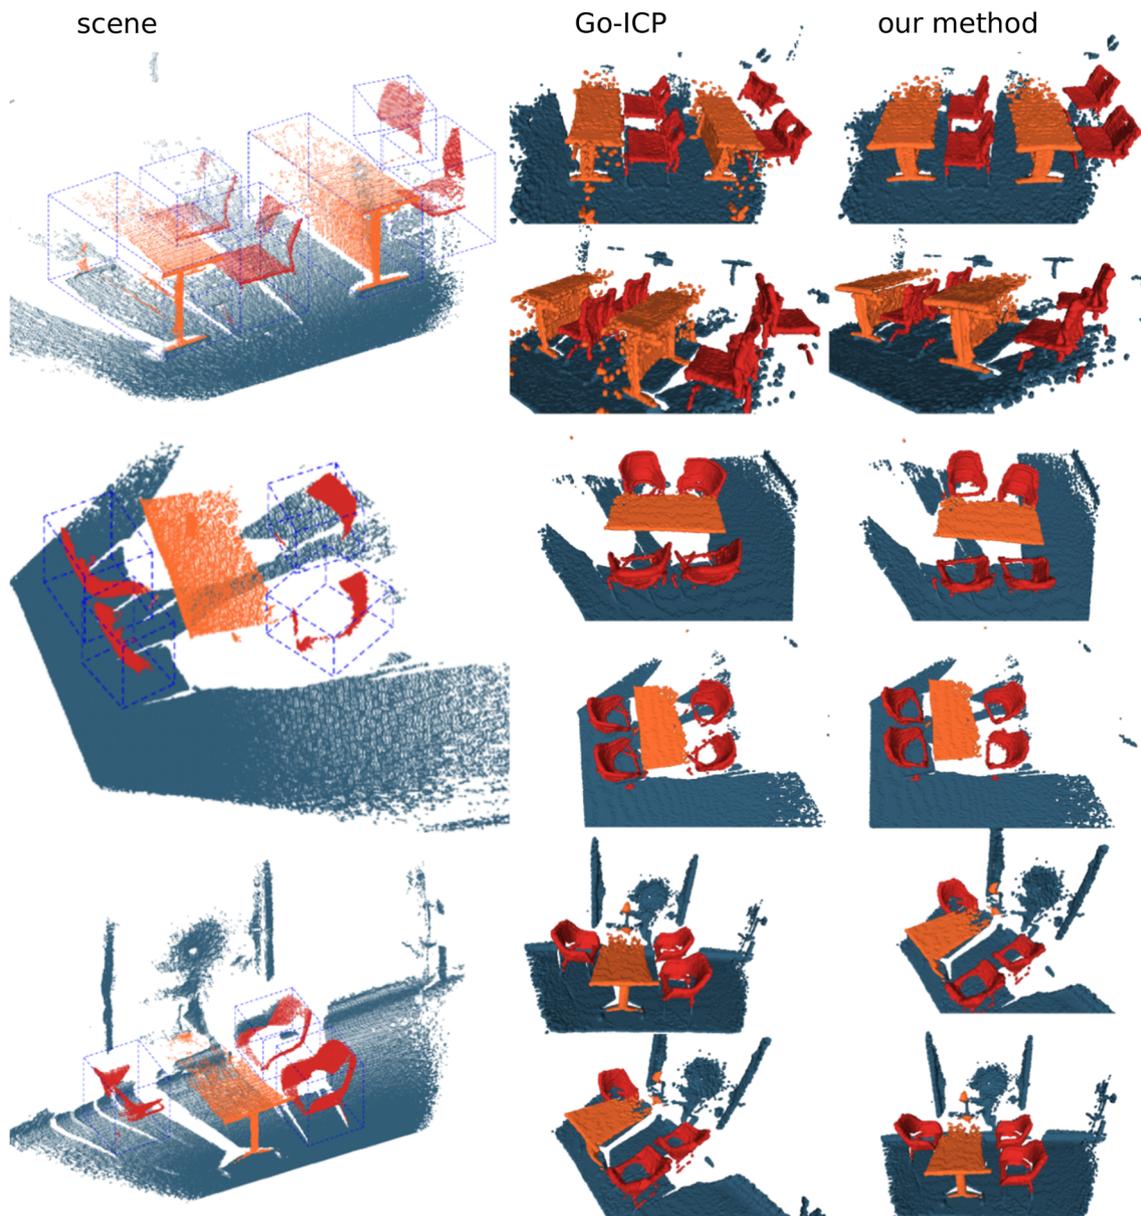


Fig. 16 Shape completion by aligning partial scans of identical objects in RGBD images



**Fig. 17** Further results on shape completion in scenes with multiple instances of the same object, this time with objects from different classes

Note that we use simple ground plane fitting and depth discontinuity-aware point clustering within object bounding boxes to isolate the partial object scans. With known position of the ground plane, we then transform the whole scene to be orthogonal to the ground plane and meet the assumption that all objects are placed upright and—in terms of relative rotation—differ only by an angle about the vertical axis. For pair-wise registration, we choose the partial shape observation with the most points as the reference with respect to which all other scans are registered. In Fig. 16, the first column shows the original scan in different orientations, the second one the partial object measurements, the third one the completion obtained by using Go-ICP as an

alignment algorithm and ransac-based symmetry detection, and the last one the result obtained by using our algorithm. As can be observed, our joint alignment strategy outperforms the comparison method and achieves more meaningful shape completion results.

Figure 17 shows three more examples on shape completion, the difference being the fact that these scenes contain multiple objects of different classes. Using the support of a semantic object classifier, our method is still able to transfer mutual information between identical objects and thus complete and improve the individual shape representations. Note, furthermore, that—while the accurate estimation of the symmetry plane orientation remains challenging—the addi-

tional estimation of symmetry parameters strongly supports the placement of a bounding box with strong hints on the object's pose.

## 7 Discussion

Symmetry detection and point set alignment over sets with small overlap are challenging problems if handled separately. Our work makes two main contributions. First, we show that those two problems become substantially easier if solved jointly. Second, we demonstrate a substantial improvement in both accuracy and success rate of the alignment by solving jointly for the symmetry parameters. The information gained from estimating symmetry and reflecting points notably makes up for otherwise missing correspondences. Surprisingly, the approach shows benefits even in the case of only partially symmetric objects. Furthermore, the approach is tuned such that both large and small overlap cases can be transparently handled. While branch-and-bound implementations are not real-time capable, we are still able to demonstrate practical usefulness, for example, in the context of back-end shape completion.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Alexandrov, S.V., Patten, T., Vincze, M.: Leveraging symmetries to improve object detection and pose estimation from range data. In: International Conference on Computer Vision Systems, pp. 397–407. Springer, Berlin (2019)
- Arun, K.S., Huang, T.S., Blostein, S.D.: Least-squares fitting of two 3-d point sets. *IEEE Trans. Pattern Anal. Mach. Intell.* **5**, 698–700 (1987)
- Avetisyan, A., Dai, A., Nießner, M.: End-to-end cad model retrieval and 9d of alignment in 3d scans. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2551–2560 (2019)
- Bergström, P., Edlund, O.: Robust registration of point sets using iteratively reweighted least squares. *Comput. Optim. Appl.* **58**(3), 543–561 (2014)
- Besl, P.J., McKay, N.D.: Method for registration of 3-d shapes. In: Sensor fusion IV: control paradigms and data structures. International Society for Optics and Photonics, vol. 1611, pp. 586–606 (1992)
- Brégier, R., Devernay, F., Leyrit, L., Crowley, J.L.: Symmetry aware evaluation of 3d object detection and pose estimation in scenes of many parts in bulk. In: Proceedings of the IEEE International Conference on Computer Vision Workshops, pp. 2209–2218 (2017)
- Bustos, Á.P., Chin, T.-J.: Guaranteed outlier removal for point cloud registration with correspondences. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**(12), 2868–2882 (2017)
- Campbell, D., Petersson, L.: Gogma: Globally-optimal gaussian mixture alignment. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5685–5694 (2016)
- Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H. et al.: Shapenet: an information-rich 3d model repository. *arXiv preprint arXiv:1512.03012* (2015)
- Chen, C.-S., Hung, Y.-P., Cheng, J.-B.: Ransac-based darces: a new approach to fast automatic registration of partially overlapping range images. *IEEE Trans. Pattern Anal. Mach. Intell.* **21**(11), 1229–1234 (1999)
- Chen, J., Wu, X., Wang, M.Y., Li, X.: 3d shape modeling using a self-developed hand-held 3d laser scanner and an efficient ht-icp point cloud registration algorithm. *Opt. Laser Technol.* **45**, 414–423 (2013)
- Chen, Y., Medioni, G.: Object modelling by registration of multiple range images. *Image Vis. Comput.* **10**(3), 145–155 (1992)
- Chetverikov, D., Stepanov, D., Krsek, P.: Robust euclidean alignment of 3d point sets: the trimmed iterative closest point algorithm. *Image Vis. Comput.* **23**(3), 299–309 (2005)
- Cicconet, M., Hildebrand, D.G., Elliott, H.: Finding mirror symmetry via registration and optimal symmetric pairwise assignment of curves: algorithm and results. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1759–1763 (2017)
- Cohen, A., Zach, C., Sinha, S.N., Pollefeys, M.: Discovering and exploiting 3d symmetries in structure from motion. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1514–1521. IEEE (2012)
- Enqvist, O., Josephson, K., Kahl, F.: Optimal correspondences from pairwise constraints. In: 2009 IEEE 12th International Conference on Computer Vision, pp. 1295–1302. IEEE (2009)
- Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* **24**(6), 381–395 (1981)
- Gelfand, N., Mitra, N.J., Guibas, L.J., Pottmann, H.: Robust global registration. In: Symposium on Geometry Processing, vol. 2, p. 5. Vienna, Austria (2005)
- Granger, S., Pennec, X.: Multi-scale em-icp: A fast and robust approach for surface registration. In: European Conference on Computer Vision, pp. 418–432. Springer (2002)
- Guo, Y., Bennamoun, M., Sohel, F., Lu, M., Wan, J., Kwok, N.M.: A comprehensive performance evaluation of 3d local feature descriptors. *Int. J. Comput. Vis.* **116**(1), 66–89 (2016)
- Hartley, R.I., Kahl, F.: Global optimization through rotation space search. *Int. J. Comput. Vis.* **82**(1), 64–79 (2009)
- Jiang, W., Xu, K., Cheng, Z.-Q., Zhang, H.: Skeleton-based intrinsic symmetry detection on point clouds. *Graph. Models* **75**(4), 177–188 (2013)
- Liu, Y., Hel-Or, H., Kaplan, C.S., Van Gool, L., et al. Computational symmetry in computer vision and computer graphics. *Found. Trends Comput. Graph. Vis.* **5**(1–2), 1–195 (2010)
- McCormac, J., Clark, R., Bloesch, M., Davison, A., Leutenegger, S.: Fusion++: Volumetric object-level slam. In: 2018 International Conference on 3D Vision (3DV), pp. 32–41. IEEE (2018)
- Mian, A.S., Bennamoun, M., Owens, R.: Three-dimensional model-based object recognition and segmentation in cluttered scenes. *IEEE Trans. Pattern Anal. Mach. Intell.* **28**(10), 1584–1601 (2006)

26. Milanese, M.: Estimation and prediction in the presence of unknown but bounded uncertainty: a survey. In: *Robustness in Identification and Control*, pp. 3–24. Springer (1989)
27. Mitra, N.J., Pauly, M., Wand, M., Ceylan, D.: Symmetry in 3d geometry: extraction and applications. In: *Computer Graphics Forum*, vol. 32, pp. 1–23. Wiley Online Library (2013)
28. Nagar, R., Raman, S.: Fast and accurate intrinsic symmetry detection. In: *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 417–434 (2018)
29. Newcombe, R.A., Izadi, S., Hilliges, O., Molyneaux, D., Kim, D., Davison, A.J., Kohli, P., Shotton, J., Hodges, S., Fitzgibbon, A.: *KinectFusion: Real-time dense surface mapping and tracking*. In: *International Symposium on Mixed and Augmented Reality (2011)*
30. Pulli, K.: Multiview registration for large data sets. In: *Second International Conference on 3-D Digital Imaging and Modeling (Cat. No. PR00062)*, pp. 160–168. IEEE (1999)
31. Qi, C.R., Liu, W., Wu, C., Su, H., Guibas, L.J.: Frustum pointnets for 3d object detection from rgb-d data. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 918–927 (2018)
32. Rock, J., Gupta, T., Thorsen, J., Gwak, J., Shin, D., Hoiem, D.: Completing 3d object shape from one depth image. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2484–2493 (2015)
33. Rusu, R.B., Blodow, N., Beetz, M.: Fast point feature histograms (fpfh) for 3d registration. In: *2009 IEEE International Conference on Robotics and Automation*, pp. 3212–3217. IEEE (2009)
34. Salas-Moreno, R.F., Newcombe, R.A., Strasdat, H., Kelly, P.H., Davison, A.J.: Slam++: Simultaneous localisation and mapping at the level of objects. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1352–1359 (2013)
35. Schiebener, D., Schmidt, A., Vahrenkamp, N., Asfour, T.: Heuristic 3d object shape completion based on symmetry and scene context. In: *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 74–81. IEEE (2016)
36. Speciale, P., Oswald, M.R., Cohen, A., Pollefeys, M.: A symmetry prior for convex variational 3d reconstruction. In: *European Conference on Computer Vision*, pp. 313–328. Springer (2016)
37. Tsin, Y., Kanade, T.: A correlation-based approach to robust point set registration. In: *European Conference on Computer Vision*, pp. 558–569. Springer (2004)
38. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: a deep representation for volumetric shapes. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1912–1920 (2015)
39. Xu, K., Zhang, H., Tagliasacchi, A., Liu, L., Li, G., Meng, M., Xiong, Y.: Partial intrinsic reflectional symmetry of 3d shapes. *ACM Trans. Graph. (TOG)* **28**(5), 138 (2009)
40. Yang, H., Antonante, P., Tzoumas, V., Carlone, L.: Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection. *IEEE Robot. Autom. Lett.* **5**(2), 1127–1134 (2020)
41. Yang, H., Shi, J., Carlone, L.: Teaser: fast and certifiable point cloud registration. *arXiv preprint arXiv:2001.07715* (2020)
42. Yang, J., Li, H., Campbell, D., Jia, Y.: Go-icp: A globally optimal solution to 3d icp point-set registration. *IEEE Trans. Pattern Anal. Mach. Intell.* **38**(11), 2241–2254 (2015)
43. Yang, J., Li, H., Jia, Y.: Go-icp: Solving 3d registration efficiently and globally optimally. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1457–1464 (2013)
44. Yang, S., Scherer, S.: Cubeslam: monocular 3-d object slam. *IEEE Trans. Robot.* **35**(4), 925–938 (2019)
45. Yew, Z.J., Lee, G.H.: Rpm-net: Robust point matching using learned features. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11824–11833 (2020)
46. Yoshiyasu, Y., Yoshida, E., Guibas, L.: Symmetry aware embedding for shape correspondence. *Comput. Graph.* **60**, 9–22 (2016)
47. Zhang, Z.: Iterative point matching for registration of free-form curves and surfaces. *Int. J. Comput. Vis.* **13**(2), 119–152 (1994)
48. Zhou, Q.-Y., Park, J., Koltun, V.: Fast global registration. In: *European Conference on Computer Vision*, pp. 766–782. Springer (2016)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Lan Hu** received the B.S. degree from Shanghai University, Shanghai, China, in 2016. She is currently pursuing the Ph.D. degree with the Shanghai Institute of Microsystem and Information Technology, Chinese Academy of Sciences, and ShanghaiTech University, Shanghai, China. Her research interests include 3-D vision, semantic SLAM, and robotics.



**Laurent Kneip** holds a Dipl.-Ing. degree from the Friedrich-Alexander University Erlangen/Nürnberg and a PhD degree from ETH Zurich, where he worked at the Autonomous Systems Lab. He also is a former recipient of the Discovery Early Career Researcher Award (DECRA) from the Australian Research Council, hosted by the Australian National University, and an associate investigator of the Australian centre of excellence for robotic vision. Dr. Kneip currently holds a position as an Associate Professor within the School of Information Science and Technology at ShanghaiTech, where he founded and directs the Mobile Perception Laboratory (MPL). Dr Kneip has published more than 70 papers mostly in top conferences and journals such as ICCV/CVPR/ECCV/ICRA/TPAMI/TRO. He is the main author of the open-source project OpenGV, and his long-term contributions in geometric computer vision have, for example, been awarded by the Marr Prize (honourable mention) at ICCV in 2017.